Meridian DealerWorks: what the work says

Assembled from audits run in 2025 and 2026. Company, people, and figures are invented and blended so no client is identifiable. The format and depth reflect what a CEO receives on day 14 of the two-week read.

The one page

The company. Vertical software and services for franchise auto dealer groups. $68M revenue, $61M ARR, about 310 FTE. Sponsor-owned, year four, four acquisitions in the hold. The sponsor's stated gate is to get EBITDA over $32M and run a process. Today the company has $24M EBITDA, $19M free cash flow, and 5.6x net leverage. The CEO is eight months in.

The constraint. The decisions sit with the CEO. The capacity sits inside functions and legacy work. Three of the four moves stall between them. Fix that once and three verdicts move. **STRONG SIGNAL**

What the drift costs. Annual payroll is about $36.6M. We found 22 to 28 percent of mapped capacity assigned to projects that support none of the four critical moves. The directional cost is **$670K to $850K per month**, sized from headcount and loaded cost.

Verdict by move

MOVE	VERDICT	IN ONE LINE
M1: Re-rate the core	YELLOW	The price sheet exists. Renewal calls still run the old catalog script, and no one owns the migration.
M2: Standardize enterprise delivery	RED	Bookings are ahead of target. The wins carry custom commitments that erode the margin the plan was priced on.
M3: Renewal-to-expansion	GREEN	The move has an accountable owner, the work is measured, and NRR has risen from 103 to 106 in two quarters. Leave it alone.
M4: Enterprise operating cadence	RED	The board pack shows the move on track. No team reports active work against it.

Three moves inside 30 days

- **Retire Legacy LMS v2.** It takes 18 percent of product capacity, about $1.9M a year, and no move claims it. Point the freed capacity at M1's migration playbook.
- **Delegate the deal desk.** Give the CRO approvals under $75K with published guardrails.
- **Name an owner for M4.** Require one artifact, the common pipeline definitions, in 30 days. If no owner can be named, defer the move and tell the board.

> **Decision.** Alignment is the constraint. The M2 and M4 findings require targeted validation before the associated high-stakes decisions. The move-level diagnosis and operating actions can proceed now.

How this document is built

- **Part A** is the value-creation thesis and the critical moves, locked with the CEO before any interviews. Everything else is read against that page.
- **Part B** is how the work lines up to the moves. One frame per move, with the same fields in the same order.
- **Part C** is the underlying theme. Three of the four moves stall in the same place, so it is treated once.
- **Part D** is the starting beliefs and which held up.
- **Part E** makes the same questions answerable every Monday.
- **Part F** is method, data, labels, and limits.

A The value-creation thesis and the critical moves

Locked with the CEO on day zero, in writing: reach $32M EBITDA by 2027 without breaking gross retention or gross margin. Convert the founding compliance-training base from catalog pricing to platform subscription. Grow enterprise on standard delivery so bookings convert to margin and live customers. Convert service renewals into expansion revenue. Run the company on one operating cadence with one CRM, common pipeline definitions, and a weekly view the CEO and CFO read the same way.

#	THE MOVE	WHAT IT FORBIDS
M1	Training catalog to platform subscription. GRR 90 to 94 and 8 percent realized price on renewal.	Discount-led renewals and save-the-logo concessions that re-anchor old pricing.
M2	Enterprise delivery standardized: three packages, services under 25 percent of first-year value, median time-to-live of 90 days.	Bespoke implementation commitments that erode margin and delay go-live.
M3	CS owns an expansion quota with playbooks by segment. NRR 103 to 109.	Expansion bought with unpriced service hours.
M4	One CRM, common pipeline definitions, and a weekly operating view whose numbers agree without reconciliation.	A reporting ritual that produces decks instead of decisions.

The forbids column names the way each number could go green while the thesis fails. Each move needs a clear failure condition. Otherwise leadership cannot tell whether it is working.

B How the work lines up to the moves

The same fields are used for every move. Run down the left column to compare. Causes shared across moves live in Part C.

M1: Re-rate the core

YELLOW

THE MOVE	Catalog to platform subscription on the founding product: 41 percent of revenue. GRR 90, down $0.6M last year. Target: GRR 94 and 8 percent realized price.
WHAT WE FOUND	Packaging and price are built and tested with three accounts. Downstream, 11 of 14 renewal reps could not describe the migration offer. All nine renewal calls sampled ran the old catalog script. Zero of the quarter's 63 renewals carried the new paper. **STRONG SIGNAL**
WHERE IT STALLS	Ownership. Product marketing believes CS owns the migration playbook. CS believes sales does. Two capable leaders each believe it is not their job. The decision that would settle it sits unowned. **STRONG SIGNAL**
WHAT IT COSTS	Every un-migrated renewal re-anchors old pricing for a year. About $2.1M of renewals rolled on catalog terms last quarter alone. At the current pace the 8 percent price move arrives a year late.
NEXT 30 / 90 / 180	30: name the owner. Kill LMS v2 and point the freed capacity here. 90: first 40 accounts migrated, scripts re-cut, discount guardrails published. 180: GRR and price realization read on the migrated cohort.

M2: Standardize enterprise delivery

RED

THE MOVE	Enterprise bookings grow on three standard delivery packages. Services stay under 25 percent of first-year value. Median time-to-live: 90 days. The margin plan was priced on standard delivery.
WHAT WE FOUND	Enterprise bookings are **$9.2M against a $7.5M target**, 23 percent ahead. Every dashboard shows this move green. And 62 percent of those wins carry custom implementation commitments negotiated into the contract, at an average 24 percent discount on services. **STRONG SIGNAL**
WHERE IT STALLS	Enterprise reps are paid on bookings. Scope exceptions have no named decider, so sales commits the work and delivery inherits it. Gross margin is 54 percent on the custom cohort and 71 percent on standard delivery. Median time-to-live is 190 days on custom work and 88 on standard. **STRONG SIGNAL**
WHAT IT COSTS	About $1.8M of margin gave way in the custom cohort this year, and the deployment backlog is 14 weeks. The dashboards do not surface this because the bookings target is being hit. **EMERGING PATTERN**
NEXT 30 / 90 / 180	30: scope exceptions get a named decider and a floor price. 90: re-price the three packages and require CRO approval for custom work above the threshold. 180: read margin and time-to-live by cohort and make a keep-or-kill decision on the custom motion.

B How the work lines up to the moves, continued

M3: Renewal-to-expansion

GREEN

THE MOVE	CS owns an expansion quota, priced into comp, with playbooks by segment. Target NRR 103 to 109.
WHAT WE FOUND	The quota is assigned, playbooks are live in two of three segments, and NRR moved from 103 to 106 in two quarters. The work is pointed where leadership decided. **STRONG SIGNAL**
WHERE IT STALLS	One watch item. Expansion in the smallest segment is running on about $18K a month of unpriced service hours, which the move forbids. The CS lead re-cut the rate card before the read ended. That response is evidence of an aligned team. **EMERGING PATTERN**
WHAT IT COSTS	Nominal if the rate-card fix holds. Re-check at the 90-day read.
NEXT 30 / 90 / 180	30: add no process. 90: verify the rate-card fix in margin. 180: confirm the third-segment playbook and the 109 trajectory.

M4: Enterprise operating cadence

RED

THE MOVE	One CRM, common pipeline definitions, and a weekly operating view the CEO and CFO read the same way. Carried in the board pack as a current priority.
WHAT WE FOUND	The board pack shows this move on track. The project inventory shows **not one team reporting real work against it.** There is no CRM consolidation workstream, pipeline definitions document, or weekly view. **STRONG SIGNAL**
WHERE IT STALLS	Nobody decided to stop the move. Its owner was never named, so it loses every capacity contest to owned work.
WHAT IT COSTS	Forecasts are assembled by hand from three systems and disagree by low single-digit millions depending on who pulls them. Until the definitions exist, M2's margin problem and M1's migration require a fresh project to track.
NEXT 30 / 90 / 180	30: name the owner and produce pipeline definitions v0. If no owner can be named, defer the move and tell the board. 90: scope one CRM and run the weekly view in two functions. 180: assemble the forecast without manual reconciliation.

C The underlying theme

Three of the four moves stall in the same place. M1 waits on a decision no one owns. M2 runs on a scope exception no one is allowed to refuse. M4 never got an owner. And the deal below dies in an approval queue. The same decision problem appears in four places. Fixing each instance separately would leave the underlying problem in place.

One deal, traced end to end

A cross-sell quote for a 340-rooftop group, followed through the company's systems and calendars.

DAY	WHERE THE DEAL WAS	WHO COULD MOVE IT
0	A rep builds a multi-product quote. The 19 percent discount needs manager approval.	Rep, then sales manager
1 to 3	The manager approves and forwards it to the deal desk under the month-two policy.	Deal desk
3 to 8	One analyst holds all non-standard terms and routes anything over $50K to the CEO.	One analyst, then CEO
8 to 11	The CEO trims the discount to 15 percent. Finance re-prices and returns the quote.	CEO, then finance
11	The customer has signed a one-year standalone renewal with the incumbent. The deal is dead.	No one

Five desks, eleven days, and two decisions. Over the trailing two quarters, **41 percent of cross-sell quotes over 30 rooftops expired or were abandoned in that queue.** STRONG SIGNAL

The ten decisions that move money, mapped

Each was mapped for who recommends, who decides, who executes, who can block, and observed median latency.

- **Six of ten** have the CEO as de facto decider, including discounts, roadmap trade-offs, and backfills that he believes he delegated.
  STRONG SIGNAL
- **Three of ten** have no named decider. That is where M1's ownership gap, M2's scope exceptions, and M4's missing owner live.
  STRONG SIGNAL
- Median latency at the CEO's desk is six days. Escalations are up 2.4x in his eight months. The organization is learning to wait for him. Leadership can publish the decision policy within 30 days: what he decides, what he does not decide, and the dollar lines.
  EMERGING PATTERN

D What we believed going in, and what held up

The read started from four beliefs about where the work goes. Three were the CEO's, one was ours. Each was tested against the work. Two held. Two did not.

1: "The re-rate is launched." Did not hold.

It is launched as pricing and unlaunched as work: zero of 63 renewals on new paper, and the playbook has no owner. The fix is an ownership decision. **STRONG SIGNAL**

2: "Cross-sell is stalled because the sales motion is failing." Did not hold.

The stall sits upstream of the sales motion. The month-two deal-desk step adds a median nine days and kills 41 percent of large quotes before any selling happens. Where quotes clear in under five days, quote-to-close conversion is above last year's. The failure concentrates in one approval queue. **STRONG SIGNAL**

> **What this changes.** A Q1 executive change was planned in the seat that owns cross-sell. Replace the executive and the queue remains. A successor inherits it, along with a $350K search and a year of onboarding. Whether the executive is right for the next chapter is a separate question. It can now be asked against a working system.

3: "Enterprise is our growth engine." Held, with a catch.

Bookings are 23 percent ahead of target. They are also being won with the custom commitments the plan forbids. The bookings instruments answer only the first question. **STRONG SIGNAL**

4: "Renewal-to-expansion is our best-run motion." Held.

The evidence confirmed it without prompting. When shown the one drift, the team fixed it inside the read. **STRONG SIGNAL**

E Make these questions answerable every Monday

This read was assembled by hand. The project inventory took four days to build because no system holds it, and the decision trail was reconstructed from CRM timestamps and calendars. Five feeds make the same questions answerable every Monday.

THE FEED	WHAT IT ANSWERS	WHAT IT TAKES
1: Reason-coded churn by product	Whether M1's retention verdict is improving each quarter.	A picklist enforced at closure in the CRM and two weeks of RevOps time. Free text stops being allowed.
2: One work inventory tagged to moves	Once every project carries a move tag, leadership can see work supporting no critical move without rebuilding the inventory.	Every project carries a move tag or "none" in the tools teams already use.
3: A decision log on ten money decisions	Latency and ownership visible weekly. The deal-desk queue in Part C emits this.	Report configuration on existing CRM timestamps. No new software.
4: Margin and time-to-live by cohort	The M2 green-number test, with standard and custom delivery side by side.	One warehouse view over billing, CRM, and the project tracker.
5: The Monday view	Per move, per team: what shipped against it last week.	The operating cadence M4 builds anyway. Teams state the work in their own words.

Items 1, 3, and 4 are configuration. Items 2 and 5 are the operating cadence M4 was already going to build. Instrumentation means the questions in this document no longer cost two weeks.

F Method, data, labels, and limits

Where the fourteen days went

- **Day 0:** the thesis and moves were locked with the CEO in writing before any interview.
- **Days 1 to 7:** two CEO sessions; seven 45-minute executive interviews; eight 20-minute interviews with operating leaders and the teams carrying M1 and M2. Every input was invited.
- **Days 8 to 11:** findings were built into decision-rights and capacity maps. Drift was costed and every finding labeled.
- **Days 12 to 14:** the readout was delivered to the CEO who commissioned it. Nothing moves further unless the CEO sends it.

The data requested, and what existed

Requested: ARR bridge by product; three years of churn by reason code; bookings by product and segment; the full project inventory; the last ten pricing exceptions and contract approvals with timestamps; the meeting-cadence inventory; the last three board decks; and comp plans for every quota-carrying role.

- Churn reason codes were free text for three of eight quarters and were normalized by hand. The findings that rest on them carry that limitation.
- No decision log existed. The route and latency in Part C were reconstructed from CRM timestamps, calendars, and email threads.
- The project inventory existed in no system and took four days plus every team lead to assemble. The assembly was itself a finding.

How findings are labeled

STRONG SIGNAL: directly observed and repeatedly evidenced. Act on it. **EMERGING PATTERN**: multiple indicators. The evidence is not yet triangulated. Weigh it. **HYPOTHESIS TO VALIDATE**: plausible and explicitly unproven. A test is named. Do not act on it alone.

Limits of this audit

- This is a comprehensive diagnosis of how the company is operating against its critical moves. It does not adjudicate which executive is right.
- It cannot determine customer intent because this audit uses deal files and churn codes, without customer interviews.
- It cannot establish seasonality or full-quarter comp effects without a full operating cycle.
- The M2 and M4 decisions require the targeted validation named in the readout.

This is a sanitized composite. The company, people, and figures are invented and blended; the method, report register, and depth reflect the day-14 artifact. Read Anatomy of an Audit, visit the TL Partners library, or write thomas@tlpartners.biz. tlpartners.biz