

FIELD BRIEF · V1.0 · MAY 2026

A Mid-Market Private Credit Fund's First 90 Days on F2

Field Brief by Filip Balenko

Application — F2 AI Deployment Strategist

Founder's Note

A note on what this is.

I'm Filip Balenko. I'm applying for F2's AI Deployment Strategist role and I built this Field Brief instead of writing a standard cover letter.

This is methodology, not product. There's no working memo generator on this site, no agentic demo, no half-rebuilt version of F2. What's here is the deliverable an AI Deployment Strategist would produce in week one to land a real F2 customer: a customer persona, a workflow decomposition, a prompt library, an output audit rubric applied to a deliberately flawed sample memo, and a 30/60/90 deployment plan.

The customer is fictional — Meridian Credit Partners, a composite modeled on publicly disclosed mid-market direct lenders in the \$1–2B AUM band. The pain points and workflow patterns are drawn from cited public sources, including F2's own published research where it grounds the customer-pain framing.

I've been the underwriter at Scale Capital's specialty credit branch, the operator building AI-driven underwriting at Extra Business Funding, and the engineer shipping agentic systems at sfv-ic.com. The deployment role asks for someone who can speak underwriter and engineer in the same meeting. I've been both.

— Filip

§ 1

Customer Persona — Meridian Credit Partners

Meridian Credit Partners is a fictional composite, modeled on the publicly disclosed strategies, fund sizes, and team shapes of mid-market sponsor-led direct lenders in the \$1–2B AUM band. Names and specific deal details are placeholders; pain points and workflow patterns are drawn from cited public sources.

Snapshot

- **Strategy** — Sponsor-led direct lending, North American mid-market
- **AUM** — \$1.2B
- **Fund II** — \$750M, 78% deployed across 14 active names
- **Fund III** — \$450M raised on a \$750M target; first close Q2 2026
- **Hold profile** — \$15M–\$60M unitranche; selective second-lien; clubs up to \$90M with rotating co-lenders
- **Sectors** — Vertical SaaS, tech-enabled business services, healthcare services (PPM, multi-site outpatient), light industrial / specialty manufacturing
- **HQ** — Midtown Manhattan, 14 FTE total
- **Vintage** — Founded 2018; first fund \$300M, fully realized 2024 at 1.6x net

Team (12 investment professionals)

- **Managing Partner** — *Steve K.*, 22 years, ex-tier-one direct lender, founder
- **Partners (2)** — sector heads (SaaS/Healthcare; Industrial/Services)
- **Principals (3)** — run live processes end-to-end
- **Associates (4)** — model build, memo draft, portfolio monitoring
- **Analysts (2)** — spreading, comp screening, data-room hygiene

Plus 2 non-investment FTE (CFO/IR/Fund Admin) outside scope.

Deal flow (LTM)

- ~40 IOIs reviewed
- ~22 second-round meetings
- ~12 advance to confirmatory diligence
- **6–8 close**
- Median IOI-to-funding: 9 weeks (sponsor-led repeats; cold processes run 12–14)

Current stack

- **Modeling** — Excel. Bespoke per analyst. No firm-wide template library.
- **CRM / pipeline** — DealCloud, lightly utilized. Partners track top-of-funnel in their heads.
- **Data rooms** — SharePoint (sponsor-managed, ~70%); Datasite (~30%)
- **Expert calls** — Tegus, occasional GLG

- **Comps / market data** — Refinitiv LPC, S&P LCD, internal closed-deal log (Excel)
- **No AI tools deployed.** Two Associates use personal ChatGPT Plus accounts. Firm has no AI policy yet.

A Tuesday with Anna, Senior Associate

- **8:15** — Inbox. Three teasers from MD-tier sponsors. Forwards one to her Principal with a 2-line read.
- **8:45** — Opens Q3 monitoring template for portfolio name *Beacon Software*. Pulls EBITDA off the borrower's reporting package, reconciles to the credit agreement definition, updates the model. Maintenance covenant test in two weeks.
- **10:30** — Spreading. Sponsor model came with the CIM — vertical SaaS, ~\$24M EBITDA. Anna rebuilds it in firm template: standardizes revenue cohorts, recategorizes addbacks, maps QoE adjustments. **Three hours of structured drudgery.**
- **13:30** — Management call. Sponsor's CFO walks through working-capital seasonality on a different live deal. Principal drives; Anna takes notes.
- **15:00** — Comp screen for Thursday's IC. Pulls four "comparable" closed deals from the firm log; builds a side-by-side leverage / pricing / structure tear sheet. The judgment call on which deals are comparable is hers — and she'll get challenged on it Thursday.
- **17:30** — IC memo draft, sections 1–3. Iterates with her Principal until 9pm.

Day's split — ~4 hours real analysis, ~3 hours structured drudgery, ~1 hour reactive. The drudgery is what F2 is buying back.

Pain points — common to mid-market direct lenders in this peer group

1. **Spreading drag.** Associates spend disproportionate hours rebuilding sponsor models in firm template — addback recategorization, KPI normalization, schedule cleanup. PwC names spreading financial statements among the repetitive, time-consuming underwriting tasks that AI-driven origination tools are now automating, alongside risk rating calculations, document processing, and compliance checks. *Source: PwC, "Strategies for banks to compete and work with private credit," 2025.*

1. **Screening turnaround.** Partners want a 1-pager 24 hours after the teaser, not five days. F2's own analysis frames IC memos as the bottleneck — teams spending 20–40 hours assembling data, formatting charts, and aligning to credit boxes — with memo turnaround speed determining who wins the allocation in competitive auctions. Or in Don Muir's framing: firms that win the bid are those that can move first. *Source: F2 blog, "How Private Credit Teams Use AI to Accelerate Underwriting," citing High Yield Harry; Don Muir (CEO, F2).*

1. **LP DDQ burden.** The average PE fund now responds to 150+ DDQs annually — up 40% from five years ago — with each questionnaire averaging 250 questions across 21 categories. ILPA DDQ completion alone runs 40–60 hours manually before AI tooling brings it to 8–12. ILPA's own framing of the historical state: customized DDQs created an "extraordinary administrative burden" on GPs and LPs alike. *Source: AutoRFP.ai analysis of ILPA DDQ 2.0; ILPA Due Diligence Questionnaire 2.0 (Nov 2021).*

1. **Institutional memory gap.** Benchmarking new deals against firm history is, in Golub Capital's framing as cited by F2, "anecdotal and inconsistent" without a deal library layer. The Q&A and notes from past IC meetings represent the firm's collective investment judgment but get buried in disorganized SharePoint folders. *Source: F2 blog citing Golub Capital; Meridian AI, "How Top PE Firms Run Investment Committees in 2025."*

Buying motion

Top-down adoption following a competitive loss is the modal entry point for AI tooling in this peer group: a Managing Partner watches a competitor execute faster on a live deal, a vendor demo lands at an industry event (SuperReturn, ALM private credit conferences, ABS East), and procurement moves top-down on a 60–90 day evaluation cycle.

Skepticism concentrates in the middle of the org. Principals worry about audit trail integrity in LP DDQ contexts.

Associates quietly worry about being deskilled. **That skepticism is the deployment job** — not the contract.

Workflow map and F2 plug-in points

Sourcing ' [Pre-screen / IOI] ' Term sheet ' [Confirmatory diligence]
' [Credit memo build] ' [IC packet] ' Docs ' Funding ' [Portfolio monitoring]

Bracketed stages = the **5 highest-leverage F2 plug-in points** at Meridian:

1. **Pre-screen / IOI** — Excel Intelligence ingests sponsor model + CIM; AutoReports generates the 24-hour 1-pager Steve has been asking for
2. **Confirmatory diligence** — Sensitivity packs, QoE addback validation against source docs, comp screens grounded in firm history (not Refinitiv averages)
3. **Credit memo build** — AutoReports drafts sections 1–3 of firm template; Associates edit upward instead of starting from blank
4. **IC packet** — Audit Mode supplies source-linked exhibits; partners click any number, see the cell
5. **Portfolio monitoring** — Quarterly actuals vs. base case, covenant headroom flags, commentary drafted in firm voice

Where F2 does not plug in for Meridian (explicitly out of scope, day 1): sourcing (relationship-driven), term-sheet negotiation (judgment), legal documentation (counsel), funding mechanics (treasury).

§2

Workflow Decomposition

How to read this section

Workflow decomposition is the deployment strategist's first analytical job. You take a customer's stated workflow, break it into discrete tasks, and identify (a) which tasks the AI platform can plausibly handle, (b) what the realistic time savings look like under conservative assumptions, and (c) which subset to launch with on day one to maximize the probability of an internal champion emerging by day 30.

The trap to avoid: treating every F2-eligible task as equally good day-one material. F2 can plausibly touch every stage of Meridian's underwriting workflow. But launching on all of them simultaneously is how deployments fail — too many surface areas to debug, too many opportunities for a bad output to land in front of a skeptical Principal, too thin a layer of internal advocacy when the inevitable first failure occurs.

The rest of this section maps F2's capabilities to Meridian's tasks, shows the time-savings math transparently, and argues a specific three-workflow launch sequence for days 1–30.

F2 capability 'Meridian task 'time savings

The mapping below pairs each of F2's stated capabilities with the specific Meridian task it touches, the conservative baseline hours it costs Meridian today, and the realistic time savings under F2.

Conservative baseline assumption: F2 publicly states up to 45 hours/week saved per analyst and 60–75% faster diligence. Both numbers represent the upper bound of customer outcomes. For deployment planning, I'm anchoring on the lower end of the speedup band (~50% reduction) with workflow-specific adjustments — higher where the task is structured (lookups, math, formula tracing) and lower where human judgment dominates (narrative, addback validation, covenant design).

Capability 1 — Excel Intelligence

- **Meridian task:** Sponsor model intake and rebuild in firm template. Addback recategorization (sponsor's "EBITDA" 'Meridian's adjusted EBITDA per credit agreement definition). KPI normalization. Schedule cleanup. Reconciliation against QoE report.
- **Baseline:** 15–25 hours per deal at the diligence stage. Anna's Tuesday spreading session was three hours on a single sub-task; full intake of a \$24M-EBITDA SaaS borrower's model runs across 2–3 working days for a Senior Associate.
- **F2 speedup:** 60–70% reduction. Aggressive end of the F2 band because spreading is structured and benefits directly from formula-level fidelity. Conservative on the upper bound because addback validation still requires Associate judgment — the model can flag a \$2.1M "one-time consulting" addback, but a human has to decide whether the underlying invoices in the data room support it.
- **Per-deal savings:** ~10–17 hours.
- **Annual savings at Meridian:** 12 deep-diligence deals × 13 hours = ~**156 hours/year**, or ~4 weeks of Associate time.

Capability 2 — Full-context analysis (comp screening & benchmarking)

- **Meridian task:** Pull comparable closed deals from firm history and market data. Build leverage / pricing / structure tear sheet. Defend the comparable set in IC.
- **Baseline:** 3–4 hours per comp screen at Meridian today, with the upstream cost of maintaining the firm's Excel-based closed-deal log eating another ~5 hours/quarter of Associate time.
- **F2 speedup:** 75–85% reduction once Meridian's deal library is loaded. Highest-confidence speedup category — this is structured data lookup against a known schema, F2's strongest ground.
- **Per-deal savings:** ~2.5–3 hours.
- **Strategic value beyond hours:** Addresses what Golub Capital, as cited in F2's own materials, calls "anecdotal and inconsistent" benchmarking. The firm-history layer is a moat once it's built — every subsequent deal benefits from incremental enrichment of the deal library.
- **Annual savings:** 12 deals × 2.75 hours = ~**33 hours/year direct**, plus removal of the quarterly deal-log maintenance burden (~20 hours/year).

Capability 3 — Auto-calculation and sensitivities

- **Meridian task:** Build sensitivity packs for confirmatory diligence and IC. Standard cases: revenue declines, margin compression, working-capital stress, covenant headroom under each. Light auto-calc of EBITDA, leverage, interest coverage, FCC.
- **Baseline:** 4–8 hours per deal depending on capital structure complexity. Unitranche with simple covenant package **4** hours; structures with PIK toggles, springing covenants, or asset-based revolvers **8** hours.
- **F2 speedup:** 70–80% reduction. Fully structured math, F2's stated 99% accuracy on core financial calcs is the relevant benchmark here, and the 1% failure mode is bounded — wrong outputs are catchable on review.
- **Per-deal savings:** ~3–6 hours.
- **Annual savings:** 12 deals × 4.5 hours = ~**54 hours/year**.

Capability 4 — AutoReports (IC memo generation)

- **Meridian task:** Draft the screening memo (1–3 pages) and full IC memo (15–25 pages with appendix). Sections include investment thesis, business overview, financial analysis, risk and mitigants, covenant package, recommendation.
- **Baseline:** F2's own analysis cites 20–40 hours of team time per memo, attributing the figure to High Yield Harry. Meridian's mid-market unitranche memos sit in the 25–35 hour range based on peer-group convention; a Senior Associate runs the draft, the Principal edits, and the cycle absorbs 1.5–2 working days of combined investment-team time.

- **F2 speedup:** 40–55% reduction on first draft, with the remaining time going to Associate and Principal edits. Lower than the structured workflows because narrative quality and persuasion remain human work — F2 produces the scaffolding, Meridian's people produce the conviction.
- **Per-deal savings:** ~12–18 hours.
- **Annual savings:** 12 deals × 15 hours = ~180 hours/year.
- **Why this is not a day-30 launch workflow:** The output goes directly in front of Steve and the IC. Bad memo output on day 25 risks killing the deployment. Defer until trust is established.

Capability 5 — Audit Mode

- **Meridian task:** Source-link every quantitative claim in IC packets and LP-facing materials. Survive partner Q&A and LP DDQ scrutiny.
- **Baseline:** Hard to bound as standalone hours. Audit and source-linking work is interleaved into memo drafting; if the underlying memo is rebuilt in F2 with Audit Mode on, the source-linking is effectively free.
- **F2 speedup:** Not a "speedup" capability. This is a **trust and defensibility** capability. The right framing for Steve and the LPs is not "saves N hours" — it's "every number F2 generates is traceable to a cell or doc, which is a precondition for any of this output going near a credit committee or LP DDQ."
- **Strategic value:** Audit Mode is the answer to Principal-level skepticism about audit trail integrity. It is a deployment **enabler** more than an efficiency gain.

Capability 6 — Portfolio monitoring (auto-calc on actuals)

- **Meridian task:** Quarterly review of each portfolio company. Pull EBITDA off borrower reporting, reconcile to credit agreement definition, update tracker, test maintenance covenants, draft one-paragraph commentary in firm voice.
- **Baseline:** 2–4 hours per name per quarter. Fund II has 14 active names, so roughly 35–55 hours of Associate time per quarter on monitoring.
- **F2 speedup:** 50–60% on the analytical work, with the commentary draft as a bonus that absorbs another chunk of formerly-manual time.
- **Per-quarter savings:** ~22 hours across the portfolio.
- **Annual savings:** ~88 hours/year.

Capability 7 — LP DDQ response (AutoReports + Audit Mode applied to operational content)

- **Meridian task:** Respond to LP DDQs during Fund III fundraising. ILPA DDQ 2.0 alone runs ~250 questions across 21 sections. Per AutoRFP's analysis of the standard, manual completion runs 40–60 hours per DDQ; AI tooling brings that to 8–12.
- **Baseline at Meridian:** ~2 weeks of Senior Associate time per quarter during active fundraising — consistent with persona pain point #3.
- **F2 speedup:** 70–80% on quantitative sections (track record, deal-by-deal attribution, risk-rating reconciliation). Lower on qualitative sections (team bios, ESG narrative).
- **Annual savings during fundraising:** Highly variable — depends on LP volume. Expect ~120–200 hours saved across a fundraising year.

Total annualized savings — conservative

Adding up the per-workflow estimates:

Spreading (Excel Intelligence)	~156 hrs/yr
Comp screening (Full-context)	~53 hrs/yr (incl. deal-log maintenance)
Sensitivities (Auto-calc)	~54 hrs/yr

IC memo drafting (AutoReports)	~180 hrs/yr
Portfolio monitoring	~88 hrs/yr
LP DDQ (during fundraising)	~150 hrs/yr (mid-point estimate)
	~681 hrs/yr

~681 hours/year = roughly 17 weeks of Associate FTE recovered, or 0.34 of a full-time-equivalent across Meridian's 6-Associate-and-Analyst layer. This is well below F2's headline "45 hrs/week saved per analyst" because (a) Meridian is a 12-person mid-market shop, not an upper-middle-market fund running 30+ active diligences, and (b) I'm using the conservative end of every speedup band.

The right way to frame this number for Steve is not "F2 saves you 681 hours." It is: **"F2 buys back roughly one-third of an Associate's time across the firm — which lets you take on two more diligences a year without adding headcount, or alternatively redirects existing capacity from drudgery to source-development and live-deal work."**

That second framing is what wins the buy. Hours-saved is a feature claim. Capacity-redirected is a business case.

Day-30 launch ranking

The principle: trust-building before partner-facing

The natural instinct is to launch on the highest-savings workflow first. By the math above, that would be IC memo drafting at ~180 hours/year. **That is the wrong call.**

The day-30 question is not "where is F2 most valuable?" — it is "which workflows, if they go well, build enough internal trust that the deployment survives its first failure?" Three constraints govern the answer:

- 1. Skepticism concentrates in the middle of the org.** Principals worry about audit integrity; Associates quietly worry about being deskilled. Day-30 launch workflows must convert at least one Principal and at least two Associates into champions.
- 2. Bad output cannot reach the IC in the trust-building window.** A flawed F2 memo presented to Steve in week three kills the deployment. Period. Any workflow whose output is partner-facing within 30 days is too high-risk.
- 3. Time savings must be visible.** Abstract "you'll save hours over the year" loses to concrete "this took Anna two days last month and ninety minutes this morning." Workflows where Associates feel the saved time directly are worth more than workflows with bigger annualized numbers but diffuse experience.

Under those constraints, the top three day-30 workflows are spreading, comp screening, and sensitivity packs — in that order.

#1 — Spreading & financial model intake (Excel Intelligence)

Highest-priority launch workflow for four reasons:

Visible time savings to the right audience. Spreading is what Anna is doing on Tuesday morning. Cutting that 3-hour session to 45 minutes is the single most viscerally felt improvement in the deployment. Associates become champions on day three, not day thirty.

Quality is verifiable without Principal involvement. A spread model is checkable against the source documents. An Associate can validate F2's output by spot-checking against the QoE report and the audited financials. No Principal needs to sign off on the workflow before it produces value, which means the deployment doesn't gate on Steve's calendar.

Low partner-facing risk. The spread model is internal scaffolding — it feeds the memo, the comps, the sensitivities. Partners don't read spreads; they read what comes out of them. If F2 mishandles a sponsor model on day 12, the cost is one Associate redoing a model. The cost is bounded.

Builds the substrate for everything else. Comps, sensitivities, and memos all consume the spread model. Getting Excel Intelligence working cleanly in week one is a precondition for the workflows that come next.

Day-30 milestone: All four Associates running F2-assisted spreading on at least one live deal each. Internal validation rate (Associate accepts F2 output without material rebuild) at 80%+ on simple cap structures.

#2 — Comp screening & deal library (Full-context analysis)

Second-priority workflow for three reasons:

Direct hit on a publicly-cited pain point. "Anecdotal and inconsistent" benchmarking is the Golub Capital framing F2 already cites in its marketing. Demonstrating that F2 closes that gap at Meridian within 30 days is a story Steve can repeat to LPs and to peer Managing Partners.

The deal library is a compounding asset. Loading Meridian's last five years of closed deals into F2 in week one is a one-time cost that pays back on every subsequent comp screen. Day 30 looks more impressive than day 5 because the library is denser.

Low partner-facing risk, moderate Associate visibility. Comp screens go into IC packets, but they are a discrete exhibit — partners challenge the comparable set, not the math underneath it. Associates feel the time savings (3 hours '30 minutes) and the quality improvement (firm-history-grounded comps vs. pulling four deals from memory and a spreadsheet).

Day-30 milestone: Last 5 years of closed deals loaded and tagged. Two live diligences using F2-generated comp screens that survived Principal review with zero structural rebuild.

#3 — Sensitivity packs (Auto-calculation)

Third-priority workflow for two reasons:

99% accuracy on core financial calcs is F2's most defensible claim. Sensitivities are pure math. If F2 fails here, F2 fails on its strongest ground — and the deployment was always going to surface that failure, better in week three than week thirty. If F2 succeeds, Meridian gets credibility on F2's headline accuracy claim, which is the foundation for trusting the platform on more judgment-heavy workflows later.

Direct support for Principal-level work without putting F2 in front of partners. Sensitivities are run by Associates, reviewed by Principals, and presented in IC by Principals. The Principal becomes the buffer between F2 output and partner audience — exactly the right structure during the trust-building window.

Day-30 milestone: Standard sensitivity templates configured (15% revenue decline + margin compression; covenant headroom test; downside FCC). At least one Principal has presented an F2-generated sensitivity pack at IC and defended it under partner Q&A.

What I deliberately did *not* launch on day 30

Three workflows with strong ROI cases that I am intentionally deferring past day 30:

IC memo drafting (AutoReports). Highest annualized savings, highest partner-facing risk. Defer to days 31–60, after spreading and sensitivities have proven F2's quantitative reliability. First F2-assisted memo at IC is the **day-45 milestone**, not the day-15 milestone. Launch sequence: F2 generates sections 1–3 (Investment Thesis, Use of Proceeds, Sources & Uses), Senior Associate edits, Principal reviews against firm template before any partner sees it. Audit Mode runs in parallel from day one to make the eventual memo launch defensible.

Portfolio monitoring. Quarterly cadence means the workflow doesn't produce visible day-30 wins unless a quarterly cycle happens to fall inside the launch window. Better as a days-60–90 expansion when there's a natural quarter-close to anchor on. The exception: if Meridian's Q4 close lands in the launch window, this moves up to day 30.

LP DDQ response. Tactical, timing-dependent. Worth pre-configuring the firm-content library (track record, deal attribution, ESG narrative, operational policies) during days 1–30 so it's ready when a DDQ arrives — but the workflow itself is "launched" the moment the next LP DDQ comes in, not on a fixed day.

What success looks like at the end of day 30

Concrete metrics for the QBR-equivalent at day 30:

- **Adoption breadth:** 4 of 4 Associates have run F2 on at least one live deal. 2 of 3 Principals have reviewed F2 output and approved it for IC packet inclusion.
- **Time-savings realized:** Tracked per-deal hours-spread for two live diligences. Target: 50%+ reduction vs. Meridian's 90-day historical baseline on comparable deal types.
- **Output quality:** Internal validation rate (output accepted without material rebuild) at 80%+ on spreading, 85%+ on comps, 90%+ on sensitivities.
- **Audit posture:** Every F2-generated number in any IC-bound exhibit is source-linked. Zero exceptions.
- **Champions identified:** At least one Associate and one Principal who would defend F2 in a hallway conversation with Steve.

Day-30 success is not "F2 is fully deployed at Meridian." Day-30 success is "F2 has earned the right to launch IC memo drafting in days 31–45."

§3

Prompt Library

How to read this section

This is the technical core of the field brief. Each prompt below is a production-grade template I'd configure for Meridian in the first 60 days, organized into the three deployment tiers from Section 2 so that the launch sequence stays intelligible.

The honest framing: I haven't seen F2's actual prompt syntax or how it parses retrieval against the deal library. These prompts are **structured asks** — they demonstrate what an experienced underwriter knows to specify, what guardrails they put on the output, and what audit hooks they leave for the reviewer. Plug them into F2's actual prompt interface and they'll need syntactic adaptation. The underwriting logic doesn't change.

Prompt design principles

Five principles hold across every prompt below. Each prompt is engineered to satisfy all five:

1. **Source-link by default.** Every quantitative claim must trace to a cell, doc page, or data-room file. Audit Mode is on, always.
2. **Force structure on the output.** Free-form responses are unaudited responses. Tables, schemas, and explicit field names produce outputs that an Associate can validate in 90 seconds instead of 9 minutes.
3. **Define the failure mode in the prompt.** "Flag any X without Y" is a better instruction than "evaluate X." Telling the model what *failure* looks like is more reliable than telling it what success looks like, especially on judgment-heavy tasks.
4. **Bound the universe.** "Last 5 years," "deals >\$25M EBITDA," "sectors in \{SaaS, healthcare services, light industrial\}." Unbounded retrieval is how F2 pulls the wrong comp set and an Associate spends 20 minutes figuring out why the leverage stat looks weird.
5. **Match firm voice explicitly.** Memo drafts, monitoring commentary, and DDQ responses each have a tone Meridian uses and a tone Meridian doesn't. The prompt names the template version.

Tier 1 — Day 1 launch prompts

Prompt 1 — Sponsor model intake and spread

Use case. Senior Associate has received the sponsor's CIM, financial model (Excel), and QoE report (PDF). Goal: ingest all three, output a standardized P&L spread in Meridian template format, ready for addback validation (Prompt 2) and sensitivity work (Prompt 4). Triggered at the start of confirmatory diligence.

Why this prompt matters at deployment. Spreading is the most viscerally felt time savings in the deployment. If F2 takes Anna's three-hour Tuesday morning spread and turns it into a 45-minute review-and-validate session, the Associate layer becomes champions on day three.

Role: You are spreading a borrower's financials for Meridian Credit Partners' confirmatory diligence on a sponsor-led unitranche transaction.

Inputs:

- Sponsor financial model (Excel) – full historical and projection workbook
- Quality of Earnings report (PDF) – issuer's QoE provider
- CIM (PDF) – sponsor's marketing materials

Task: Build a standardized P&L spread in Meridian Template v4 format. Cover the last 3 fiscal years (FY-3, FY-2, FY-1) plus TTM and the next 5 projection years.

Output structure (table):

- Row labels: Revenue (by segment if disclosed) ' Gross Profit ' SG&A ' Reported EBITDA ' Management Adjustments ' Adjusted EBITDA ' D&A ' EBIT ' Interest ' Taxes ' Net Income
- Columns: FY-3A | FY-2A | FY-1A | TTM | FY+1E | FY+2E | FY+3E | FY+4E | FY+5E
- Each cell must include a source link (workbook tab + cell reference, or PDF page + section)

EBITDA definition: Use Meridian's standard credit-agreement EBITDA definition (see firm template v4 §3.2). If the sponsor model uses a different definition, build BOTH columns side by side and flag the delta.

Constraints:

- Currency: report in USD. If source data is multi-currency, flag and use prevailing rate at period end; cite the rate source.
- Fiscal year: if borrower's fiscal year is non-calendar, label clearly and note in a footer.
- Do not reclassify revenue or expense lines without explicit instruction.
- If a line item appears in QoE but not the sponsor model (or vice versa), surface in a "Reconciliation Items" table beneath the main spread.

Confidence flag: Tag each TTM and projection figure with HIGH / MEDIUM / LOW confidence based on source-doc support strength. LOW confidence requires an explicit one-line rationale.

Expected output shape. Standardized 11-row × 9-column P&L table with source links in every cell. Reconciliation Items table beneath, listing any sponsor-model-vs-QoE deltas. Footer with currency, fiscal year, and EBITDA definition notes. Confidence tags on TTM and projection columns.

Edge cases

- **Multiple EBITDA definitions in source documents.** Sponsor model defines EBITDA inclusive of stock-based comp; QoE excludes it; Meridian's credit agreement defines it net of capitalized software development. The prompt forces side-by-side reconciliation rather than picking one silently.
- **Non-standard fiscal year.** Borrower runs a 4-4-5 retail calendar or a non-calendar FY ending in March. Output must label clearly to avoid downstream IC confusion.
- **Addbacks embedded in revenue lines.** Sponsor pre-bakes a 3% revenue uplift for "pricing initiatives launched Q4." This needs to surface as an addback in the projection columns, not a clean revenue line.
- **Foreign currency exposure.** Borrower has a UK or EU subsidiary reported in GBP/EUR. Translation method (period-end vs. average rate) must be explicit and consistent with the projection columns.
- **Stub periods.** Acquisition closing mid-quarter creates stub financials that don't roll up cleanly into the historical spread. Surface separately, don't average into a full-year figure.

Audit checklist

1. **Source link present on every cell.** Audit Mode should produce zero exceptions. Any cell without a source link is a deployment failure, not a content failure.
2. **Sum-check on revenue and EBITDA.** Sum of revenue segments equals total revenue in every column. Reported EBITDA + Management Adjustments equals Adjusted EBITDA in every column. Math is verifiable in 30 seconds.
3. **EBITDA definition flag is present and correct.** The reviewer can tell at a glance whether F2 used Meridian's credit-agreement EBITDA, the sponsor's EBITDA, or both side-by-side. If both, the delta line is populated.

Prompt 2 — Addback identification and validation

Use case. After Prompt 1 produces the spread, Associate runs Prompt 2 to validate every management adjustment against data-room evidence. This is where sponsor models tend to inflate run-rate EBITDA via undocumented addbacks; the deployment-strategist value is in catching these before they reach the credit memo.

Why this prompt matters at deployment. Addback validation is judgment-heavy work that bores Associates and lands in IC as a partner attack vector ("I don't believe the \$2.1M consulting addback"). F2 doing the structured part — categorize, size, evidence-check — lets the Associate spend their time on the actual judgment call.

Role: You are validating management addbacks proposed in the borrower's adjusted EBITDA reconciliation.

Inputs:

- Output of Prompt 1 (Meridian-template P&L spread)
- Sponsor model addback bridge (Excel tab)
- QoE report addback section
- Data room – invoices, contracts, board materials, prior-period reporting

Task: For every addback in the sponsor's bridge from Reported EBITDA to Adjusted EBITDA, output the validation table below.

Output structure (table):

- Addback Name | Category | \$ Amount | % of Run-Rate EBITDA | Recurring/One-Time Classification (sponsor) | Recurring/One-Time Classification (Meridian view) | Supporting Evidence | Confidence | Action Flag

Categories (use exactly):

- One-time professional fees
- Restructuring / severance
- Legal / settlement
- M&A transaction costs
- Pro forma cost savings (run-rate)
- Pro forma revenue (run-rate)
- Owner compensation normalization
- Stock-based compensation
- Other (specify)

Validation rules:

- For each addback >5% of run-rate EBITDA: require explicit doc-cite to invoice, contract, board minute, or QoE supporting schedule. No cite = Action Flag = "INSUFFICIENT EVIDENCE."
- For addbacks tagged "Pro forma cost savings" or "Pro forma revenue": require evidence the initiative is signed/contracted/launched, not just planned. Unsigned = Action Flag = "PRO FORMA UNSUPPORTED."
- For addbacks classified by sponsor as one-time: cross-check against prior 3 fiscal years. Recurring 2+ years = Action Flag = "RECURRING DESPITE ONE-TIME CLAIM."
- Sum of all addbacks must reconcile to sponsor's stated total. If reconciling delta exists, flag.

Output also: a one-paragraph "Adjusted EBITDA Adjustment Summary" stating which addbacks Meridian should haircut, by how much, and why. Quantify the proposed haircut to run-rate Adjusted EBITDA.

Expected output shape. Addback table with 9 columns, one row per management adjustment. Beneath the table:

reconciliation check vs. sponsor total, plus a one-paragraph Meridian haircut recommendation with quantified delta to Adjusted EBITDA.

Edge cases

- **Addbacks without supporting documents.** Sponsor claims \$1.4M of "synergies from prior acquisition" without a memo or schedule in the data room. Prompt forces "INSUFFICIENT EVIDENCE" flag rather than silent acceptance.

- **Recurring addbacks falsely categorized as one-time.** "Strategic consulting fees" appearing in three consecutive years' addback bridges. Cross-check rule catches this.
- **Pro forma addbacks for unconsummated initiatives.** Sponsor adds back \$3M of "annualized cost savings from headcount reduction announced last quarter." Validation rule requires evidence of actual termination, not announcement.
- **Owner compensation normalization in non-PE-owned targets.** Founder-owned business with \$2M of personal expenses run through the company. Legitimate addback, but the supporting evidence threshold is higher than for sponsor-led targets.
- **QoE adjustments not in sponsor bridge.** QoE provider proposes additional addbacks the sponsor didn't include. These should surface in "Other (specify)" with a flag asking whether Meridian wants to credit them.

Audit checklist

1. **Every addback >5% of run-rate has a doc citation.** Spot-check at random — pull three addbacks above the threshold and verify the cited document actually supports the addback amount and category. If even one citation is wrong or hallucinated, halt deployment of this prompt and ticket to F2 engineering.
2. **Recurring-vs-one-time logic is defensible.** Pull the prior-period addback bridges and confirm F2's "RECURRING DESPITE ONE-TIME CLAIM" flags are correct. False positives here are recoverable; false negatives are how a bad addback survives to IC.
3. **Reconciliation closes.** Sum of addbacks in F2's table equals sponsor's stated Adjusted EBITDA bridge total within \$25K. Reconciling deltas are explicitly surfaced.

Prompt 3 — Comp screen against firm history

Use case. Associate building the comp tear sheet for IC. Goal: surface 5–7 truly comparable closed Meridian deals from the firm's deal library, present leverage / pricing / structure side by side, flag structural outliers in the proposed deal vs. the comp set. Direct hit on Golub Capital's "anecdotal and inconsistent" benchmarking pain point.

Why this prompt matters at deployment. F2's full-context analysis lives or dies on whether it can find the right comparables. A bad comp set in IC is worse than no comp set — partners challenge the comparability and the credibility of the whole memo erodes. Getting this prompt right inside the first 30 days demonstrates F2 can handle Meridian's institutional memory layer.

Role: You are building a comparable-deals tear sheet for an IC packet at Meridian Credit Partners. The comp set must be drawn from Meridian's closed deal library, not external market data.

Inputs:

- Subject deal: borrower name, sector (specify GICS sub-industry), proposed capital structure (unitranche / 2nd lien / split), proposed leverage, proposed pricing (SOFR + spread, OID, fees), proposed covenant package, EBITDA size, sponsor identity
- Meridian deal library: all closed deals 2021–present, tagged by sector,

size, structure

Task: Identify 5-7 closed Meridian deals most comparable to the subject deal. Output a side-by-side tear sheet plus a structural outlier analysis.

Comparability criteria (in priority order):

1. Same GICS sub-industry, OR adjacent sub-industry within shared end-market exposure
2. EBITDA size within 0.5x-2.0x of subject deal
3. Same capital structure (unitranche to unitranche, not unitranche to 2nd lien)
4. Same sponsor archetype (mega-cap, MM PE, lower-MM PE, founder-owned)
5. Closing date within last 24 months (rate environment match)

If fewer than 5 deals satisfy criteria 1-4, expand to adjacent sectors and flag the comparability gap.

Output structure (table):

- Columns: Subject | Comp 1 | Comp 2 | ... | Comp N
 - Rows: Borrower | Sector | Sponsor | Closing Date | EBITDA at Close | Total Leverage | Senior Leverage | Structure | SOFR Spread | OID | Upfront Fee | Maintenance Covenant | Covenant Cushion | Equity Cure Rights | Hold

Outlier analysis (text below table):

- Identify any subject-deal terms that fall outside the inter-quartile range of the comp set. Quantify the deviation.
- For each outlier, propose: (a) is this a structural concern? (b) what would bring it in line? (c) is there a deal-specific rationale for the deviation that justifies acceptance?

Comparability disclosure:

- For each comp included, write one sentence stating why this deal qualifies. If a criterion was relaxed, name the relaxed criterion explicitly.

Date range: Disclose at top of output ("Comp set drawn from N closed Meridian deals between [Date] and [Date]").

Expected output shape. Side-by-side tear sheet table (subject + 5–7 comps × 14 rows of attributes). Below the table: outlier analysis identifying any subject-deal terms outside the comp set's inter-quartile range, with deal-specific rationale. Comparability disclosure (one sentence per comp). Date range at top.

Edge cases

- **Insufficient firm history in sector.** Meridian's first deal in a niche vertical (e.g., specialty insurance services). Prompt forces relaxation of the sector criterion with explicit disclosure rather than silent substitution.
- **Sponsor-led vs. non-sponsor mixing.** Sponsor-led deals price tighter than non-sponsor-led equivalents. Prompt's sponsor-archetype criterion catches this; partner Q&A in IC will catch it if the prompt doesn't.
- **Pre-rate-cycle vs. post-rate-cycle.** Comparing a Q4 2021 unitranche (SOFR + 5.50%) to a Q1 2026 unitranche (SOFR + 5.75% but on a fundamentally different base rate) without flagging the rate context is misleading. 24-month closing-date constraint addresses this.

- **Recent deals with limited post-funding data.** Closed last quarter — too early to show performance against base case. Acceptable for structural comp but should not anchor the leverage discussion.
- **Restructured or repaid deals.** Comps that have since amended (covenant relief, repricing) or repaid early. Disclose status alongside original terms; both data points matter.

Audit checklist

1. **Each comp's inclusion rationale is explicit and defensible.** Pull the comparability disclosure. If F2's rationale is "similar size and sector" without naming the specific GICS sub-industry or EBITDA range, that's insufficient. Push back to F2 engineering on retrieval logic.
2. **Outlier flags trigger correctly.** Compute the IQR for each tear-sheet row independently and verify F2's flags match. Particularly important on leverage and pricing — these are the two attributes IC will challenge.
3. **Date range is disclosed and matches the rate environment.** If F2 includes a Q4 2021 comp without flagging the rate context, the deployment hasn't internalized the rate-cycle constraint. Ticket to engineering.

Prompt 4 — Sensitivity pack — downside cases

Use case. Confirmatory diligence into IC. Need standard Meridian downside cases applied to the base-case model:

revenue decline, margin compression, working-capital stress, with covenant headroom under each. F2's strongest claim — 99% accuracy on core financial calcs — gets stress-tested here.

Why this prompt matters at deployment. Sensitivities are pure math with structured outputs. If F2 fails here, F2 fails on its strongest ground, and the deployment surfaces a fundamental issue early. If F2 succeeds, Meridian gains credibility on the headline accuracy claim and the platform earns the right to handle more judgment-heavy work in Tier 2.

Role: You are running standard downside sensitivities for a Meridian Credit Partners IC packet on a unitranche transaction.

Inputs:

- Output of Prompt 1 (Meridian-template P&L spread, base case)
- Output of Prompt 2 (validated Adjusted EBITDA after Meridian haircuts)
- Proposed deal terms: facility size, structure, pricing, amortization, maintenance covenants
- Meridian Sensitivity Library v3 (firm-standard stress cases)

Task: Apply each sensitivity to the base-case projections, recompute through the model, and output covenant headroom and breach analysis.

Standard sensitivities to run:

1. Base Case (no stress) — for reference
2. Downside 1: -15% revenue for 18 months, then 50% recovery to base
3. Downside 2: -200bps gross margin compression for 24 months

4. Downside 3: combined -10% revenue and -100bps GM compression for 12 months
5. Working Capital Stress: 30-day extension of DSO, 15-day compression of DPO
6. Stress Combo: Downside 1 + Working Capital Stress concurrent

For each sensitivity, compute and output:

- Revenue, Adjusted EBITDA, Adjusted EBITDA margin (each year of stress + recovery)
- Total Leverage (Total Debt / LTM Adj. EBITDA) at each test date
- Senior Leverage at each test date
- Interest Coverage Ratio (Adj. EBITDA / Cash Interest)
- Fixed Charge Coverage Ratio (per credit agreement definition)
- Free Cash Flow (defined as Adj. EBITDA - Cash Interest - Cash Taxes - Capex - Working Capital)
- Covenant headroom in % terms (Maintenance Cov / LTM EBITDA / actual ratio)
- BREACH FLAG (Boolean) at each test date for each covenant

Output structure:

- Master sensitivity table: rows are scenarios, columns are time periods, cells are the 8 metrics above
- Breach summary: list every (scenario × test date × covenant) breach, with % cushion deficit
- One-paragraph commentary: which sensitivity poses the greatest covenant risk, how many quarters from close does the first breach occur, what structural feature mitigates (if any)

Constraints:

- LTM EBITDA convention: 4-quarter trailing, calculated per credit agreement EBITDA definition
- Covenant test method: match credit agreement (LTM, not annualized, unless agreement specifies otherwise)
- Working capital stress applies only to operating WC, not financing
- Tax rate: hold flat at borrower's stated effective rate; do not stress
- If projections include sponsor's pro forma cost-save addbacks, run a parallel "no synergy" version of each sensitivity

Expected output shape. Master sensitivity table (6 scenarios × 5 forward years × 8 metrics each). Breach summary listing every covenant violation by scenario and test date. Commentary paragraph identifying the binding constraint.

Edge cases

- **Multi-year ramps with timing assumptions.** Borrower's base case assumes a 6-month sales-cycle elongation in year 2. Stress should compound on the already-stressed projection, not the original base case.
- **Revolver utilization in stress.** Working capital stress draws revolver; revolver draw increases interest expense; interest coverage tightens further. Prompt should compute the second-order effect, not hold revolver flat.
- **Non-recurring vs. recurring revenue mix.** SaaS borrower with 80% recurring revenue stresses differently than a project-based services business. Stress should apply asymmetrically if revenue mix is disclosed.
- **Covenant step-downs over time.** Some agreements step covenants tighter in years 3–4. Sensitivity must apply the right covenant level at each test date.
- **PIK toggles or springing covenants.** PIK-toggle features shift cash interest under stress; springing covenants only test if revolver is drawn above threshold. Both behaviors must be modeled, not flattened.

Audit checklist

1. **Base case math reconciles to Prompt 1 spread.** Pull base-case Year 1 EBITDA and Free Cash Flow, verify match with the underlying spread. Any delta = wrong starting point and every downstream sensitivity is invalid.
2. **Stress is applied to the right line items.** "Revenue down 15%" should flow through to gross profit at base-case GM%, not to EBITDA directly. Margin compression should hit GP, not SG&A. Spot-check on one sensitivity.
3. **Covenant test method matches the credit agreement.** If the agreement specifies LTM, F2 uses LTM. If it specifies annualized for the first year post-close, F2 honors that. Mismatch is a deployment failure with high partner-attack-vector risk.

Tier 2 — Days 31–60 prompts

Prompt 5 — Covenant design with cushion math

Use case. Term-sheet stage. Principal designing the financial covenant package for a new deal. Goal: propose maintenance covenants with appropriate cushion to base case, structured to give Meridian early warning without over-tightening.

Role: You are proposing a financial covenant package for a Meridian unitranche term sheet.

Inputs:

- Output of Prompt 1 (P&L spread, base case)
- Output of Prompt 4 (downside sensitivity outputs)
- Borrower's last 8 quarters of EBITDA (volatility input)
- Meridian Covenant Library – firm-standard covenant constructions by sector and EBITDA volatility band
- Proposed deal: total leverage, senior leverage, structure, hold

Task: Propose a financial covenant package with explicit cushion math against the base case.

Required covenants to design:

1. Maximum Total Leverage (LTM)
2. Minimum Fixed Charge Coverage (LTM)
3. Optional: Minimum Liquidity (if borrower has working capital seasonality or revolver-dependence)

For each covenant, output:

- Proposed Level at Close (e.g., 5.50x Total Leverage)
- Step-down schedule (if applicable) – quarterly tightening over Year 1 / 2
- Cushion to Base Case at Close: % cushion = (Covenant Level - Base Case Ratio) / Base Case Ratio
- Cushion to Downside 1 at first test date: same calc, vs. Downside scenario
- First-breach quarter under each downside scenario

Cushion targets (Meridian convention):

- Stable EBITDA business (last 8Q std dev <10% of mean): 25% cushion to base case
- Moderate volatility (10-20%): 30% cushion
- High volatility (>20%): 35-40% cushion + minimum liquidity covenant

Equity cure rights:

- Propose terms: cap on cures per fund life, cap per covenant period, treatment of cure dollars (EBITDA add vs. debt reduction), mulligan clauses
- Reference Meridian's standard cure construction unless deal-specific rationale to deviate

Output:

- Proposed covenant table with all numerical inputs and outputs
- One-paragraph rationale explaining cushion sizing relative to volatility band
- Flag any covenant where proposed level is tighter than Meridian's comparable closed deals (per Prompt 3 comp set) – this is a negotiation signal

Expected output shape. Covenant proposal table (3 covenants × close level, step-downs, cushion math, breach analysis). Rationale paragraph. Flag if proposal is tighter than firm precedent.

Edge cases

- **Volatile EBITDA businesses.** Standard 25% cushion is insufficient for a borrower with 30%+ EBITDA volatility. Prompt forces wider cushion plus liquidity floor.
- **Seasonal businesses.** Holiday-driven retailer with Q4 EBITDA = 60% of full year. LTM testing washes seasonality but quarterly testing destroys it. Prompt must honor LTM convention or propose an alternative (e.g., trailing 4Q rolling).
- **Capex-heavy growth phases.** Borrower planning \$20M of growth capex in Year 1. Standard FCC covenant should accommodate; prompt should propose a covenant holiday for capex above maintenance threshold.
- **Equity cure abuse.** Sponsor proposes 4 cures per fund life with no cap per covenant period. Prompt should flag deviation from Meridian standard (typically 2 cures total, max 1 per 4 quarters).

Audit checklist

1. **Cushion math references the correct base case.** The denominator in the cushion calc must match the EBITDA from Prompt 1's validated spread, not the sponsor's unvalidated number.
2. **Step-downs follow firm convention.** If Meridian's standard is 0.25x quarterly tightening over 4 quarters, F2's proposal matches. If F2 proposes faster step-downs without rationale, push back.
3. **Equity cure rights are explicit and bounded.** Cap, cap-per-period, cure dollar treatment, mulligan terms — all four must be in the proposal. Silent omission means the borrower's counsel writes the cure terms in the credit agreement and Meridian inherits whatever they propose.

Prompt 6 — Screening memo (24-hour 1-pager)

Use case. Steve's pain point #2 — partners want a 1-pager 24 hours after the teaser, not five days. Screening memo is short, fast, and explicitly preliminary. Different output entirely from the full IC memo.

Role: You are drafting a 24-hour screening memo for Meridian Credit Partners following receipt of a sponsor teaser.

Inputs:

- Sponsor teaser (1-2 page PDF)
- CIM (if provided at teaser stage; flag if absent)
- Borrower public information (if borrower-named in teaser): website, press, basic financial disclosures
- Meridian sector strategy notes (firm-internal, by sector)
- Output of Prompt 3 (comp screen) IF sufficient teaser-level data exists; otherwise skip

Task: Produce a one-page screening memo addressed to the firm's Investment Committee, structured for 5-minute partner read.

Output structure (single page, ~400 words):

HEADER:

- Borrower (or "Project [Codename]" if anonymized)
- Sponsor
- Sector / Sub-Industry
- Estimated Deal Size and Structure
- Date / Author / Confidence Level (HIGH/MEDIUM/LOW)

SECTION 1 – Initial Read (~80 words):

- One paragraph: business overview, sponsor thesis, why this comes to Meridian (relationship-driven, broad auction, club deal, etc.)

SECTION 2 – Three Diligence Priorities (~120 words):

- Three numbered priorities, each 1-2 sentences. Specific to this deal, not generic ("verify EBITDA quality" is generic; "verify whether the \$4.2M consulting adback represents a recurring spend masked as restructuring" is specific).

SECTION 3 – Comparability Read (~60 words):

- One paragraph referencing 1-2 closest comparable Meridian deals (from Prompt 3 if available). State expected leverage range and pricing range based on comp set. Explicit confidence flag if comps are weak.

SECTION 4 – Recommendation (~80 words):

- PURSUE / PASS / GATHER MORE INFO with one-paragraph rationale
- If PURSUE: name 2-3 confirmatory items needed before formal IOI
- If PASS: name the dispositive concern
- If GATHER MORE INFO: name the specific information needed

Constraints:

- Hard cap 1 page (Meridian Template v4 letter size, 11pt font, 1" margins)
- No appendices, no charts, no tables – pure prose
- Voice: direct, hedged where appropriate, no salesmanship
- Source links on every quantitative claim, even at screening stage

(Audit Mode default-on)

Expected output shape. One-page memo, ~400 words, four sections, header block, source links inline.

Edge cases

- **Missing financials in teaser.** Sponsor teaser typically gives revenue range and EBITDA, sometimes only revenue. Prompt should flag absence and default to LOW confidence on quantitative reads.
- **Sponsor unknown to firm.** First-time sponsor with no Meridian deal history. Prompt should call out this as a diligence priority, not gloss over it.
- **Sector outside firm experience.** Teaser is for an oil services rollup; Meridian doesn't do oil services. Recommendation should likely be PASS with the dispositive concern named explicitly.
- **Anonymized teasers.** Sponsor sends "Project Falcon" without naming the borrower. Comparability read becomes weaker; confidence flag should reflect this.
- **Returning sponsor / repeat borrower.** Sponsor Meridian has done 3 prior deals with. Prompt should reference firm history with this sponsor (last deal terms, performance) — turns a generic screen into a relationship-aware screen.

Audit checklist

1. **Initial read flags every assumption.** If F2 says "the business appears to have ~70% recurring revenue," the source must be cited. If the source is the teaser itself rather than verified, that's a confidence flag, not an assertion.
2. **Diligence priorities are specific to this deal.** Reviewer pulls the three priorities and asks: would these priorities be different for a different deal in the same sector? If no, F2 produced generic content. Push back.
3. **Recommendation includes confidence level.** PURSUE/PASS/GATHER doesn't stand alone — it pairs with HIGH/MEDIUM/LOW confidence and a named driver. Recommendation without confidence is a partner-trust failure waiting to happen.

Prompt 7 — IC memo sections 1–3

Use case. First F2-assisted IC memo at credit committee. Day-45 milestone. Generates Investment Thesis (§1), Use of Proceeds (§2), Sources & Uses (§3) using Meridian template v4. Senior Associate edits, Principal reviews before any partner sees output.

Why this prompt matters at deployment. This is the prompt where F2 either earns or loses partner trust. Failure here ends the deployment. The reason it shows up in Tier 2 and not Tier 1 is precisely because the failure mode is catastrophic — a flawed memo at IC kills the platform's credibility before it has a chance to demonstrate value elsewhere.

```
Role: You are drafting Sections 1-3 of an IC memo for Meridian Credit Partners, using firm Template v4. The Senior Associate will edit your output; the Principal will review before partner distribution. You are producing
```

scaffolding, not final language.

Inputs:

- Outputs of Prompts 1, 2, 3, 4 (spread, addbacks, comps, sensitivities)
- Term sheet (executed) or proposed term sheet
- Sponsor materials: CIM, management presentations, financial model, QoE
- Meridian Template v4 (firm-standard IC memo format, including required exhibits and source-citation conventions)
- Meridian Voice Guide v2 (firm tone reference: direct, hedged where warranted, no marketing language, no superlatives without quantitative support)

Task: Generate Sections 1, 2, and 3 of the IC memo.

SECTION 1 — INVESTMENT THESIS (~600 words, 3 sub-sections):

§1.1 Transaction Overview (~150 words)

- Borrower description in Meridian's voice (not sponsor's pitch language)
- Transaction type (LBO, refinancing, acquisition financing, dividend recap)
- Meridian's proposed role (sole lender, lead, club participant)
- Proposed facility, hold, and pricing summary

§1.2 Investment Thesis (~300 words)

- Three numbered investment merits, each 1 paragraph
- Each merit must tie to a quantitative claim (cite source)
- Voice: state the merit; do not editorialize
- Structure: [Merit Name] | [Quantitative Support] | [Why This Matters to a Direct Lender]

§1.3 Key Risks Preview (~150 words)

- Three top risks named explicitly (full risk section appears later in memo)
- Each paired with one-sentence mitigant
- This is a preview, not exhaustive

SECTION 2 — USE OF PROCEEDS (~200 words):

- Itemized table: each use, \$ amount, description
- Reconcile to total Sources of Funds
- If transaction includes refinancing, name the existing instruments being refinanced and their terms
- If transaction includes a dividend, state the dividend amount and source (sponsor recap distinguished from operating cash distribution)

SECTION 3 — SOURCES & USES (~150 words plus table):

Sources of Funds (table):

- Each source: instrument name, \$ amount, % of total cap structure, lead arranger
- Total must reconcile to Total Uses

Sources commentary (~100 words):

- Capital structure summary: total leverage, senior leverage, equity contribution %
- Comparison to comp set leverage (reference Prompt 3 output)
- Any structural notes: PIK toggles, prepayment penalties, MFN protection

Constraints (apply to all three sections):

- Source link on every quantitative claim. Zero exceptions.
- Voice: Meridian Template v4. Do not import sponsor pitch language. Do not use words like "compelling," "best-in-class," "industry-leading" without quantitative substantiation.
- Hedge language is preferred over false precision. "Approximately 40%" is better than "exactly 41.3%" if the underlying precision is not source-supported.
- Section 1 (Investment Thesis) must explicitly reflect Meridian's frame as a lender, not as an equity investor. Equity upside is irrelevant; downside protection is the frame.

Expected output shape. Three sections, ~950 words total, with Audit Mode source links on every quantitative claim.

Tables in Sections 2 and 3.

Edge cases

- **Sponsor pitch language bleeds into investment thesis.** Sponsor's CIM uses words like "transformational acquisition opportunity." F2 should not import this language into Section 1. Voice guide constraint addresses; reviewer must verify.
- **Sources & Uses with multiple closing scenarios.** Term sheet contemplates an "upsizing scenario" where sponsor commits more equity. S&U should show both scenarios side-by-side, not pick one silently.
- **Refinancing transactions.** Use of Proceeds must name existing debt being refinanced — not just "refinance existing debt." Counterparty, instrument, original maturity, current rate.
- **Dividend recaps.** Politically sensitive at IC. Section 2 should state dividend amount explicitly and not bury in "general corporate purposes."
- **Investment thesis written from equity lens.** Common failure mode for AI-drafted memos. Lender thesis is "this borrower can service this debt under reasonable downside cases." Equity thesis is "this borrower can grow EBITDA 3x." Don't mix.

Audit checklist

1. **Every quantitative claim has a source link.** Spot-check three claims at random across the three sections. Any un sourced number is a deployment failure with partner-attack-vector risk.
2. **Investment thesis voice matches Template v4.** Pull a recent partner-approved memo from firm history. Compare voice. F2 output should be indistinguishable in tone from the historical memo. Voice drift is the most common AI failure mode in IC contexts.
3. **Sources & Uses adds to 100% and capitalizes the right closing fees.** Standard test: do Sources equal Uses? Are arrangement fees, OID, and legal fees in the Uses table or being treated as off-balance-sheet items? Off-balance treatment without explicit rationale is wrong.

Prompt 8 — Risk and mitigants section drafting

Use case. Section 6 (or equivalent) of the IC memo. Risks ranked, mitigants tied to deal structure. Single most-at-tacked section in IC partner Q&A.

Role: You are drafting the Risk and Mitigants section of a Meridian IC memo. This section will be the most-scrutinized part of partner Q&A. Your job is not to minimize risk; it is to surface the right risks and pair each with a specific mitigant tied to deal structure.

Inputs:

- Outputs of Prompts 1, 2, 4, 5 (spread, addbacks, sensitivities, covenants)
- Diligence findings (data room review, management calls, customer references, expert calls)
- Term sheet
- Meridian's standard risk taxonomy (firm-internal: business risk, financial risk, structural risk, market risk, sponsor risk, ESG risk)

Task: Generate a ranked risk-and-mitigants table covering the top 5-7 risks. Do not list more than 7 risks. Forced prioritization is a feature.

Output structure (ranked table):

For each risk (5-7 rows):

- Risk Name
- Category (use exact taxonomy above)
- Severity (HIGH / MEDIUM / LOW)
- Likelihood (HIGH / MEDIUM / LOW)
- Risk Description (1-2 sentences, specific to this deal)
- Mitigant – Structural (covenant, collateral, equity contribution, reporting requirement)
- Mitigant – Diligence-based (validated by diligence finding X)
- Residual Risk after Mitigants (HIGH / MEDIUM / LOW)
- Reference: cite the doc, sensitivity, or finding that surfaced this risk

Ranking logic:

- Severity x Likelihood, with ties broken by Residual Risk
- Surface 1-2 HIGH severity / HIGH likelihood items at top, even if they threaten the deal thesis. Burying these is a fail mode.

Below the table, output:

- "Top 3 Concerns" – narrative paragraph (~100 words) summarizing the top three risks for the IC, in plain language
- "Deal-Breaker Watch List" – risks that would cause Meridian to pass IF they materialize during confirmatory diligence (1-3 items)

Constraints:

- No generic risks ("execution risk," "market risk") without deal-specific framing.
- Mitigants must reference a SPECIFIC covenant, collateral package element, reporting requirement, or diligence finding. "We will monitor" is not a mitigant.
- Do not editorialize the residual risk to make the deal look safer than the diligence supports. If Residual Risk is HIGH after mitigants, say so.
- Voice: direct, technical, no hedging language ("possibly," "potentially") unless the underlying probability is genuinely ambiguous.

Expected output shape. Ranked risk table with 7 columns × 5-7 rows. Below: Top 3 Concerns narrative paragraph, Deal-Breaker Watch List (1-3 items).

Edge cases

- **Pollyanna failure (hiding risks).** AI-drafted risks often understate severity. Prompt forces severity/likelihood/residual triple to make masking visible.
- **Generic risks without deal-specific framing.** "Customer concentration" is generic; "Customer concentration with top 3 customers representing 47% of revenue, two of which are non-contracted month-to-month" is specific. Prompt's no-generic-risks constraint addresses.
- **Mitigants disconnected from structure.** "Strong management team" is not a mitigant for financial risk. Prompt forces mitigant to reference a specific covenant, collateral element, reporting requirement, or finding.
- **Risk count creep.** Listing 12 risks signals indiscrimination. Hard cap at 7 forces prioritization, which is what IC actually wants.

Audit checklist

1. **Each risk has both severity AND likelihood.** Single-axis risk lists (just "high risk") are insufficient. F2 must produce both, and Residual Risk must reflect the mitigants applied.
 2. **Mitigants reference a specific structural feature.** Pull each mitigant. Does it name a covenant ("Fixed Charge Coverage covenant at 1.10x"), a collateral element ("First lien on all assets including IP"), or a reporting requirement ("Monthly KPI reporting with covenant test commentary")? If not, push back.
 3. **No more than 7 risks.** Counting check. If F2 produced 9 risks because there are "9 important things to mention," F2 hasn't internalized prioritization. This is a prompt-iteration ticket, not an engineering ticket.
-

Tier 3 — Days 61–90+ prompts

Prompt 9 — Portfolio monitoring quarterly review

Use case. Quarterly portfolio review for a single Fund II name. Compare actuals to base case, flag deviations >10%, draft commentary in firm voice for inclusion in quarterly LP letter.

Role: You are running quarterly portfolio monitoring on a Meridian Fund II borrower. Your output feeds the firm's quarterly LP letter and the internal risk-rating committee.

Inputs:

- Borrower's Q[X] reporting package (P&L, balance sheet, cash flow, KPI dashboard, management commentary)
- Original IC memo at funding (base case projections)
- Last 4 quarters of variance reports (trend context)
- Current credit agreement (covenant definitions, test methodology)
- Meridian Quarterly Voice Guide (firm-internal: factual, no hedge language, covenant-test results stated explicitly)

Task: Produce a quarterly variance report and one-paragraph commentary.

Output structure (table – Variance Report):

For each KPI (Revenue, Adj. EBITDA, Adj. EBITDA Margin, Free Cash Flow, Total Leverage, FCC, Liquidity):

- Q[X] Actual
- Q[X] Base Case (from IC memo at funding)
- Variance (\$ and %)
- Variance Trigger Flag (TRUE if absolute % deviation >10% OR if direction conflicts with prior 2 quarters' trend)
- Trend over last 4 quarters

Covenant test results (separate table):

- Each maintenance covenant: Required Level | Actual Level | % Cushion | Test Status (PASS / WATCH / BREACH)
- WATCH = within 10% of breach. Surface explicitly.

Commentary (~120 words, single paragraph):

- Lead with covenant status (PASS / WATCH / BREACH)
- One sentence on revenue trajectory vs. base case
- One sentence on EBITDA / margin trajectory vs. base case
- One sentence on liquidity / leverage status
- One sentence on any management-disclosed event (M&A, customer loss, refinancing) that affects forward read
- Closing sentence: implied risk-rating action (UPGRADE / HOLD / DOWNGRADE / WATCHLIST)

Voice constraints (per Meridian Quarterly Voice Guide):

- No hedge language ("appears to," "seems to") – borrower either is or isn't
- State covenant results explicitly, never "approximately compliant"
- If WATCH or BREACH on covenant, lead with that
- No editorializing on management quality

Expected output shape. Two tables (variance + covenant tests) and one ~120-word commentary paragraph in firm voice with implied risk-rating action.

Edge cases

- **Missing reporting from borrower.** Borrower late on Q3 reporting package. Prompt should flag missing data explicitly and not interpolate.
- **Quarter-end timing differences.** Borrower fiscal Q3 ends Sept 30; Meridian fiscal Q3 ends Dec 31. Variance report must align periods clearly.
- **Non-financial KPIs without thresholds.** NPS, customer count, pipeline coverage — these are in management dashboards but don't have clear breach thresholds. Surface as context, not as variance triggers.
- **Covenant amendments mid-quarter.** Amendment in October that loosened FCC from 1.20x to 1.10x. Test must use the amended level for Q4 onward, not the original.
- **Refinancing or amendment in progress.** Borrower negotiating amendment. Commentary should note the in-progress event without speculating on outcome.

Audit checklist

1. **Variance triggers are correctly calibrated.** Pull the variance report and verify any KPI flagged as triggered actually exceeds the 10% threshold. Borderline cases (11-12% variance) should trigger; clean misses (8% variance) should not.
2. **Covenant test method matches the credit agreement.** LTM vs. annualized, EBITDA definition, equity cure adjustments. Same audit logic as Prompt 4.
3. **Commentary tone matches firm convention.** Pull a partner-approved quarterly LP letter from firm history. Voice should match — no hedge language, covenant status leads, implied risk-rating action stated.

Prompt 10 — LP DDQ response generation

Use case. ILPA DDQ 2.0 received from a prospective Fund III LP. ~250 questions across 21 sections. Goal:

pre-populate quantitative sections from firm content library, leave qualitative sections for human authorship, flag low-confidence responses.

Role: You are pre-populating an ILPA DDQ 2.0 response for Meridian Credit Partners' Fund III fundraising. Your output is a starting draft for the IR team. Quantitative sections should be answered from firm sources with citations; qualitative sections should be flagged for human authorship.

Inputs:

- ILPA DDQ 2.0 (received format: PDF / Excel / Word / portal)
- Meridian Firm Content Library:
 - * Track record (closed deals, performance metrics, IRR/MOIC by vintage)
 - * Team biographies and tenure
 - * Operational policies (compliance, valuation, ESG, risk management)
 - * Prior DDQ responses (last 8 quarters)
 - * Audited financial statements
- Meridian Fund III Marketing Materials (PPM, executive summary)

Task: Generate draft responses to all 21 sections of ILPA DDQ 2.0.

Output structure (per question):

- Question ID (per ILPA DDQ taxonomy)
- Question Text
- Draft Response
- Source: cite firm doc + section, OR prior DDQ response, OR "REQUIRES HUMAN INPUT"
- Confidence: HIGH / MEDIUM / LOW
- Flag for IR Review: TRUE if MEDIUM or LOW confidence

Categorization rules:

QUANTITATIVE SECTIONS (Track Record, Investment Performance, Fund Economics):

- Pull from audited financials and deal-by-deal attribution
- Cite specific document and page
- HIGH confidence default; LOW confidence only if calculation method differs across sources

QUALITATIVE-OPERATIONAL SECTIONS (Compliance, Risk Management, Valuation, ESG):

- Pull from firm policy library + prior DDQ responses
- MEDIUM confidence default; reference any policy update since last DDQ
- If policy doesn't exist, output "REQUIRES HUMAN INPUT – no firm policy documented" and FLAG

QUALITATIVE-NARRATIVE SECTIONS (Investment Strategy, Team, Market Differentiation):

- Pull from PPM and Fund III marketing
- LOW confidence default – voice and emphasis matter, human authorship preferred
- Output draft language but flag for IR rewrite

NEW-ESG SECTIONS (where firm doesn't formally track):

- Output: "REQUIRES HUMAN INPUT – Meridian does not currently track [metric]. Recommend either: (a) implement tracking before Fund III final close, or (b) provide directional response acknowledging gap"
- Always FLAG

Format output to match LP's preferred submission format (PDF / Excel / Word / portal). If portal-based, output as Excel with question IDs in column A for manual upload.

Constraints:

- No fabricated specifics. If firm content library doesn't support a numerical claim, output "REQUIRES HUMAN INPUT" rather than estimate.
- Track-record data must reconcile to audited financials. If DDQ asks for net IRR by vintage and audited financials show gross IRR, flag the gap and require human resolution.
- If question asks about a discontinued strategy or wound-down fund, reference firm history accurately; do not omit.

Expected output shape. Question-by-question response document (~250 rows) with confidence flags and review flags. Format matches LP's submission preference.

Edge cases

- **New ESG questions firm doesn't formally track.** ILPA 2.0 added portfolio-company carbon footprint, board diversity policies, supplier diversity tracking. Many mid-market firms don't track these. Prompt's "REQUIRES HUMAN INPUT" output is the right behavior — better to flag than to fabricate.
- **Track record questions with discontinued strategies.** Meridian's Fund I had a co-investment sleeve that was wound down. DDQ asks about active strategies; full disclosure is required even for wound-down vehicles.
- **Fund III questions referencing Fund II realized vs. unrealized.** Fund II is 78% deployed and 22% realized. Net IRR including unrealized marks is materially different from net IRR on realized exits only. Prompt must distinguish; many DDQ responses get this wrong silently.
- **Format mismatch.** LP sends Excel template with proprietary column structure. F2 output must match the template exactly, not produce a Word document the IR team manually transposes.
- **Prior-DDQ inconsistencies.** Last quarter's DDQ said "Risk Committee meets monthly"; current draft says "Risk Committee meets bi-weekly." Either previous DDQ was wrong or policy changed. Either way, flag.

Audit checklist

1. **Every quantitative response cites a source document and page.** Pull 5 random quantitative responses; verify citation accuracy. One wrong citation in a DDQ context is a fund-credibility event with the LP.
2. **Low-confidence responses are flagged for human review.** Spot-check qualitative-narrative section responses. F2 should default these to LOW with FLAG=TRUE. If F2 marks them HIGH and ships without flag, the IR team will catch it but the prompt is mis-tuned.
3. **Format matches LP submission preference.** If LP requires Excel with their question IDs, F2 output is Excel with their question IDs. Format mismatch is the most-cited reason DDQ responses get rejected as "incomplete."

§4

Output Audit Rubric

How to read this section

The audit rubric is the deployment strategist's most-used tool. Every F2 output that touches an IC packet, an LP DDQ, or an internal credit decision passes through some version of this rubric before it ships — sometimes as a 30-second mental checklist, sometimes as a formal review for high-stakes outputs. The rubric is not an academic framework; it's a working artifact.

This section does three things:

1. Defines the five audit axes and what passing each looks like in practice
2. Walks through a deliberately flawed sample IC memo excerpt to demonstrate the rubric catching three specific failures
3. Shows how each caught failure decomposes into a prompt-side fix and an engineering ticket — the JD's "audit reports, flag issues to engineering" bullet rendered as operational practice

The five audit axes

Axis 1 — Source grounding

The standard. Every quantitative claim in the output traces to a specific cell in the source workbook, a specific page in a source PDF, or a specific entry in the firm's deal library. Audit Mode is the platform feature; source grounding is the deployment-strategist's verification that Audit Mode is producing complete and accurate citations.

What passing looks like. A reviewer can click any number in any output and land on the underlying source within one click. Citation includes document name, page or cell, and (for cell references) the named tab. Aggregations and computed values cite their input cells, not just an output cell.

What failure looks like.

- Number appears without citation
- Citation points to a document that doesn't exist or doesn't contain the cited number
- Citation is a vague "QoE report" without page/section
- Computed number cites the cell containing the computed result rather than the inputs that produced it

Why this axis is first. A source-grounding failure invalidates every downstream check. If the underlying number is wrong, the math being correct is irrelevant. Source grounding is the foundation; the other four axes assume it has passed.

Axis 2 — Calculation fidelity

The standard. Math is reproducible. Formulas are explicit. Conventions match the credit agreement and Meridian template. Aggregations sum correctly. Year-over-year and ratio calculations are internally consistent across the document.

What passing looks like. Reviewer pulls any computed metric, recomputes from the cited inputs, and gets the same answer to within rounding. Conventions are stated explicitly when they could be ambiguous (e.g., "LTM EBITDA per credit agreement §3.2 definition," not just "LTM EBITDA").

What failure looks like.

- Sums don't add up
- Ratios computed against wrong base (e.g., gross EBITDA instead of adjusted)
- LTM convention inconsistent across the document (LTM stated for one metric, annualized for another, no flag)
- Currency conversion or fiscal-year alignment errors
- Off-by-one period errors (Year 1 numbers labeled as TTM, etc.)

Why this axis is second. F2's most defensible marketed claim is 99% accuracy on core financial calcs. Calculation fidelity is the rubric's check on that claim. A failure here is a high-confidence engineering ticket — if F2 can't compute correctly, the deployment doesn't survive.

Axis 3 — Completeness

The standard. The output covers every required section of the firm template, includes every required exhibit, and addresses every standard analytical workstream. Missing sections are either explicitly marked as "not applicable" with rationale or completed.

What passing looks like. Reviewer pulls the template checklist (Meridian Template v4 has, for instance, 14 required sections in a full IC memo) and ticks off each one against the F2 output. Every box is ticked, or the omission is justified.

What failure looks like.

- Required section silently omitted
- Required exhibit (e.g., Sources & Uses table, sensitivity pack) missing
- "TBD" or placeholder text in shipped output
- Section header present but content thin (e.g., "Risks" section with two generic bullets when template requires 5-7 ranked risks with mitigants)

Why this axis sits in the middle. Completeness failures are easier to catch than the first two axes (you're checking presence/absence, not correctness) but they're also the failure mode most likely to embarrass the deployment. A partner pulling up an F2-drafted memo and finding the Sensitivities section missing is a credibility event; the partner doesn't need to verify any math to know the output is incomplete.

Axis 4 — Risk surfacing

The standard. The output proactively identifies concerns the reader needs to know about. It does not summarize what the borrower or sponsor said; it interrogates it. Adverse facts are surfaced with appropriate severity, not buried.

What passing looks like. A reviewer reading only the F2 output (without access to underlying documents) can identify the top three things that could go wrong with the deal. Risks are ranked by severity and likelihood, not just listed. Adverse trends in source data trigger explicit callouts.

What failure looks like.

- Output reads as borrower advocacy ("strong management team," "favorable market dynamics")
- Adverse data (declining margins, customer concentration, working capital deterioration) appears in the data but is not flagged in narrative
- Risk section lists generic risks without deal-specific framing
- Mitigants disconnected from structural features (vague "we will monitor" rather than specific covenant or collateral element)

Why this axis is fourth. Risk surfacing is the most judgment-heavy axis and the hardest to make pass/fail. It's also the axis where AI-drafted content most consistently fails — the Pollyanna failure mode is well-documented across language models. The deployment strategist's job is to push back hard here, including on outputs that pass the first three axes cleanly.

Axis 5 — Tone calibration

The standard. The output reads in Meridian's voice, not the sponsor's voice and not the AI's default voice.

Hedge language is used appropriately. Marketing language is absent. The output would be indistinguishable from a partner-approved memo from firm history if voice were the only attribute considered.

What passing looks like. Reviewer compares F2 output to a recent partner-approved memo on a similar deal type. Voice is consistent. No "compelling," "best-in-class," or "transformational" without quantitative substantiation. Hedging is calibrated to genuine uncertainty, not used as a stylistic default.

What failure looks like.

- Sponsor pitch language imported into investment thesis ("compelling acquisition opportunity")
- AI-default hedging applied uniformly ("the company appears to potentially have...")
- Marketing superlatives without quantitative support
- Inconsistent voice across sections (Section 1 reads as sponsor pitch, Section 6 reads as AI summary, neither reads as Meridian)

Why this axis is last. Tone is the most cosmetic axis but it's also the partner-trust axis. A memo that fails tone calibration but passes the first four can survive deployment with extensive editing. A memo that fails the first four cannot survive at any tone.

The deliberately flawed sample

Below is an excerpt from a fictional F2-drafted IC memo for "Project Cypress" — a \$35M unitranche financing for a vertical SaaS borrower. The excerpt covers the Investment Thesis (§1.2) and the leading paragraphs of the Financial Analysis (§4) sections. It is constructed to fail three audit axes specifically: source grounding, calculation fidelity, and risk surfacing.

The sample is realistic in length and density for a Tier-2-prompt output (Prompt 7, IC memo sections 1–3). Read it cold, then watch the rubric apply.

[SAMPLE — F2 DRAFT, PRE-AUDIT]

§1.2 Investment Thesis

Cypress represents a compelling lending opportunity in the vertical SaaS sector, supported by best-in-class unit economics and a sponsor with a strong track record of value creation. The company has demonstrated robust revenue growth, with TTM revenue of approximately \$48M reflecting 34% year-over-year growth from FY2024. Adjusted EBITDA of \$14.2M implies a 30% margin, which is industry-leading among comparable assets in our deal library.

The investment thesis rests on three pillars:

Predictable, contracted revenue base. Cypress generates approximately 78% of revenue from multi-year contracts with average remaining duration of 2.4 years, providing significant cash flow visibility for debt service. Net revenue retention of 112% indicates strong customer expansion dynamics.

Defensible market position. Cypress operates in the construction project management software vertical, where it competes with Procore and a number of smaller players. Management indicates Cypress has won over 60% of competitive deals in the past 12 months, reflecting product superiority and customer references.

Aligned sponsor with sector expertise. The sponsor, Marlin Equity Partners, has deployed over \$2B across vertical SaaS investments and has owned Cypress since 2022. Sponsor equity contribution of \$52M represents 38% of total capitalization, well above the 30% threshold Meridian typically requires.

§4 Financial Analysis (excerpt — opening paragraphs)

Cypress's financial profile supports the proposed unitranche structure. TTM revenue of \$48M and Adjusted EBITDA of \$14.2M produce total leverage of 4.1x and senior leverage of 4.1x at close, well within Meridian's risk tolerance

for sponsor-led SaaS transactions. Free cash flow conversion of approximately 85% (Adjusted EBITDA less capex less change in NWC, divided by Adjusted EBITDA) provides ample coverage for the proposed \$2.4M of annual cash interest.

Adjusted EBITDA includes \$3.1M of management addbacks, primarily related to one-time professional fees associated with the 2024 acquisition of a smaller competitor and run-rate cost savings from completed integration activities. Both addback categories are well-supported by data room evidence and consistent with Meridian's standard addback acceptance criteria.

The sensitivity pack (Exhibit 4-A) demonstrates resilient performance under standard downside scenarios. Under a 15% revenue decline scenario, adjusted EBITDA compresses to approximately \$11M, and total leverage tests at 5.2x — within the proposed maintenance covenant of 5.50x.

[END SAMPLE]

The audit, applied

The rubric runs in axis order. I'm catching three failures across three different axes.

What the audit doesn't catch

A note on the rubric's limits, because over-claiming is itself a deployment-strategist failure mode.

The rubric catches three categories of failure cleanly: missing citations, math errors, and judgment-shaped issues like risk-surfacing weakness. It does not catch:

1. **Subtle factual errors.** If F2 cites a real source that contains a real number, but the number is wrong because the source itself is wrong (e.g., a sponsor's CIM contains an error in TTM revenue and F2 faithfully reproduces it), the rubric will not catch it. Source grounding passes; calculation fidelity passes; the underlying error survives.

1. **Wrong-comp-set selection at the retrieval layer.** If F2 pulls a comp set from the deal library that includes deals an experienced underwriter would not consider truly comparable, the rubric won't catch it without a separate retrieval-quality check. Prompt 3's audit checklist addresses this at prompt level; the broader rubric does not.

1. **Strategic misreads.** If the deal is structurally fine but represents a bad strategic decision for Meridian (e.g., adds concentration risk to an already-overweighted sector), no quantitative or narrative axis surfaces this. Strategic read is partner-judgment work that no AI output should replace.

The audit rubric is a defensive tool. It catches the failures that would embarrass the deployment if they reached IC.

It is not a substitute for partner judgment, and the deployment strategist who frames it as such is overreaching.

How the rubric operates day-to-day

For Meridian specifically, here's how the rubric runs in practice:

Tier 1 outputs (spreading, comps, sensitivities, addbacks). Associates run the rubric themselves on every output, using the per-prompt audit checklists from Section 3. Failures are documented in a shared log. Pattern-failures (the same axis failing on 3+ outputs in a week) get escalated to me.

Tier 2 outputs (memo sections, covenant proposals, risk sections). The deployment strategist runs the full rubric on every output before it reaches a Principal. Principals run an abbreviated rubric (axes 4 and 5 — risk surfacing and tone) before partner distribution.

Tier 3 outputs (DDQ responses, monitoring commentary). IR team and Associates run rubric in parallel — IR for tone calibration and completeness, Associates for source grounding and calculation fidelity.

Engineering feedback loop. Every rubric failure produces a ticket in the format above. Tickets aggregate weekly. I run a 30-minute call with F2 engineering every other week to walk through the top 3-5 tickets, prioritize platform fixes, and confirm prompt-side workarounds where engineering fixes are out-of-scope or longer-timeline. The Voice of Customer bullet in the JD is exactly this loop.

This cadence is what turns the audit rubric from an abstract framework into the operational backbone of the deployment.

§5

30/60/90 Deployment Plan for Meridian

How to read this section

This is where Sections 1–4 turn into a calendar. The persona told us who the customer is. The workflow decomposition told us what to launch with and why. The prompt library is the technical content the deployment ships. The audit rubric is the operating framework that catches failures before they reach partners. The 30/60/90 is the timeline that ties all four into a sequence that produces a successful deployment by day 90.

Three principles govern the plan:

1. **Build trust before reaching for value.** Every workstream is sequenced to produce a visible win for the next-most-skeptical stakeholder. Day-30 wins target Associates. Day-60 wins target Principals. Day-90 wins target the Managing Partner and the LPs.
2. **Risk-management dominates day-1–30 design.** The first 30 days exist to demonstrate that F2 doesn't break in obvious, visible ways. Ambitious launches in this window are how deployments fail.
3. **Realistic uncertainty is named, not hidden.** Each phase identifies what could go wrong, what the mitigation is, and what the contingency is. A deployment plan that doesn't name its own failure modes is selling, not planning.

Pre-deployment week (Day -7 to Day 0)

Before day one, three workstreams have to be in place. None of these are billable F2 deployment work — they're customer-success preconditions that determine whether the deployment can even start.

Customer success handoff. Sales closed Meridian; the deployment strategist needs the full discovery package.

Specifically: the buying-motion narrative (who championed it, who pushed back, what concerns surfaced in the procurement cycle), the technical fit conversation notes, any contractual commitments F2 made about specific outcomes, the list of decision-makers and their concerns. This is a 2-hour internal meeting at F2, not a customer-facing event.

Stakeholder mapping. Before I introduce myself to Meridian, I need a stakeholder map naming every person I'll work with, their role, their concern, and their decision-making weight. For Meridian: Steve K. (MP, sponsor of the deployment, decision authority on continuation), the two sector-head Partners (passive in the buying motion, will

weigh in at day-90 QBR), three Principals (skeptical, audit-trail-focused, primary review layer for memo work), four Associates (will become champions if Tier 1 lands; will become resisters if it doesn't), two Analysts (most directly affected by Excel Intelligence, lowest political weight). Two non-investment FTE (CFO and IR/Fund Admin) come in at day 30+.

Technical preconditions. SSO requirements (Meridian uses Okta), data residency questions for the LP DDQ workstream, retention policy alignment to Fund III's LP-A obligations.

Days 1–14 — Foundation

The first two weeks are configuration, ingestion, and stakeholder calibration. No live deal work in this window.

Day 1 — Kickoff. Two-hour kickoff call with Meridian's full investment team. Steve introduces me to the team. I take 90 minutes of questions. The question I'm watching for: who in the room is asking risk questions? That person becomes my primary day-30 ally on the Principal layer.

Days 2–5 — SSO, retention, infrastructure. Configuration work. Okta SSO integration. Retention policy configuration to match Fund III's LP DDQ commitments (typically: 7-year retention on all deal-related materials). Data residency review.

Days 3–7 — IC template ingestion and voice calibration. Ingest Meridian Template v4 and Voice Guide v2. Generate test memo sections from old, anonymized Meridian deals, compare F2's output to the partner-approved historical memo, identify voice drift, iterate. Voice calibration is the single most under-budgeted workstream in AI deployments at investment firms. I budget 25 hours for voice calibration in days 3–7.

Days 7–10 — Deal library ingestion. Load Meridian's last 3-5 years of closed deals into F2's deal library. Each closed deal needs to be tagged with: GICS sub-industry, EBITDA at close, capital structure type, sponsor archetype, closing date, hold size, exit status. I do this with one Associate.

Days 10–14 — First Tier 1 prompt configuration. Configure Prompts 1, 3, and 4 (spreading, comps, sensitivities) for live use. Run each prompt against 2-3 historical deals. Acceptance rate target: 80% on Tier 1 prompts by end of day 14.

Day 14 milestone

- SSO live, retention policy confirmed, data residency cleared
- Voice calibration signed off by Steve + one Principal
- Deal library loaded and tagged: ~40-60 closed deals
- Tier 1 prompts (1, 3, 4) configured with 80%+ acceptance rate on historical test cases
- Internal champions identified: at least one Associate and one Principal

Days 15–30 — First live deals

Days 15–30 are the highest-risk window of the deployment. Two Associates run F2 on at least one live deal each.

Daily 15-minute standups. Every output passes through the Section 4 audit rubric.

Days 15–18 — Live deal #1 (Tier 1 only). The first live deal runs Prompt 1 (spreading), Prompt 2 (addback validation), Prompt 3 (comp screen), and Prompt 4 (sensitivities). The drafting Associate (Anna) runs the prompts. I run the audit rubric on every output before it reaches her Principal. The rubric runs in real time.

Days 18–25 — Live deal #2 + lessons-learned iteration. By day 18, I should have enough audit data from deal #1 to identify which prompts need iteration. I update Prompts 1-4 in a versioned change log (Meridian Prompt Library v1.1) and re-run the updated prompts on deal #2.

Days 25–30 — Day-30 milestone preparation. Materials: a short memo (3-4 pages) titled "Day-30 Deployment Status — Meridian." Content: adoption metrics, time-savings realized, output quality (acceptance rates by prompt, audit failure rates by axis), engineering tickets opened and resolution status, prompt iteration log, named risks for days 31–60.

Day 30 milestone

- Tier 1 prompts running on 2+ live deals with 80%+ acceptance rate
- 4 of 4 Associates have run F2 on at least one deal
- 2 of 3 Principals have reviewed F2 output and approved for IC packet inclusion
- Time-savings realized: 50%+ reduction on spreading vs. baseline, 75%+ on comp screening

- Audit failure rate trending: source grounding 5%, calculation fidelity 2%, completeness 10%
- At least 3 engineering tickets logged with F2 platform team, with at least 1 resolved
- Internal champions confirmed: 1 Associate, 1 Principal

Days 31–60 — Tier 2 launch and IC memo milestone

This is the phase where the deployment graduates from internal scaffolding to partner-facing output. The day-45 milestone — first F2-assisted IC memo at credit committee — is the single most important event of the entire 90-day deployment.

Days 31–37 — Onboard remaining investment professionals. Days 31-37 onboard the remaining 6 investment professionals (the third Principal, the two Partners, the Managing Partner, two Analysts). The two Partners and the MP don't run F2 themselves. They review F2 output.

Days 31–40 — Tier 2 prompt configuration and dry runs. Configure Prompts 5, 6, 7, 8 (covenant design, screening memo, IC memo §1-3, risk and mitigants). Voice calibration on Tier 2 is even more important than on Tier 1, because Tier 2 outputs are partner-facing. By day 40, all four Tier 2 prompts have a 70%+ acceptance rate on historical test cases.

Days 40–45 — First live IC memo with F2 assistance. Day-45 milestone: first F2-assisted IC memo at credit committee. The Senior Associate (Anna) drafts the memo using F2-generated §1-3 and §6. She edits all four sections heavily. Her Principal reviews the full memo. Three layers of audit: I run the full audit rubric on the F2 outputs at the Associate-edit stage and again after Principal review.

What success looks like at day 45. Partners have substantive Q&A on the deal, not on the memo's quality. Citations hold up under spot-checking. No partner asks "did F2 write this section?" in a tone that suggests skepticism.

Days 45–55 — Meridian-specific prompt library v1. The lessons from the first IC memo plus the cumulative iterations from days 1-45 produce **Meridian Prompt Library v1.0**. This is the document I leave behind if I'm hit by a bus on day 56.

Days 55–60 — First LP DDQ response. By day 55, Meridian has likely received at least one Fund III LP DDQ. I run Prompt 10 on it. Every quantitative claim has to be source-linked because LPs DO spot-check, especially during an active fundraising.

Day 60 milestone

- 12 of 12 investment professionals onboarded
- Tier 2 prompts (5, 6, 7, 8) live with 70%+ acceptance rate
- First F2-assisted IC memo successfully presented at credit committee
- Meridian Prompt Library v1.0 documented and shared
- First LP DDQ response delivered with F2 assistance
- Audit failure rates continuing to decline: source grounding 8%, calculation fidelity 8%, completeness 6%
- 5+ engineering tickets logged, at least 2 resolved

Days 61–90 — Expansion and QBR

The last 30 days are about three things: extending into Tier 3 workflows, completing the first quarterly cadence, and preparing for the day-90 QBR.

Days 61–70 — Portfolio monitoring launch. Fund II has 14 active names; each gets a Prompt 9 (portfolio monitoring quarterly review) run against its Q[X] reporting package. ~28 hours of Associate time across the portfolio (vs. 35-55 hours historical baseline). The monitoring run produces the firm's quarterly LP letter draft.

Days 70–80 — DDQ scaling and remaining LP responses. Multiple LP DDQs in flight. Volume target by day 80: 5-8 LP DDQs responded to with F2 assistance.

Days 80–88 — QBR preparation. QBR materials: metrics dashboard (deals processed, hours saved, MD review-cycle reduction, IC pushback rate, monitoring cycle time, DDQ turnaround time), audit failure summary (3-month roll-up by axis), engineering ticket log, prompt library v2.0, day-90 commitments for days 91-180.

Day 90 — QBR. 90-minute meeting. Steve, both Partners, three Principals, plus me. I present the dashboard in 30 minutes. Discussion in 60 minutes.

The decision the QBR drives: continued deployment authorization. Three possible outcomes:

Full expansion. Metrics meet or exceed targets. The deployment continues with expanded scope.

Continuation with constraints. Metrics are mixed. Specific workflows are working well; others are paused.

Pause. Metrics underperform materially. The deployment pauses pending fundamental rework. This is the failure outcome and it's named honestly.

Day 90 milestone

- All 10 prompts deployed with documented acceptance rates
- 4-6 IC memos completed with F2 assistance
- Q[X] portfolio monitoring cycle completed: 14 names reviewed in ~50% of historical time
- 5-8 LP DDQ responses delivered
- Total time saved: ~150-200 hours over 90 days, extrapolating to roughly 600-800 hours annualized
- Meridian Prompt Library v2.0 documented and operational
- Continued deployment authorization received

Risks and mitigations

Six named risks, ranked by likelihood × severity.

Risk 1 — Senior partner skepticism crystallizes around a specific failure

Likelihood: HIGH. Severity: CRITICAL.

The single most-likely deployment-killing event is a specific F2 failure that lands in front of Steve or one of the Partners and confirms a pre-existing skepticism.

Leading indicator. Audit failure rates trending up rather than down across days 14-30. Specific failures clustering around partner-visible workstreams.

Mitigation. Three layers of audit before any output reaches a partner. Aggressive prompt iteration on any failure mode that recurs more than twice. Voice-of-customer escalation to F2 engineering on any structural failure that prompt-side fixes can't address.

Contingency. If a partner-visible failure occurs, I escalate to Steve directly within 24 hours, present the root-cause analysis, the prompt-side fix, the engineering ticket, and the timeline to resolution. Honest accounting of what went wrong is the only viable response.

Risk 2 — Data leakage in DDQ contexts

Likelihood: MEDIUM. Severity: HIGH.

LP DDQs require firm-confidential information to be processed by F2. If F2's processing environment is not airtight to Meridian's standards, or if an LP discovers Meridian uses F2 in DDQ responses without disclosure, the firm's LP relationships are at risk.

Leading indicator. Any LP asks about Meridian's use of AI tools in operational responses. F2's data residency or processing-environment posture changes mid-deployment.

Mitigation. Pre-deployment data residency review. Explicit retention policy aligned to LP-A obligations. Conservative default on what data F2 ingests for DDQ workstreams.

Contingency. If an LP raises concerns, Meridian responds with the data-handling documentation prepared during pre-deployment, names F2's use explicitly, and offers to walk through the audit trail on any specific response.

Risk 3 — Model fatigue / Associate disengagement

Likelihood: MEDIUM. Severity: MEDIUM.

After 60-90 days, Associates may begin trusting F2 outputs without rigorous review. Audit failure rates plateau or rise.

Leading indicator. Audit failure rates flat or rising across days 45-75. Associate self-reports indicating they "trust F2 by now" on workflows where the rubric should still run.

Mitigation. Ongoing audit cadence (every output, no exceptions). Quarterly recalibration: at days 90, 180, 270, run a "rubric refresh."

Contingency. If model fatigue becomes a documented pattern, I run a 1-day "audit reset" workshop with all four Associates, walking through 5-10 specific failure cases to recalibrate vigilance.

Risk 4 — Voice drift across prompt iterations

Likelihood: MEDIUM. Severity: MEDIUM.

Each prompt iteration to fix a specific failure can introduce subtle voice drift elsewhere.

Leading indicator. Principals start editing F2-generated memo sections more heavily across days 60-90.

Mitigation. Voice regression testing built into the prompt iteration cycle. Every prompt update must produce output indistinguishable from the historical Template v4 baseline on a held-out test set of 3-5 historical memos.

Contingency. If voice drift surfaces, schedule a 2-day voice recalibration sprint with Steve and one Principal.

Risk 5 — Customer-side resource constraints

Likelihood: MEDIUM. Severity: MEDIUM.

The deployment requires meaningful customer-side time investment. If the firm's deal flow accelerates mid-deployment, customer-side capacity may evaporate.

Leading indicator. Standup attendance drops. Voice calibration reviews delay by more than 48 hours.

Mitigation. Front-load the highest-customer-time-cost work in days 1-14 when Meridian has full attention on the deployment.

Contingency. If customer-side capacity evaporates, I extend the timeline rather than compromising on quality.

Risk 6 — F2 platform-side stability or capability changes

Likelihood: LOW (in 90-day window). Severity: HIGH.

F2 is an early-stage company; product changes mid-deployment are possible.

Leading indicator. F2 internal communications about pending platform changes.

Mitigation. Direct relationship with F2 product/eng leadership. Early warning on breaking changes.

Contingency. If a breaking change ships mid-deployment, escalate to F2 leadership within 24 hours, propose a customer-side workaround for the affected window.

What this plan deliberately doesn't promise

I'm not promising the headline 45 hrs/week saved per analyst that F2 markets — that's an upper-bound number from F2's marketing materials, not a Meridian-specific projection.

I'm not promising IC pushback rate goes to zero on F2-assisted memos. Partners challenge memos. The right metric is "pushback focused on the deal substance, not the memo quality."

I'm not promising every workflow scales linearly to all sectors. SaaS deals will work better in Tier 1 than industrial deals.

I'm not promising no engineering tickets stay open at day 90. Naming them honestly in the QBR is better than pretending they're resolved.