

GUIDED TRANSFORMATION JOURNEY

AI-Native Transformation Roadmap

Operationalizing AI

Corvus Mechanical

Sample document. Corvus Mechanical, its people, and its numbers are fictional, composed to illustrate a Tier8 Guided Transformation Journey deliverable. No real company is referenced or depicted.

Prepared for Corvus Mechanical's leadership team, including Dana Whitfield, COO, from Corvus Mechanical's Guided Transformation Journey | July 2026

PATH: OPERATIONAL TRANSFORMATION

TRACK: EXPLORING

LIVING DOCUMENT, V5

CONTENTS

Table of Contents

1. **Executive Summary:** the three findings driving this plan, and the first move
 2. **The Transformation Path:** what was recommended, what wasn't, and why
 3. **Where Corvus Is Headed: 2029:** Dana's vision, played back
 4. **Where Corvus Is Today: The Baseline:** the four workflows this plan fixes, and what each costs today
 5. **What Corvus Has to Work With: Capability Check:** the readiness read behind why this plan sequences the way it does
 6. **The Three Redesigns That Matter Most:** what changes, why, who owns it, and what each one sets up
 7. **The Sequence:** three horizons, with owners and measurable gains
 8. **Conclusion:** what this plan asks for, and what it sets up
 9. **Assumption Ledger:** every number that can be corrected
 10. **A Living Document:** how amendments work
- Appendix A. Redesign Economics:** the cost and payback numbers behind each redesign
-

ORIENTATION

How to Read This Roadmap

This roadmap comes out of a Guided Transformation Journey: a set of structured working sessions, not a survey. Those sessions mapped how Corvus Mechanical runs today against where Dana says the business is heading by 2029. Every finding, number, and quote below comes from that record. It is a planning tool built from that record, not a guarantee, and whether it fits Corvus's business is a call only Corvus can make.

The journey routes every company onto one of three paths, and the evidence from the sessions decides which path, not preference. Operational Transformation fits when the constraint on growth sits inside how the business runs today. Strategic Expansion fits when operations are already solid and the real ceiling is market position. Dual is reserved for organizations with the sponsorship and the bandwidth to run both paths at once. Corvus's own evidence points to Operational Transformation. Section 2 lays out that evidence in full, and what it was weighed against.

THE OPTION THIS PLAN DOES NOT RECOMMEND

One more option deserves naming here, because many companies reach for it by default, and this plan does not recommend it to Corvus: bolting AI tools onto an operation that isn't ready to use them well. It looks like progress. It produces far less. It just adds one more disconnected tool that nobody owns, on top of the informal AI use the record already shows.

Every recommendation below moves toward AI-native architecture instead: tools built into a connected, governed way of working, because that is what actually compounds over time. Section 6 names, redesign by redesign, the specific readiness gap that rules out bolting AI on now for Corvus's operation. None of the three redesigns can safely carry AI yet, based on Section 5's readiness read. Forcing one in now would add an ungoverned tool to a company that, by its own account, has never assigned anyone to own how AI gets used.

A note on how this is written. Corvus's leadership, not Dana alone, is the intended reader: the CEO and board pushing this effort forward will read this too and act on it alongside her. Where a finding, a number, or a quote is specifically Dana's own, reflecting the vantage point she has as the person who lived these sessions, she is named directly instead of folded into a generic "you." The sessions themselves had one respondent, Dana as COO. Section 5 recommends bringing the fuller leadership team into a team assessment before the 90-day mark, and Section 10 explains how to bring this document to them for review.

Last, a word on the title. This is not a claim that Corvus is AI-native today. Section 5's readiness read says plainly it isn't yet: nothing AI-related at Corvus is owned, governed, or trained right now. This plan does not lead with AI tools before the operation is ready to use them well. It is the roadmap for getting there.

Each of the three redesigns in Section 6 does double duty. It fixes something broken today, and it builds a piece of the foundation an AI-native version of that same function will eventually run on. That connection is named directly, redesign by redesign in Section 6 and horizon by horizon in Section 7. The reader should not have to infer it.

SAMPLE

SECTION 01

Executive Summary

Corvus's evidence puts it on the Operational Transformation path. When Corvus's operation was tested against doubled demand, everything that broke was inside its own walls. The two national accounts Corvus lost were lost on coverage, consistency, and proposal speed, with operating data named as the deciding difference.

Three findings drive everything in this plan:

FINDING 01**Corvus's operation runs on knowledge that lives in people instead of systems, and that knowledge has a retirement clock on it.**

Nine coordinators dispatch 140 technicians from memory. Three estimators rebuild every quote from scratch. Twelve senior technicians carry the diagnostic knowledge of the company, and five of them retire within four years. This is the single thread connecting Corvus's dispatch capacity limit, its 6 to 12 day quotes, and its three-year technician ramp.

FINDING 02**Corvus's costliest losses are loops, not steps.**

Roughly 8% of jobs need a second truck roll because the right part wasn't aboard, an estimated 170 repeat visits every week. Work orders loop backward at billing when labor, parts, and notes don't match, feeding a 9-day invoice lag and 61-day days sales outstanding. Quotes loop through missing-information chases that consume most of their 6 to 12 day calendar.

Sized in first-pass terms, and classified the way the CFO will read them: the windshield loop, the drive time between jobs, ties up an estimated 1,050 technician-hours a week. If Horizon 1's instrumentation confirms that estimate, and one third of it proves recoverable, that is capacity released equivalent to roughly 9 technicians without hiring, time whose loaded cost runs \$1.1 to \$1.5 million a year on ledgered rate assumptions. That figure is capacity, not cash. It becomes money only through a conversion decision this plan names in Section 6, and no return is claimed until that decision is made.

The second-roll loop is sized separately, roughly 340 technician-hours a week on its own assumptions, and deliberately never added to the windshield figure. Repeat-visit travel is travel, and an unknown share of it already sits inside the windshield estimate. Horizon 1's measurement design separates the two. These are first-pass figures on labeled assumptions, built to be corrected, but the order of magnitude is why this plan starts where it does.

FINDING 03**Corvus's readiness to fix this is early-stage, and the plan respects that.**

Nothing AI-related at Corvus is owned, governed, or trained today, and its strongest people have the least slack. So this roadmap starts with ownership, measurement, and two contained pilots, not with systems. One readiness gap is now load-bearing rather than descriptive: because both the dispatch pilot and knowledge capture would put customer and building data into new tooling, closing the data-governance gap is a precondition to both. It is sequenced in Horizon 1, before capture or AI processing begins. Dana told us what a two-front, big-bang effort cost Corvus in 2023. This plan is built to be the opposite of that.

THE FIRST MOVE

Name the transformation owner and start the five measurements in Section 7. **Owner: Dana and the CEO, within 30 days.** Everything else sequences behind it.

Each of the three redesigns below also does double duty. It fixes something broken today, and it builds a piece of the foundation Corvus's AI-native destination will eventually need: structured dispatch data, a structured quote library, and a structured diagnostic library. Section 6 names that connection explicitly for each one.

SECTION 02

The Transformation Path

Recommended and confirmed: Operational Transformation. The constraint on Corvus's growth is delivery capacity and speed, not market demand.

Here is the evidence. Doubled demand breaks dispatch first: coordinator knowledge that lives in people rather than systems becomes a hard capacity limit, per Dana's account. Parts and senior-technician escalations break next. The national accounts were lost on coverage, consistency, and proposal speed: the competing bid was backed by operating data where Corvus offered experience, per Dana.

Deliberately not recommended now: Strategic Expansion. Corvus's uptime-priced predictive maintenance ambition is real, and it stays in this plan, but it runs on exactly the data the operation doesn't capture yet. Chasing it in parallel would split one sponsor, with no named owner and zero slack, across two fronts, against a fresh memory of what that costs.

REASSESSMENT TRIGGER

When structured service history exists for Corvus's top accounts and dispatch data is trustworthy, expected within 9 to 18 months on this plan, the predictive maintenance opportunity gets formally re-evaluated. It is the prize for doing operations first, not the casualty of it.

SECTION 03

Where Corvus Is Headed: 2029

Dana's vision, played back.

Provenance note: this section was rebuilt from the session record rather than quoted from a live transcript. Confirm or correct any of it at the first amendment, and it regenerates.

Routine dispatch runs itself. Corvus's people own the exceptions and the customer commitments. Three coordinators handle judgment calls and relationships instead of nine doing manual matching, and every technician arrives with the job's history, likely parts, and the building's quirks in hand. Quotes go out in 48 hours because estimators judge price and risk instead of assembling from scratch, and Corvus wins the speed deals again. The approval threshold sits at \$150,000 because the numbers underneath quotes have earned that trust, and the approval queue is no longer where deals wait. The knowledge of Corvus's retiring senior technicians lives in a system new technicians actually use, and ramp time drops from three years to one. Leadership sees daily margin by branch and by line of business without asking. And Corvus sells what it never could before: maintenance contracts priced on uptime, because the data finally exists to price them.

That is the human-ownership structure this plan protects. Dispatch runs itself for routine calls, but Corvus's coordinators still own every exception and every customer relationship. Quotes assemble themselves for known configurations, but Corvus's estimators still own the judgment calls on price and risk. Nothing in this plan asks Corvus's people to disappear from the business. It asks the routine work to stop needing all of their time.

What the customer notices: a facility manager gets a faster response from a technician who walks in already knowing the building, and renewal conversations are backed by operating data instead of relationship assurances alone.

SECTION 04

Where Corvus Is Today: The Baseline

Four workflows sit underneath everything in this plan. For each one, this section names the loop cost: what today's broken pattern actually costs Corvus by repeating, uncorrected.

Service dispatch

A facility manager's call becomes a work order, then a coordinator's judgment call. Which of 140 technicians can really handle this job? That call is made from memory and old messages, in a scheduling system everyone works around. The handoff from intake to committed assignment is the longest sit in the chain. Technicians average 3.1 jobs a day and an estimated 90 minutes of daily windshield time, the drive time between jobs. That estimate comes from Corvus's reconstructed vision record and has never been instrumented. About 8% of jobs need a second truck roll because the right part wasn't aboard, by Dana's own estimate. Volume runs roughly 2,100 to 2,200 jobs per five-day week across six branches.

THE LOOP COST

At 140 technicians and roughly 90 minutes each, windshield time alone runs on the order of 1,050 technician-hours per week. That is arithmetic on an estimate, not a measured figure, so it needs instrumenting before anyone treats it as precise. Second rolls add an estimated 170 repeat visits weekly (8% of roughly 2,150 jobs). These two figures overlap by construction: repeat-visit travel is travel, and an unknown share of it already sits inside the windshield estimate. They are sized separately in this document and never summed. Horizon 1's measurement design separates routine travel from repeat-visit travel, so each loop gets its own honest number. Cycle time by call type is not measured at all today, and that absence is itself a finding.

Install quoting

A quote begins as detective work. The request package arrives incomplete, and the estimator chases sales, field staff, vendors, and the customer to fill the gaps. The traced instance took 9 business days, inside Dana's stated range of 6 to 12. In her words: "the actual judgment work was a fraction of the total time; queueing and information retrieval ate most of the calendar." Everything over \$50,000 then waits on Dana or the CFO personally, a queue Dana calls "a running joke internally." Corvus loses roughly 1 in 4 deals purely on response speed. Volume runs to dozens per month, a figure Dana called "not defensible" pending a data pull. The share requiring personal approval is unknown.

THE LOOP COST

Not computable in dollars from the record as it stands.

What is known: this is a structural, recurring delay across dozens of quotes a month, and no system, only the personal attention of Dana and the CFO, catches which ones need it. Measuring actual quote volume and the true share above \$50,000 is this plan's first-90-day action.

Field diagnostic knowledge

A second-year technician hits an unfamiliar failure and calls a senior technician, who answers from memory between jobs. The exchange is valuable, and, as Dana put it, "it disappears the moment the call ends." Twelve senior technicians carry this knowledge, and five of them retire within four years. Ramp time to independence runs three years. Escalation call volume has never been sized, though Dana estimates it at tens of calls a week.

THE LOOP COST

Not computable in dollars from the record as it stands.

What is known: the retirement clock (5 of 12 in four years) and the total absence of any written record of what those twelve people know. Escalation call volume is this plan's second first-90-day measurement.

Billing and margin visibility

Work orders loop backward at every mismatch: labor, parts, purchase orders, notes, or signoff. Invoices trail job completion by 9 days on average, and days sales outstanding runs 61 days, both per Corvus's own billing figures. Margin review, as Dana describes it, "argues about coding before discussing the business," and the number leadership trusts is weeks old by the time anyone sees it.

THE LOOP COST

Not computable in dollars from the record as it stands.

What is known: the structural pattern of backward loops at every mismatch, and a 9-day gap between finishing a job and billing for it that compounds directly into the 61-day collection cycle.

The pattern across all four

Capable people are compensating for missing structure, at real personal cost. This is not praise. In Dana's words: "it is capable people compensating for broken flow, and it is exhausting them." The people carrying the workarounds see the brittleness clearly. The workarounds themselves insulate leadership

from seeing it just as clearly. And in two places, a number this plan needs does not exist anywhere in Corvus's business today: dispatch cycle time by call type, and escalation call volume. Those absences are findings in their own right, and closing them costs measurement discipline, not new systems.

SAMPLE

SECTION 05

What Corvus Has to Work With: Capability Check

This section reads Corvus's readiness across five pillars. Those pillars determine how safely, and how far, AI can be built into any of the redesigns below. The five pillars are: who's accountable for how AI tools get used, and how well the organization absorbs real change. Also: how embedded AI already is in daily work, what controls exist over what goes into these tools, and how prepared Corvus's people are to use them well.

Each pillar lands in one of four bands, from earliest to most mature: Initial, Developing, Advancing, or Leading. The band reflects what's actually in place today, not effort or intent. The evidence column cites exactly what earned that score.

AREA	WHERE CORVUS STANDS	THE EVIDENCE
Responsible Use & Accountability	INITIAL	AI ownership "never been assigned, formally or informally." One data-rule email went out and was forgotten. There is no AI-specific process to catch a wrong output, though Corvus's normal review layers are real.
Adoption & Change	INITIAL	Concrete urgency and a confirmed path exist, but only one true sponsor: Dana. The CEO backs the effort but delegates it, the CFO is skeptical, and the 2023 ERP scar is still fresh.
Integration into Work	INITIAL	No AI is embedded in any real workflow today. Corvus's operational data lives in truck cabs, in people's heads, and in a scheduling system everyone works around.
Security & Risk	INITIAL	Role-based access and trusted IT exist, real capability to build on. But Corvus has zero visibility into shadow AI use, an honest "I don't know" from Dana, and its operational-risk instinct is sharp: Dana named dispatch as the risk immediately, and for the right reason.
Enablement & Skills	INITIAL	No training exists, and no tools have been formally provided. Corvus's best people are also its busiest, which leaves little room to learn something new on top of what they already carry.

Read this the way we do: early is not unready. Three things in this record don't show up in the scores.

First, Corvus's review discipline and access controls are real capability to build on. Second, a handful of Corvus's people adopted AI tools with no mandate, which, in Dana's own calibration, "tells me they see a use for it." Third, and largest, the knowledge in Corvus's coordinators' and senior technicians' heads is not just a risk. It is codifiable intellectual property no competitor holds, and this plan treats capturing it as building an asset, not a rescue.

ON THE CRITICAL PATH

One row of this table is now on the critical path. Redesigns 1 and 3 in Section 6 would each put customer, building, and equipment data into new tooling, and this record shows no AI-specific data controls today. Closing that gap is a signed precondition to any capture or AI processing in either redesign. The capability read above is load-bearing, not just descriptive.

Caveat, stated plainly: this capability read reflects one leader's view, Dana's. Several answers were honest unknowns. The team assessment included with this journey, up to six leaders, would replace those unknowns with data, and would surface where Corvus's leadership team sees the same questions differently. Recommended before the 90-day mark.

SECTION 06

The Three Redesigns That Matter Most

Ranked by what today's loops cost, filtered by what Corvus can actually start within 90 days.

The full cost, payback, and assumption numbers behind each redesign are in Appendix A. This section covers what changes, why this one, and what each one sets up.

REDESIGN 1

Dispatch and the job-context pack

What changes: every completed job leaves behind a structured context pack: what was found, what was done, parts used, the building's quirks, captured at work-order close. Corvus's coordinators' matching rules get codified from their heads into explicit, written criteria. Then, in one pilot branch, routing assistance works from that data, with coordinators still handling every exception.

Why this one first: it attacks the largest estimated bleed, windshield hours, instrumented in Horizon 1 before any investment goes past pilot prep. It also attacks the separately estimated second-roll rate, and the hardest capacity ceiling: coordinator tribal knowledge, by Dana's own doubled-demand answer. It generates the structured operating data every later move needs, including quoting context, parts prediction, and eventually uptime contracts. This is also where Dana said AI risk concentrates, which is why the pilot stays one branch with human exception handling, not a company-wide switch.

Setting up AI-native: The destination mirrors what Dana described: routine dispatch runs itself, Corvus's coordinators own exceptions and relationships, and every technician arrives with the job's history, likely parts, and building quirks already in hand.

What blocks it today, per Section 5, is Integration into Work (Corvus's dispatch data lives in truck cabs and coordinators' heads, not a system) and Responsible Use & Accountability (nobody yet owns how AI would touch a dispatch decision). This redesign is the action that closes the Integration gap. Once every job leaves behind a structured context pack instead of nothing, that data becomes the trustworthy input a routing-assistance tool would eventually need. Nothing here builds that tool company-wide yet. Section 5's governance gap has to close first, which is why routing assistance stays a one-branch pilot, gated behind a signed data-preconditions package (Horizon 1, action 7).

Owner: the operations leader for the pilot branch, with the transformation owner (Section 7) accountable overall.

Measurable gains: second-roll rate in the pilot branch, baseline measured in Horizon 1, then reduced against a target set from that baseline. Windshield minutes per technician per day, baseline measured, then reduced. Coordinator hours spent on routine matching, sampled before and after.

Exposure: what Corvus risks is the pilot's lifecycle cost, one branch's coordinator attention, and a bounded operational risk. Routing suggestions can influence a dispatch decision even with coordinators deciding, and a bad assisted dispatch in the pilot branch is a real customer-facing miss. That risk is bounded by one-branch scope, coordinator final say on every assignment, a defined exception protocol, and a rollback that is one decision, not a project. Worst case: the routing assist is switched off, and the context pack remains, which has standalone value on its own.

Conditions to be true before routing assistance goes live: the data preconditions package signed (vendor review, retention rules, access scope), since the context pack and routing tooling put customer, building, and equipment data into new tooling. Context pack captured on more than 80% of pilot-branch jobs. Matching rules documented and coordinator-endorsed. An exception protocol defined. A rollback that is one decision.

WHAT NOT TO DO

Do not buy a routing system first and reshape dispatch around it. That is the 2023 shape. Data and rules come first, tooling after, one branch at a time.

REDESIGN 2

Quoting speed

What changes: a quote library, built from Corvus's historical quotes, turns repeat configurations into assembly instead of authorship. The request package gets a completeness standard, what sales must include before an estimator touches it, killing the detective-work loop at its source. Vendor pricing gets cached or put under response agreements. The approval threshold rises in steps, \$50,000 to \$75,000 to \$100,000, as quote-data trust builds, with Dana and the CFO agreeing the evidence bar for each step in advance.

Why: this is Corvus's revenue-speed loop. A quarter of lost deals die on response time, and Dana's own read of those losses is that the faster competitor defines the job first. It is also the CFO's shortest path from skeptic to sponsor: the threshold steps only happen when the numbers earn them, which makes trust measurable instead of assumed.

Setting up AI-native: The destination is Corvus's own 2029 picture almost word for word: quotes out in 48 hours, because estimators judge price and risk instead of assembling every quote from scratch, and Corvus wins the speed deals again.

What blocks it today, per Section 5, is the same Integration gap: no structured quote data exists, so every quote gets rebuilt from scratch by three estimators. This redesign is the action that builds that structure. Once repeat configurations move from authorship to assembly, that library becomes the trustworthy input an eventually assisted quoting tool would need. Nothing here deploys AI-assisted quoting yet, and nothing lets the approval threshold jump to \$150,000 on ambition. Each step moves only once a signed, numerically defined evidence bar clears, because trust in the underlying data has to be earned, not assumed, and Section 5 shows nobody yet owns how AI would touch a pricing decision.

Owner: Corvus's estimating and sales leadership. The CFO co-owns the threshold schedule.

Measurable gains: quote cycle reduced against the Horizon 1 baseline, with 48 hours as the 2029 standard for library-covered configurations. Win and loss rate on speed, measured from a log that starts in Horizon 1. Share of quotes needing personal approval, measured then reduced.

THE EVIDENCE BAR FOR THE FIRST THRESHOLD STEP (\$50,000 TO \$75,000)

Proposed here for the CFO to finalize and sign before anything moves: at least 20 completed, library-covered quotes in or near the band. Estimator-priced margin within 2 points of actual on those jobs. A pricing-error exception rate under 5%. Zero deals lost to a pricing error. The CFO keeps spot-audit rights throughout.

Bar feasibility, stated plainly: at dozens of quotes a month across all sizes and configurations, reaching 20 completed cases in the right band realistically takes roughly 4 to 10 months, not a rolling 90 days. The bar's window is the accumulation period itself: it closes when the 20th qualifying job lands. If any condition breaks during that stretch, the threshold stays put and the count resets.

Later steps, numbered now so nothing is signed blind. Each is finalized and signed by the CFO before the step before it activates, never negotiated mid-flight. \$75,000 to \$100,000: at least 15 completed, library-covered quotes in or near the band (volume thins at higher values), margin within 2 points, exception rate under 4%, zero pricing-error losses, expected time-to-testable roughly 8 to 14 months from step activation. \$100,000 to \$150,000: at least 12 completed in-band quotes, margin within 1.5 points, exception rate under 3%, zero pricing-error losses, expected time-to-testable 12 to 24 months, which likely places this step in Horizon 3 or beyond. This plan says so rather than pretending otherwise. Every number here is a proposal for the CFO to tighten or reset. A baseline makes a bar checkable; only the bar itself moves the control.

Exposure: cash exposure is limited to estimator and admin time. Two named risks remain even behind human review. An assembled quote can carry a wrong assumption to a customer, bounded by estimator review of every quote and the CFO's spot-audit rights. And each threshold step, once active, moves approval errors in that band from Dana's desk to estimator judgment, bounded by the signed evidence bar, the exception-rate ceiling, and single-decision reversion.

Conditions to be true: quote library seeded from the last 18 to 24 months of quotes. Completeness standard enforced at intake. A vendor pricing agreement or cache for the top vendors by quote frequency.

WHAT NOT TO DO

Do not let the threshold jump to \$150,000 on ambition. Corvus's 2029 number is earned in steps, or the trust it represents is fiction.

REDESIGN 3

Diagnostic knowledge capture

What changes: the escalation call stops evaporating. Calls from junior technicians to senior technicians get logged, and, once the data preconditions below are signed, the exchange gets captured at the moment it happens: recorded and transcribed into a searchable, growing library organized by symptom, equipment, and building. Senior technicians review and correct entries rather than authoring documentation from scratch. New-technician onboarding starts drawing from the library.

Why: five of the twelve heads holding Corvus's diagnostic capability retire within four years, ramp time is three years, and this is the risk Dana ranked highest in her own account. It compounds: the library feeds the dispatch context pack, shortens escalations, and is the long-lead ingredient of the uptime-contract ambition.

Setting up AI-native: The destination matches what Dana described earlier: the knowledge of Corvus's retiring senior technicians lives in a system new technicians actually use, and ramp time drops from three years to one.

What blocks it today, per Section 5, is Security & Risk (zero AI-specific data controls exist), and this move puts real customer, building, and equipment information (from field calls) into new tooling for the first time. This redesign is the action that closes the gap, but deliberately in stages. It starts with a call log in Horizon 1 that captures counts and topics only, no call content at all. It does not begin recording or transcribing anything until the data-preconditions package is signed:

vendor review, retention rules, access scope, and a written rule for what may and may not be recorded. That gate exists because Security & Risk has to close before any content capture, not alongside it.

DATA PRECONDITIONS, BEFORE ANY CAPTURE BEGINS. THIS IS A GATE, NOT A PARALLEL TASK.

Recording and transcribing field calls puts customer, building, and equipment information into new tooling, and Corvus's readiness record shows no AI-specific data controls today. Four things have to be in place before the first call is recorded. A vendor review: where the data lives, who can access it, whether it trains anyone else's models. Retention rules, and access scope inside Corvus. An explicit written rule for what may and may not be recorded. Technicians whose calls are captured acknowledge that rule directly. This package is owned by the transformation owner with IT, and capture in Horizon 2 starts only after it is signed. The call log itself, counts, durations, topics, no content, starts in Horizon 1 regardless, because it captures no call content and it buys the baseline.

Owner: Corvus's field service leadership, with one senior technician as a named champion. The record suggests real enthusiasm here, worth recruiting deliberately.

Measurable gains: escalation calls per week, measured for the first time, baseline in Horizon 1. Junior technician time-to-resolution on escalated faults, sampled. Library entries endorsed by senior technicians. Ramp milestones for the next new-hire cohort against the three-year baseline.

Exposure: tooling and senior hours, plus a real data exposure: field calls contain customer and building information, and putting them into new tooling is a new risk surface. That is why the preconditions gate comes before capture rather than alongside it. The exposure is bounded by the signed vendor review, the recording rule, retention limits, and access scope, not eliminated. Worst case: the library underwhelms, and Corvus still holds its first written diagnostic record before the retirements, inside controls it did not have before.

Conditions to be true: the data preconditions package signed. Capture effortless at the point of the call, no after-the-fact paperwork for seniors, or it will not happen. The library searchable in the field. Seniors' review time scheduled, not squeezed: their slack is zero, so this is a staffing decision, not an ask.

WHAT NOT TO DO

Do not launch this as a documentation program. Corvus's own record shows the condition that rules this out: senior technicians have zero slack, so documentation they must author from scratch will consistently get deprioritized behind the day job. Capture the exchange that already happens; never add homework. And do not start recording before the gate is signed. A shortcut here spends the trust this whole plan runs on.

The backlog (ranked, not now)

Billing and accounts-receivable loop redesign. Much of its mismatch loop dies upstream once work-order close discipline lands in Redesign 1, so revisit at the Horizon 2 close. Daily margin visibility follows billing hygiene. Parts stocking prediction follows 6 to 12 months of context-pack data. Predictive maintenance contracts wait for the Section 2 reassessment trigger. The full \$150,000 approval threshold sits at the end of Redesign 2's staircase.

SAMPLE

SECTION 07

The Sequence

SAMPLE

1

Horizon 1: first 90 days. Own it, measure it, prepare it.

#	ACTION	OWNER	DONE MEANS
1	Name the transformation owner. Corvus's best candidates are its busiest people, which makes this a backfill decision, not a search. Until it's made, this plan has no engine, and Corvus's own record says so.	Dana + CEO	A named person with dedicated hours, by day 30
2	Assign AI accountability, and reissue the data rule as something people can actually find. Cover the external AI tools already in informal use directly, since the record shows proposal content moving through them today with no rule governing it. This is a few hours of drafting, a leadership session, and follow-through to acknowledgment. Sized honestly: the two cheapest governance gaps to close, but not free.	Transformation owner	Named AI owner; rule acknowledged by all office staff, including current informal AI users
3	Start five measurements: dispatch cycle by call type; windshield time, by sampling first, designed to separate routine travel from repeat-visit travel so the two loop sizes stop overlapping (telematics only if sampling proves insufficient); an escalation-call log (counts and topics, no content); quote volume, cycle time, approval share, and in-band completion lag; and win and loss reasons on lost bids.	Each metric assigned to the manager already running that area; transformation owner collects	Baselines in hand by day 60
4	Codify coordinator matching rules and define the context pack. Pick the pilot branch.	Ops leader	Documented rules coordinators endorse
5	Seed the quote library from historical quotes. Draft the completeness standard.	Estimating lead	Library covers Corvus's top repeat configurations
6	Run the team assessment, up to 6 leaders, included in this journey. Replaces this plan's single-respondent unknowns with data before Horizon 2 decisions.	Dana	4 or more leaders completed
7	Data preconditions packages for both data-into-new-tooling moves: (a) knowledge capture, meaning vendor review, retention rules, access scope, and the written recording rule (Redesign 3's gate); (b) context-pack and routing tooling, meaning vendor review, retention, and access scope (Redesign 1's gate).	Transformation owner + IT	Both packages signed. (a) before any call content is captured. (b) before the Horizon 2 routing pilot activates.

Stop doing: personal approvals in the \$50,000 to \$75,000 band, once the staircase's first evidence bar clears (Redesign 2), with the CFO signing the bar before anything moves. What starts now is the measurement that makes the bar checkable. Estimators building repeat-configuration quotes from scratch, once the library covers them. Treating escalation calls as free: at tens of calls a week and a minute or two each, logging from day one costs an estimated one to three field hours a week. That cost is counted, not waved off, and the baseline it buys is what makes Redesign 3's case checkable.

Deliberately absent from Horizon 1: any AI system purchase or deployment. Corvus's readiness record says build ownership and data first, and this plan believes the record.

Where this points: the destination three horizons out is routing, quoting, and diagnostic support that runs on trustworthy data, not on people's memory. The gap blocking it today, per Section 5, is that nobody owns how AI gets used and nothing has been trained: none of Horizon 1 can be AI work yet, on purpose. So the action here is to turn three things that live only in people's heads or scattered spreadsheets into named ownership and structured data instead. That is exactly the raw material Section 6's AI-native destinations for all three redesigns are waiting on.

IF ACTION 1 SLIPS PAST DAY 30

Action 3's escalation log and quote data pull proceed anyway, since they run on existing managers and cost hours, not a project. The CEO is notified of the delay, since owner assignment is the one dependency the rest of the sequence needs to move. Actions 4 through 7 hold until the owner lands. Corvus's record shows a pattern of unassigned gaps getting absorbed by heroics rather than surfaced; this sequence is built to surface this one instead.

2 **Horizon 2: months 3 to 9. Two pilots and a staircase.**

The dispatch pilot goes live in one branch: context pack at close, routing assistance with coordinator exceptions, activated only after its data preconditions package (action 7b) is signed, and instrumented against baselines. The pilot-branch ops leader owns it; the transformation owner stays accountable. The quote library and completeness standard are in force, and the threshold moves to \$75,000 once the first evidence bar clears its accumulation window, likely extending past this horizon's start given the 5 to 10 month time-to-testable. The staircase runs on the bar's clock, not the calendar's. A vendor pricing cache covers the top vendors, owned by estimating, with the CFO co-owning the staircase.

Knowledge capture goes live on escalation calls, but only after the Horizon 1 data preconditions package is signed, with senior review time scheduled. Field service leadership owns it, with the named senior-technician champion. The capacity-conversion decision for any released dispatch hours gets made here, explicitly, by Dana and the ops leader: absorb demand, strengthen coverage, or slow hiring. Until that decision is made, released hours stay classified as capacity, not cash. Training is scoped to the teams

touching these three changes only, not a company-wide program, owned by the transformation owner. This plan reassesses itself against Horizon 1 baselines here: the measurements may re-rank the redesigns, and this plan expects that.

Where this points: the destination is the same one named throughout, routing assistance, assisted quoting, and diagnostic pattern matching, each running on data Corvus can trust. The gap closing this horizon is the missing structure itself: context-pack data, quote-library data, and confirmed escalation-call data, none of which existed before Horizon 1. The action is that by month 9, both pilots run as disciplined, human-owned processes generating exactly that data. That is not yet AI-native architecture. It is the foundation Section 6 says has to exist before any of the three destinations becomes a build worth doing rather than a guess.

3 Horizon 3: months 9 to 18. Scale what earned it.

Ownership carries forward from Horizon 2, with the transformation owner still accountable for the whole. The dispatch model scales branch by branch as pilot metrics hold, each branch's own ops leader owning their rollout, and the coordinator role gets formally redefined toward the 2029 exceptions-and-relationships shape. Forty-eight hour quotes become the standard for library-covered configurations, owned by estimating. The threshold moves to \$100,000, then \$150,000, as each staircase bar clears on its own accumulation window, owned by the CFO, and the approval queue stops being where deals wait. The billing loop gets redesigned using work-order close discipline, with Corvus's billing and accounts-receivable lead drawn into the plan for the first time here. Margin visibility moves from monthly toward daily, owned by the CFO and finance team. The predictive maintenance reassessment, the Section 2 trigger, fires here if the data has accumulated, owned by Dana, and the expansion path gets its honest look with a real foundation under it.

HOW THIS DIFFERS FROM 2023, EXPLICITLY

One workflow at a time, measured against baselines Corvus sets, reversible at each step, no system replacement, and the operation keeps running throughout. The 2023 effort tried to change everything at once from the outside in. This sequence changes one thing at a time from the inside out.

Where this points: the destination is a connected, AI-native version of dispatch, quoting, or diagnostics, whichever Corvus is closest to earning by this point. The remaining gap, if any, is named plainly: Section 5's governance and adoption pillars have to have closed alongside the operational fixes, not automatically because 18 months passed. The action this horizon takes is not to build that AI-native version yet. It is to have the foundation and the honest data in place to make that call for real, and only then.

SECTION 08

Conclusion

Three things are true about Corvus today. Its growth is capped by delivery capacity and speed, not by demand, so this plan builds inside the operation Corvus already has. Its knowledge lives in people instead of systems, and five of its twelve most senior technicians retire within four years with nothing written down. And Corvus's readiness to use AI safely is early-stage: nothing today is owned, governed, or trained, while its strongest people carry the least slack of anyone in the company.

None of the three redesigns in Section 6 asks Corvus to buy a system before honest data and discipline exist. Each one asks for one named owner, one honest measurement, and a contained pilot: one branch, one call log, one staircase step at a time. That match is deliberate. The record shows a company with one real sponsor, a CEO who backs the effort but delegates it, a skeptical CFO, and a fresh memory of what a two-front, big-bang effort costs. This plan does not ask for more than Corvus's operation can currently absorb.

The payoff is not just fixing what is broken today. Each redesign also builds the structured data and the ownership an AI-native version of that same function will eventually need. Dispatch data trustworthy enough to support routing assistance. A quote library structured enough to support assisted pricing. A diagnostic library structured enough to eventually help underwrite the uptime-priced maintenance contracts in Corvus's 2029 picture. Corvus does not have to choose between fixing operations now and moving toward AI-native work later. This plan is built so the first becomes the path to the second.

The ask now is simple. Read this roadmap as a leadership team, not individually, and correct anything that's wrong; it regenerates. Where it holds up, the first 90 days in Section 7 start the clock. This plan is built to be corrected and re-run as circumstances change, not filed away once it's read.

SECTION 09

Assumption Ledger

Every number below can be corrected, and this roadmap regenerates. This page is the front door for corrections.

SAMPLE

#	ASSUMPTION	BASIS	IF WRONG
1	Roughly 90 minutes a day of windshield time per technician, including an unknown share of repeat-visit travel	Corvus's reconstructed vision record; never instrumented	If materially lower, Redesign 1 re-ranks against quoting; instrument in Horizon 1 before investing past pilot prep
2	Roughly 2,100 to 2,200 jobs a week	Corvus's figure	Scales all dispatch arithmetic proportionally
3	Quote volume runs to "dozens" a month, within a factor of 2	Dana's figure, called "not defensible" pending a data pull	Changes quoting gain size, not direction; the Horizon 1 data pull resolves it
4	A meaningful share of quotes are repeat-configuration	Inferred from Corvus's service mix; not confirmed	The library yields less if wrong; the completeness standard carries the redesign either way
5	Escalation calls run "tens per week"	Observed, not measured, by Dana's own flag	Capture still matters for retirement risk regardless; sequencing may shift
6	Senior technicians engage if capture costs them near nothing	Inferred from the record, plus one enthusiasm signal	If wrong, this becomes a leadership ask, not a design fix; revisit the staffing of review time
7	Each second roll costs roughly 2 hours (travel plus setup, no re-diagnosis)	Assumed, for arithmetic; drives the roughly 340 hours a week loop size	A real cost that's higher only strengthens the case for Redesign 1
8	Roughly 1 in 4 lost deals die on response speed	Dana's estimate, never measured	Changes Redesign 2's ranking weight, not its direction; the Horizon 1 win and loss log replaces it with data
9	An 8% second-roll rate	Dana's estimate, and the baseline Horizon 1 measures	If the measured baseline differs, the target resets against it; the direction holds
10	A loaded technician cost of \$65 to \$85 an hour	Assumed; not asked in the sessions	Every loaded-cost figure scales linearly; send Corvus's actual loaded rate and the economics regenerate
11	Roughly one third of windshield time is recoverable	Assumed pending Horizon 1 instrumentation	Capacity figures for Redesign 1 scale with it

#	ASSUMPTION	BASIS	IF WRONG
12	Lifecycle cost ranges: dispatch pilot roughly \$25,000 to \$60,000 to build plus \$10,000 to \$20,000 a year ongoing; knowledge capture roughly \$5,000 to \$15,000 a year in tooling plus senior review hours plus cleanup and preconditions work; quote library 60 to 100 hours to build plus 30 to 50 hours of cleanup plus 4 to 8 hours a month of maintenance	Rough order, no vendor quotes; deliberately includes ongoing ownership, not build alone	Payback windows stretch or shrink; the evidence bars and sequence are unaffected
13	The pilot branch carries roughly 1 in 6 of total job volume	Assumes six roughly equal branches	Pilot payback scales with the branch's actual share
14	Annualization basis: roughly 40 hours a technician per week, 50 weeks a year	A stated convention	All annual figures scale with it
15	The pilot planning target moves second rolls from 8% to 5%	A mechanism basis: the context pack and parts prediction attack the named cause, the right part not being aboard. The committed target is set from the Horizon 1 baseline; 8-to-5 is only the planning figure the payback illustration uses	The payback illustration rescales; the conditional framing in Section 6 already absorbs this
16	An unknown share of repeat-visit travel sits inside the windshield estimate	Structural: both figures involve travel	This is why the two loops are never summed; Horizon 1's sampling design separates them
17	No cash or profit is claimed from released capacity until a named conversion decision is made in Horizon 2	A taxonomy discipline: capacity, cash avoided, and profit captured are different claims	If no conversion decision is ever made, the capacity figures stand, but no financial return does
18	Average install value and install gross margin are unmeasured, named variables	Required by Redesign 2's profit formula	The Horizon 1 data pull turns the formula into a number

#	ASSUMPTION	BASIS	IF WRONG
19	Evidence-bar accumulation: in-band, library-covered completed quotes accumulate at roughly 3 to 8 a month, with a 1 to 3 month lag from quote to completion, putting time-to-testable at roughly 4 to 10 months for the first step. Higher bands accumulate slower: roughly 8 to 14 months for \$100,000, 12 to 24 months for \$150,000	Derived from Corvus's "dozens" a month total volume and typical install timelines; all assumptions, no data	The bars' windows restate from the Horizon 1 data pull; the bars' numbers do not change, only their clocks
20	Escalation logging costs roughly 1 to 2 minutes a call, at tens of calls a week, or 1 to 3 field hours a week total	Sized deliberately, not waved off, given this organization's constraint of zero slack	If the real cost is higher, the logging design simplifies before the ask grows
21	An absorbed technician-hour contributes \$25 to \$45 over loaded cost	Assumed; Corvus's billed rates weren't asked in the sessions; Horizon 1 pulls the actual billed rates by line of business	The Redesign 1 payback range scales linearly; send billed rates and it regenerates
22	The quote-cycle planning target moves 6 to 12 days to 5 days in Horizon 2	A mechanism basis: the traced instance showed judgment work as a fraction of the calendar, with the rest lost to queueing and retrieval, exactly what the library and completeness standard remove. The committed target is set from the Horizon 1 baseline	The planning figure rescales against the measured baseline; the direction holds

SECTION 10

A Living Document

This roadmap regenerates. Send a correction, for example “we forgot the compliance signoff step” or “windshield time measured at 55 minutes,” and an updated version comes back with the change carried through findings, sequencing, and arithmetic. The cover of each version names what changed.

This plan is meant to be reviewed by Corvus’s leadership team directly, not relayed secondhand. Where they disagree with the record, that is not a flaw in the plan. Disagreement is data, and it is exactly what the amendment loop exists to absorb. The team assessment seats included in this journey, up to six leaders, are the fastest way to turn one leader’s honest unknowns into the leadership team’s answer.

Reassessment cadence: this plan expects to be re-run against Corvus’s Horizon 1 baselines at day 90, and against pilot results at month 9. Tier8 holds the full record behind this plan, so a reassessment would not mean starting over from scratch. If it would help to run the day-90 or month-9 check-in together, or to revisit the plan later as circumstances change, Corvus is welcome to reach out. That’s an open door, not something already scheduled or built into this engagement.

APPENDIX A

Redesign Economics

The full cost, payback, and assumption numbers behind each of the three redesigns in Section 6.

HOW TO READ THE ECONOMICS BELOW

Every labor figure is classified as exactly one of three things, because the CFO will read them differently, and should. Capacity released is hours and technician-equivalents: real, but not money until a named decision converts it. Cash cost avoided is claimed only when a specific avoided cost is named. Incremental gross profit is claimed only through a named revenue mechanism, with its variables shown. Where a figure has no honest classification yet, the table says “none claimed” rather than rounding up.

The bracketed numbers after each redesign’s economics heading point to the specific entries in Section 9 that redesign’s math rests on. Check those entries to see exactly what’s assumed, and what changes if it’s wrong.

REDESIGN 1

Dispatch and the job-context pack

First-pass economics [Assumption Ledger items 10, 11, 14, 16, 17, 21]

CLASSIFICATION	FIRST-PASS FIGURE	WHAT CONVERTS IT
Capacity released	Roughly 350 technician-hours a week, about 9 technician-equivalents, if instrumentation confirms the windshield estimate and one third proves recoverable. The loaded cost of that time is \$1.1 to \$1.5 million a year. Second-roll reduction at the planning target releases a further roughly 130 of the roughly 340 second-roll hours a week, sized separately and never summed with the windshield figure (item 16)	A named Horizon 2 decision: absorb routine demand currently scheduled out, strengthen emergency-response coverage, or slow future hiring. No cash or return is claimed until that decision is made (item 17)
Cash cost avoided	Vehicle and fuel costs on avoided repeat visits: real, but small, sized in Horizon 1. No other avoided cost is claimed	Automatic once the second-roll rate falls
Incremental gross profit	Formula only: (released hours absorbed into billable work) times (billed rate minus loaded cost). Billed rates get pulled in Horizon 1; absorption is the Horizon 2 conversion decision	The same conversion decision, plus demand. The record supports demand existing (routine work scheduled out, lost coverage bids), but this plan measures before it promises

PILOT COST, FULL LIFECYCLE [ITEM 12]

One branch, an estimated \$25,000 to \$60,000 to stand up (tooling, context-pack setup, coordinator time), plus an ongoing ownership cost of roughly \$10,000 to \$20,000 a year in maintenance, data-quality review, and rule upkeep. A build-only number would understate this by construction.

FIRST-PASS PAYBACK, STATED CONDITIONALLY AND SHOWN IN FULL

The pilot branch carries roughly 1 in 6 of total job volume, so its share of recoverable windshield hours is roughly 58 hours a week (350 divided by 6). This is windshield hours only; second-roll hours are excluded until Horizon 1 separates the overlap (item 16). If those hours are absorbed into billable work at an assumed contribution of \$25 to \$45 an hour over loaded cost (item 21), they contribute roughly \$73,000 to \$131,000 a year. That is 58 hours times about 50 weeks times \$25 to \$45. The pilot's build cost is \$25,000 to \$60,000, net of its \$10,000 to \$20,000 a year ongoing cost (item 12). Against that cost, payback lands in roughly 3 to 14 months on these planning assumptions. Every one of those assumptions is ledgered and correctable. If the capacity is released but never converted, the pilot's return is capacity, data, and a lower second-roll rate, and this plan does not restate that as cash.

ASSUMPTIONS

[1] the windshield estimate is roughly right, instrumented in Horizon 1; if the real figure is far lower, Redesign 1 re-ranks against quoting. [2] job volume of roughly 2,100 a week holds. [16] the overlap between windshield and second-roll travel stays unresolved until Horizon 1 separates them.

REDESIGN 2

Quoting speed

First-pass economics [Assumption Ledger items 3, 8, 12, 18]

CLASSIFICATION	FIRST-PASS FIGURE	WHAT CONVERTS IT
Capacity released	Estimator hours freed as repeat configurations move from authorship to assembly. Modest, at dozens of quotes a month, and unquantified until the Horizon 1 data pull	More quote throughput from the same three estimators; feeds the profit formula below
Cash cost avoided	None claimed. A specific avoided cost, for example a hire not made in estimating, would be named here if the record supported one; it does not yet	N/A
Incremental gross profit	Formula only: (quotes a month) times (average install value) times (install gross margin) times (share lost on speed) times (recovery fraction). The record holds "dozens" of quotes a month and roughly 1 in 4 speed losses; average install value and install margin are unmeasured, pulled in Horizon 1, at which point this formula becomes a real number	Winning back a share of the speed-lost deals, on the mechanism the record names: response time

COST SIDE, FULL LIFECYCLE [ITEM 12]

An estimated 60 to 100 hours of estimator and admin time to build the library. Add 30 to 50 hours of historical-quote cleanup and validation before the library is trustworthy. Add ongoing maintenance of roughly 4 to 8 hours a month with a named owner. Vendor pricing arrangements cost negotiation attention, not money.

FIRST-PASS PAYBACK

Not computable until the Horizon 1 data pull lands items 3 and 18, and this plan does not invent the missing variables. What can be said now: the cost side above is bounded and small, so recovering even a modest share of speed-lost installs at plausible values clears it quickly. The formula becomes a real number the day the pull lands.

ASSUMPTIONS

[3] quote volume of “dozens” a month is right within a factor of 2, pulled in Horizon 1. [4] a meaningful share of quotes are repeat-configuration; if most are genuinely novel, the library yields less and the completeness standard carries the gain instead. [19] in-band accumulation and completion lag as ledgered.

REDESIGN 3

Diagnostic knowledge capture

First-pass economics [Assumption Ledger items 5, 12, 20]

CLASSIFICATION	FIRST-PASS FIGURE	WHAT CONVERTS IT
Capacity released	None claimed until the escalation log sizes the flow. The Horizon 1 baseline decides whether shortened escalations release measurable senior and junior hours	The Horizon 1 log, then a reduced time-to-resolution
Cash cost avoided	None claimed. A specific avoided cost, for example avoided contract-labor coverage for retiring expertise, would need the record to name it; it does not yet	N/A
Incremental gross profit	None claimed. The later mechanism is named but unquantified: a faster ramp puts new technicians on billable work sooner, and the library underwrites the uptime-contract ambition named in Section 2	Ramp measurement across the next cohort; the Section 2 reassessment trigger

COST SIDE, FULL LIFECYCLE [ITEM 12]

Capture and transcription tooling at roughly \$5,000 to \$15,000 a year. Add 2 to 4 scheduled senior-technician hours weekly for review, plus transcript cleanup and validation time before entries are trustworthy. Add ongoing library curation with a named owner, plus the data-preconditions work itself. The measurable stake this cost sits against is the replacement cost of five retiring careers of diagnostic knowledge and a three-year ramp. This plan does not convert that stake into an invented dollar figure.

ASSUMPTIONS

[5] escalation volume is “tens per week” as observed; if it’s far lower, capture still matters for retirement risk, but sequencing may shift. [6] senior technicians will engage if capture costs them near nothing and credits them visibly.

Prepared by Tier8 | RAISE framework | Built from Corvus’s journey record and labeled where it assumes or reconstructs. Correct anything; it regenerates.