


AI Field Guide

Agentic AI Transformation for Asset Managers

CrossCountry's perspective on where AI is changing Asset Management operations, and how to begin.



A note on why we built this.

Every conversation about AI in Asset Management seems to start in the same place. A firm has grown, the portfolio is more complex than it was three years ago, the investor base expects more and expects it faster, and the team running operations is roughly the size it has always been. Somewhere in the first hour, someone says a version of the same sentence: "We know AI can help, we just don't know where to start."

This guide is our answer to that sentence.

We built it because the honest answer to "where do we start" is not a strategy deck or a two-year roadmap. It is a specific piece of work, in a specific corner of your operations, that produces something real in a matter of weeks. For firm leadership, that matters too: executives start with use cases that drive material outcomes, the ROI from each funds what follows, and the program earns its mandate at every stage. The progression from a single agent to a full cluster forms the foundation of everything that follows.

That progression works because of who is behind it. The teams at CrossCountry who build these solutions come from inside Asset Management operations, they have closed NAV under audit pressure, rebuilt liquidity models against a deadline, and stood up governance programs from nothing. They know where the hard problems live because they have lived in them.

A field guide is meant to be used. This one documents twelve areas where AI can improve operations for Asset Management firms. Find the part that matches the problem in front of you. Each one shows where the work breaks down today, what you can put in place in the first week, and what the fully built-out version looks like.

Start with one. Once you have seen what it does, the path forward builds itself.

Our AI and Asset Management Leadership



Tom Alexander

Tom Alexander

Head of AI Innovation & Transformation



Bruce Klein

Bruce Klein

Co-Lead of Asset Management Practice



Sean Sinclair

Sean Sinclair

Co-Lead of Asset Management Practice



Anjali Khullar

Anjali Khullar


Partner, Asset Management

Four terms appear throughout this guide. Together they describe how every solution moves from a single agent to a full cluster.




Jumpstart Agent

The first agent deployed. It runs in weeks on your own data to establish a baseline and produce a concrete diagnostic output before you commit to the full build.



Week 1 Output

The tangible deliverable the Jumpstart produces early: a data quality scorecard, a variance memo, or a risk report. Evidence of value on your data before the full cluster is built.



Agentic Cluster

What the Jumpstart expands into: a coordinated group of agents automating a full operations domain. Every agent that follows draws from the same data foundation, and each cluster builds on what came before.



Human in the Loop

The control principle running through every stage. Qualified professionals review AI output at defined points, every action leaves an audit trail, and controls are built in before an agent goes live, so the work accelerates without handing over judgment or defensibility.

How a solution grows from a single agent to a full cluster

The control that governs it all

AI lets Asset Managers grow the portfolio without growing the back office.

Asset Management operations are scaling faster than the teams meant to handle them. Portfolios grow, fund structures multiply, and investors expect faster, more accurate reporting each cycle, while headcount holds flat. Most finance and operations teams are absorbing that pressure with the same manual workflows they have run for years, and the strain shows in close cycles that will not compress and exceptions that surface too late to act on. **AI changes that equation in four ways.**

1

AI absorbs the volume that used to require new headcount.

The repetitive, high-volume work that grows with every fund and every cycle, reconciliations, allocations, report production, moves to agents that scale with the portfolio instead of with the team.

2

Exceptions surface early, while there is still time to act.

Instead of explaining a break after the fact, teams see the outlier, the variance, or the covenant signal as it emerges, turning period-end firefighting into a standing, continuously monitored process.

3

The rebuild-from-scratch work stops repeating every cycle.

Normalization, data cleanup, and package assembly that teams redo from zero each period become a maintained foundation, so capacity goes to judgment and analysis rather than reconstruction.

4

Operations becomes a strategic asset to the firm.

As AI handles the operational load, the finance and data infrastructure that once consumed team capacity becomes a live resource for deal and operating teams to surface portfolio signals, model scenarios, and answer questions on demand rather than on a reporting lag.

**In a governed environment, AI earns trust only when control is built in from the start.
The next page shows how CrossCountry does exactly that.**

All AI output is governed before it drives a single decision.

Governance is where most Asset Managers want to begin, and for good reason: in a regulated environment, an AI output is only as useful as it is defensible. CrossCountry builds control into every agent as it is designed.

Each use case passes three consecutive questions before it goes live, and the answers determine the controls, testing, and evidence that agent carries.

What is the AI doing?

Every use case is sorted by what the agent actually does: ingesting data, transforming it, exercising judgment, or monitoring for exceptions. Each kind of work carries different risks and calls for different controls.

How much do you rely on it?

Some use cases let AI assist while a person validates the result. Others let AI output drive a decision directly, without a human re-performing the work. The second kind demands stronger controls, testing, and documentation than the first.


Can you stand behind it?

Every use case generates traceable evidence of what the AI was asked, what it produced, how it was reviewed, and what decision followed. The audit trail is designed into the workflow.

Governance is the first thing built into every agent, and it works alongside a set of other capabilities that every CrossCountry cluster shares. The next page shows what they are.

Every agentic cluster shares the same seven core capabilities.

Every agentic cluster CrossCountry builds shares the same core capabilities, whatever operational domain it automates. They determine how each agent takes in data, does its work, proves that work is right, and stays governed and accountable throughout. Together they are the foundation every cluster is built on.



These twelve clusters map to where Asset Management experiences the most strain in operations.

Each cluster automates a domain where manual workflows break down as the portfolio grows. They group into three areas of the operating model: the integrity of the numbers, the reporting and monitoring built on them, and the governance and controls that keep both defensible. Every cluster is built on the governance foundation on page 5.

Revenue, Valuation, and Data Integrity

- 1 Expense Allocation and Fee Analysis
- 2 NAV Calculation and Fund Economics
- 3 Data Foundation and Governance
- 4 Portfolio Valuation and Governance

Reporting, Performance, and Liquidity

- 5 Close and Consolidation
- 6 Reporting and Stakeholder Communications
- 7 Portfolio Performance Monitoring and Early Warning
- 8 Risk-Adjusted Liquidity and Capital Planning

Governance, Compliance, and Treasury

- 9 Operational Risks and Controls Monitoring
- 10 Investor and Legal Obligations Management
- 11 Regulatory and Compliance Operations
- 12 Treasury and Banking Operations

Every agentic cluster answers the same five questions, in the same order.

Every agentic cluster in this guide is documented the same way, so once you know the structure you can move through all twelve quickly. Here is how each page is built.

The template page is titled 'Agentic Cluster # [blank]' and 'Agentic Cluster Name [blank]'. It features five numbered sections:

- 1 Sound Familiar?** A section with a horizontal line for input.
- 2 What's breaking down** A section with three horizontal lines for input.
- 3 Where you start, and how it grows** A section with a blue header. It contains a diagram showing 'The Jumpstart Agent' (orange box) connected by a dashed orange arrow to 'The Full Agentic Cluster' (orange box). Below this, there are fields for 'Agent name' and 'Week 1 output', each with a horizontal line. To the right, there are three blue boxes labeled 'Agent Group #1', 'Agent Group #2', and 'Agent Group #3', each with a horizontal line.
- 4 What you get** A section with four horizontal lines for input.
- 5 How this cluster is governed** A section with four horizontal lines for input.

At the bottom left, it says '15 © CrossCountry Consulting'. At the bottom right, there is a small logo.

1 “Sound familiar?”

A phrase you might actually hear in your own organization. If it rings true, the rest of the page is for you.

2 “What's breaking down”

Start here. This frames the problem in your operations.

3 “Where you start, and how it grows”


The heart of every page: one Jumpstart agent, proving value in weeks, then expanding into the full cluster. If you read one section, read this one.

4 “What you get”

The outcomes you keep, framed as what changes for your team.

5 “How this cluster is governed”

The answer to page 5, applied to this cluster. Every page carries it, because governance is the foundation.



Agentic Cluster Detail

Agentic Clusters 1-4: Revenue, Valuation, and Data Integrity

Expense Allocation and Fee Analysis


Sound Familiar?

“Fee and expense allocations take days to reconcile by hand.”

What’s
Breaking
Down

When fee and expense allocations run through fragmented, manual workflows, they become difficult to reconcile, and the complexity of guarantee-fee tracking and offset calculations only deepens the problem. LPs are scrutinizing fee practices more closely each year, so an allocation decision without a documented basis behind it becomes a dispute waiting to happen, one that often sits unresolved until an LP advisory committee forces the issue.

Where you start, and how it grows



What you get

- Time back.** The reconciliation cycle between Corp Finance and Fund Finance now runs on its own, ending the manual disputes that used to consume the team.
- Fewer surprises.** G-fee offset exceptions and rule drift surface as they happen, so they can be corrected before they compound into misstatements or LP-visible errors.
- A standing basis.** Every allocation carries a documented methodology LPs can review line by line, so each cycle begins from a foundation instead of a reconstruction.

How this cluster is governed

- Documented methodology** preempts the LP fairness challenges that surface in advisory committee reviews and side-letter audits.
- Audit-ready summaries** give funds evidence to stand behind in an SEC exam, as fee-practice scrutiny intensifies.
- Continuous monitoring** flags G-fee offset exceptions and rule drift before they reach LP statements.



NAV Calculation and Fund Economics

Sound Familiar? “NAV still comes down to a spreadsheet and a late-night check.”

What’s Breaking Down

NAV still depends on spreadsheets and manual checks, so data issues surface late in the reporting cycle, when they are hardest to fix. Carry, waterfall, equalization, and clawback mechanics are difficult to calculate and harder to explain, and for retail and wealth vehicles, daily NAV production adds volume that manual processes cannot absorb. As fund structures grow more complex, LPs, investors, and auditors are all demanding more transparency at once, and late-cycle adjustments cascade into downstream reporting and audit findings that cannot be cleared without rebuilding the calculation itself.

Where you start, and how it grows

The Jumpstart Agent


NAV Input Validation and Movement Attribution Agent

You start with a single agent that validates NAV inputs from fund admin and portfolio systems, flags data gaps before calculation, and establishes a baseline with exception-based review for outlier valuations, FX movements, and accrual anomalies.

Week 1 Output: A validated NAV baseline by fund and investor, an exception inventory, and a draft movement attribution memo covering valuations, cash, FX, and fee accruals.

The Full Agentic Cluster

From that baseline, the cluster expands to automate the full NAV lifecycle, from data ingestion through fund economics modeling, fund admin oversight, and audit-ready attribution.



What you get

Accuracy you can defend. Every period-over-period NAV movement carries a documented driver, so carry, waterfall, and equalization calculations give auditors and LPs a basis they can review rather than a figure to challenge.

Fewer late surprises. Input validation and exception-based review catch data issues before calculation, ending the late-period adjustments that cascade into downstream reporting and disrupt fund close.

One record instead of a rebuild. LP re-up conversations, daily NAV for retail and wealth vehicles, and audit review all draw on the same validated NAV record, rather than one reconstructed and defended after the fact.

How this cluster is governed

Consistent methodology applied across every fund and investor eliminates the spreadsheet errors that drive NAV misstatement and allocation errors.

Shadow NAV reconciliation on fund administrator outputs ends blind reliance on third-party calculations, with vendor performance trends escalated for governance-level oversight.

Full movement attribution preempts the LP and auditor questions that surface when NAV changes cannot be explained during diligence or audit.



Data Foundation and Governance

Sound Familiar?

“Nobody agrees which system has the ‘right’ number.”

What’s Breaking Down

Disparate systems and inconsistent data definitions across funds and portfolios create a constant manual cleanup burden that consumes analyst capacity. Weak governance limits scalable analytics and AI adoption, and thin data lineage makes it hard to trace a reported figure back to its source. Every reporting cycle, every regulatory submission, and every AI initiative depends on data that cannot yet be trusted at scale, and that is the foundation the CDO and technology leaders are accountable for, built on systems the firm already owns but has not yet governed to the scale the business now requires.

Where you start, and how it grows

The Jumpstart Agent

Data Quality and Lineage Baseline Agent

You start with a single agent that profiles data across fund and portfolio source systems, including financials, tax, diligence, and performance, establishes initial field mappings and quality gates, and identifies the highest-impact data exceptions.

Week 1 Output: A data quality scorecard with lineage gaps, a prioritized exception inventory, and a governance roadmap with master data and access control requirements.

The Full Agentic Cluster

From that baseline, the cluster expands to establish trusted, governed data across the systems the firm already runs, with lineage and policy enforcement built in.

Agent Group #1
Ingestion and Normalization

Agent Group #2
Standardization and Mastering

Agent Group #3
Lineage, Governance, and Access

What you get

Trusted data, built once. Reusable, governed data across funds and portfolios eliminates the duplicate cleanup every analyst absorbs today and cuts the manual preparation behind reporting, regulatory submissions, and analytics.

A foundation for what comes next. Advanced analytics and AI adoption run on the systems the firm already owns, scaling without compounding governance risk, and reported results gain confidence through end-to-end lineage.

One standing source for audit, diligence, and AI. Governed, traceable extracts replace ad hoc data pulls, so audit, diligence, and downstream AI workloads draw on a standing source rather than a reconstruction every reporting cycle.

How this cluster is governed

Standardized field mappings and quality gates ensure every downstream analytics or AI workload runs on governed, traceable data rather than ad hoc extracts.

End-to-end lineage from source through transformation to reported output satisfies regulator and auditor expectations for traceability and prevents findings on data governance.

Role-based access and full usage logging protect LP PII, MNPI, and deal data across AI and analytics workloads, giving security leaders defensible control over who can access what.



Portfolio Valuation and Governance

What's
Breaking
Down

Portfolio valuation cycles run on inconsistent, non-standard PortCo financials and spreadsheet-driven workflows that require manual normalization every quarter. Valuation Committees, ICs, auditors, and LPs all expect clear, defensible explanations of what changed and why, and when those cannot be produced quickly, the result is audit rework, late-period adjustments, and IC and LP challenges. Even after a valuation is finalized, assembling the Valuation Committee, IC, LP, and audit presentation around it is a separate manual effort, rebuilt from scratch every cycle.

Sound Familiar? "Every valuation committee cycle, we're still chasing the same manual inputs from the same three people."

Where you start, and how it grows

The Jumpstart Agent


Valuation Input and Variance Baseline Agent

You start with a single agent that ingests valuation inputs, normalizes them to a consistent financial reporting structure, flags data quality exceptions before the cycle begins, rolls forward prior-period models with updated inputs, and generates a variance report with driver explanations.

Week 1 Output: A normalized PortCo input set, an exception inventory, and a draft quarter-over-quarter variance report with what-changed-and-why narrative.

The Full Agentic Cluster

From that baseline, the cluster expands to automate the full valuation lifecycle, from data intake through Valuation Committee sign-off, investor reporting, and audit defense.



What you get

Faster cycles, less rework. Eliminating manual normalization and quarterly rebuild compresses and stabilizes valuation cycles, and clear lineage from PortCo source data to final conclusions satisfies auditor and Valuation Committee scrutiny.

Fewer challenges to defend. Explainable valuation drivers and assumption changes reduce IC and LP challenge, while consistent methodology applied across the portfolio strengthens IC and Valuation Committee confidence.

Presentation-ready drafts. Valuation conclusions flow into draft Valuation Committee, IC, LP, and audit presentations, removing the manual reassembly every cycle, so the committee reviews and finalizes rather than rebuilds. Sign-off evidence is captured at every approval step.


How this cluster is governed

Source-to-conclusion reconciliation and consistent input normalization eliminate the methodology drift that produces valuation errors across periods and across the portfolio.

Embedded sign-off evidence and end-to-end audit trails address the documentation gaps that drive audit adjustments and restatement exposure during fund close.

Quarter-over-quarter variance attribution preempts the unexplained-valuation questions and manual reassembly that erode confidence in IC reviews and LP communications.





Agentic Cluster Detail

Agentic Clusters 5-8: Reporting, Performance, and Liquidity

Close and Consolidation

Sound Familiar?

“Close always slips a day or two when someone is out.”

What’s Breaking Down

Spreadsheet-driven close processes are error-prone and hard to scale across funds. Manual reconciliations and journal entries consume capacity, introduce risk, and produce unpredictable close cycles that delay every downstream report and decision. Persistent audit adjustments signal control gaps that cannot be cleared without redesigning the close itself.

Where you start, and how it grows

The Jumpstart Agent

Reconciliation and Exception Baseline Agent

You start with a single agent that matches transactions across subledgers and bank statements, flags exceptions by materiality, routes unmatched items to the close team, and surfaces the bottlenecks driving close-cycle length and audit adjustment exposure.

Week 1 Output: A reconciliation exception inventory by materiality, a close bottleneck heatmap by entity, and a prioritized automation roadmap.

The Full Agentic Cluster

From that baseline, the cluster expands to automate the close end-to-end, with embedded controls and audit-ready evidence at every step.

Agent Group #1

Reconciliation and Journal Entries

Agent Group #2

Workflow and Exception Management

Agent Group #3

Controls and Audit

What you get

Time back at period-end. Compressing and stabilizing the close lets finance teams spend period-end on analysis and decisions rather than chasing reconciliation exceptions, and surfaces issues before they cascade into downstream reporting.

Fewer audit surprises. Embedded audit-ready evidence and approval sign-offs throughout the close eliminate end-of-cycle reconstruction and reduce the findings and material weaknesses tied to spreadsheet-driven controls.

A close that scales. The model handles additional funds and entities without adding headcount in proportion, so the close no longer stretches every time the business grows.

How this cluster is governed

Reconciliation and pattern-based journal-entry drafting eliminate the manual touchpoints that introduce error into the close and produce restatement exposure.

Embedded sign-off evidence and consolidated audit packages address the control-documentation gaps that surface during external audit and create material weakness exposure.

Bottleneck escalation and checklist automation prevent the cascading delays that push downstream reporting and decisions behind schedule every period.



Reporting and Stakeholder Communications

Sound Familiar?

“We rebuild the same investor report from scratch every time someone asks.”

What’s
Breaking
Down

Ad hoc stakeholder requests routinely require rebuilding reports from scratch, because each audience's report is assembled by hand from separate source pulls with no shared data layer underneath. Board, IC, LP, management, compliance, ILPA, and marketing reports are built separately every cycle, so inconsistent data, formatting, and narratives across audiences erode credibility while heavy manual Excel and PowerPoint effort drains finance and IR capacity. The deeper problem is structural: when every audience has its own data pull and its own template, version conflicts and errors are inevitable no matter how carefully the team works.

Where you start, and how it grows

The Jumpstart Agent

Unified Reporting Data Set and Cross-Source Validation Agent

You start with a single agent that aggregates data and domain-specific content from fund accounting, portfolio, performance, value creation, and liquidity systems, validates cross-source consistency, and surfaces the version conflicts driving rework across reports.

Week 1 Output: A unified reporting data set, a cross-source consistency exception report quantifying every discrepancy across audiences, and a baseline narrative template library by audience.

The Full Agentic Cluster

From that baseline, the cluster expands to turn one validated data set into consistent, audience-tailored narratives delivered through a controlled, auditable workflow.

Agent Group #1

Data and Narrative
Unification

Agent Group #2

Drafting and
Validation

Agent Group #3

Workflow and
Distribution

What you get

One source, every audience. Consistent narratives across management, investor, and regulatory audiences come from a single source of truth, and LP and Board reporting lands on a compressed cycle without growing the team.

Fewer discrepancies to explain. Cross-source validation and source-to-report lineage cut the conflicting figures and regulatory omissions that surface when reports are built separately, so LPs stop flagging discrepancies before the fund can explain them.

No rebuild for every request. Ad hoc stakeholder requests are answered from the standing data set rather than rebuilt from scratch, and every communication carries a documented approval chain and version history back to its source.

How this cluster is governed

Cross-source validation against a unified data set eliminates the conflicting figures that surface across LP, Board, and regulatory reports built separately today.

Source-of-record validation and version control reduce the omission and transcription errors that produce regulatory findings on ILPA and SEC submissions.

Documented approval chains and version history give every stakeholder communication a traceable path back to source data, closing the gap between what different audiences see.



Portfolio Performance Monitoring and Early Warning

Sound Familiar? "By the time we see a performance issue in the report, it's already two weeks old."

What's
Breaking
Down

Portfolio performance is tracked in static dashboards with limited real-time visibility, so issues surface late or between reporting cycles. Without a systematic way to prioritize attention across PortCos or to link operational signals to valuation implications, escalation paths are reactive and ownership is unclear. Write-downs, covenant breaches, and liquidity surprises catch teams off guard, because monitoring and decision-making are disconnected: by the time a notification reaches the right person, the window to act has usually closed.

Where you start, and how it grows

The Jumpstart Agent

Portfolio Risk Score and Heatmap Baseline Agent

You start with a single agent that ingests management reporting packages, KPIs, covenant compliance data, and cash flow actuals from PortCo source systems, normalizes metrics across the portfolio for comparability, and detects trends and anomalies against plan.

Week 1 Output: A portfolio performance heatmap with composite risk scores, ranked intervention priorities, and a baseline early-warning trigger inventory.

The Full Agentic Cluster

From that baseline, the cluster expands to operationalize monitoring with linked alerting, valuation impact analysis, and ownership routing.

Agent Group #1

Ongoing Ingestion

Agent Group #2

Risk Scoring and
Detection

Agent Group #3

Alerting and Action

What you get

Problems caught between cycles. Monitoring between reporting cycles surfaces performance, liquidity, and operational risks earlier, and composite risk scoring ranks where to focus attention across the portfolio.

Fewer surprises at IC and for LPs. Linking performance signals directly to valuation implications, and triggering early-warning alerts to deal leads with clear drivers and next steps, reduces the write-downs, covenant breaches, and deterioration that would otherwise catch IC and LPs off guard.

Less reliance on manual review. Systematic anomaly detection replaces manual portfolio reviews, so issues are surfaced by the system rather than found by whoever happens to look.

How this cluster is governed

Regular KPI and covenant monitoring surfaces deterioration well ahead of quarterly reporting cycles, enabling structured intervention before issues become material.

Cross-portfolio correlation analysis identifies the systemic exposures that single-asset monitoring misses, addressing concentration dynamics that produce LP-visible drawdowns.

Covenant headroom and cash-burn monitoring trigger alerts before a breach, eliminating the late escalations that disrupt fund-level liquidity planning.



Risk-Adjusted Liquidity and Capital Planning

Sound Familiar?

“We model liquidity scenarios in spreadsheets and hope nothing moves before we finish.”

What's Breaking Down

Liquidity planning is reactive and scenario modeling is manual, which limits confidence in capital allocation decisions. Valuation changes, leverage, and exit timing are not consistently linked to forward-looking cash needs. Increasing fund complexity from continuation vehicles, evergreen funds, and private wealth structures heightens liquidity risk without scalable tooling, and stress-testing across multiple funds at once remains out of reach.

Where you start, and how it grows

The Jumpstart Agent

Forward-Looking Liquidity Scenario Baseline Agent

You start with a single agent that models forward-looking liquidity scenarios tied to current valuations, leverage covenants, and portfolio cash flow projections, then projects capital call and distribution timing based on deal pipeline status, fund terms, and historical deployment patterns.

Week 1 Output: A multi-fund liquidity scenario baseline, a capital call and distribution timing projection, and a starter sensitivity matrix for Treasury review.

The Full Agentic Cluster

From that baseline, the cluster expands to connect valuation, leverage, and operating assumptions into continuously refreshed liquidity intelligence.

Agent Group #1

Scenario Modeling

Agent Group #2

Capital Activity
Forecasting

Agent Group #3

Sensitivity and
Threshold Monitoring

What you get

Decisions without the manual rebuild. IC, Treasury, and Fund Leadership get scenario outputs they can act on directly, with capital allocation and liquidity decisions tied to current valuations and leverage positions rather than rebuilt by hand each time.

Constraints caught earlier. Liquidity constraints and funding risks surface earlier across funds and complex vehicle structures, and threshold-breach alerts arrive with driver attribution and recommended actions before a shortfall becomes urgent.

Stress-testing that finally scales. Liquidity can be stress-tested under adverse market and valuation scenarios across multiple funds at once, closing the gap between valuation assumptions, leverage, and forward-looking cash needs.


How this cluster is governed

Regularly refreshed scenario modeling tied to valuation and leverage signals identifies potential shortfalls early, eliminating the reactive scrambles around capital call timing.

Multi-fund sensitivity analysis quantifies commitment capacity under adverse scenarios, preventing the over-commitment that forces asset sales at unfavorable valuations.

Capital call timing optimization with pipeline and fund-term linkage reduces the poorly timed calls that strain LP relationships and increase default exposure.





Agentic Cluster Detail

Agentic Clusters 9-12: Governance, Compliance, and Treasury

Operational Risks and Controls Monitoring

Sound Familiar?

“We find control gaps during the audit, not before.”

What’s Breaking Down

Control breakdowns in fund and portfolio operations surface late, usually during quarter-end reporting, or an audit. Exceptions live in scattered spreadsheets, ownership of the fix is unclear, and the same issues recur because no one is accountable for closing them. The structural gap is that operations, finance, and compliance each watch their own process, so there is no consolidated view of exceptions against defined thresholds and no routing to a named owner. Problems get explained after the fact, once an LP or auditor raises them, rather than caught while there is still time to act.

Where you start, and how it grows

The Jumpstart Agent

Exception Detection and Scoring Agent

You start with a single agent that ingests operational, financial, and process signals from fund and portfolio workflows, detects exceptions against defined control thresholds, classifies each by severity, and applies consistent ownership-assignment rules to establish the routing logic for full cluster deployment.

Week 1 Output: An exception inventory with severity classification, ownership assignments, and a baseline remediation tracking framework.

The Full Agentic Cluster

From that baseline, the cluster expands to operationalize controls monitoring with embedded ownership, escalation, and audit-ready evidence.

Agent Group #1
Detection and Classification

Agent Group #2
Ownership and Remediation

Agent Group #3
Oversight and Audit Evidence

What you get

Exceptions caught as they happen. Control exceptions are caught when they occur rather than surfacing late in scattered spreadsheets.

Clear ownership and resolution. Every exception gets a named owner and a tracked path to resolution across funds and operations, replacing the period-end scrambles with a standing, monitored queue.

A documented control environment. Auditors see timestamps, actions, evidence, and outcomes on every exception, and third-party fund administrators are held to their SLAs with tracked performance and exception trends.

How this cluster is governed

Scheduled, threshold-based detection catches the control breakdowns that would otherwise surface unnoticed.

SLA and exception monitoring on the administrator relationship surfaces recurring performance issues at the governance level, where they can be raised before they affect downstream reporting.

A timestamped issue log with full evidence trails replaces after-the-fact explanations, giving auditors and LPs a documented record instead of assurances.



Investor and Legal Obligations Management

Sound Familiar?

“Side-letter and MFN terms live in someone’s head.”

What’s Breaking Down

Side letter terms, Most Favored Nation (MFN) elections, and other investor-specific obligations sit in scattered PDFs and spreadsheets, read and interpreted by hand each time they apply. A single LP’s terms can touch fee and carry calculations, reporting cadence, and distribution waterfalls, but nothing links the obligation to the processes it governs. Commitments are checked from memory, coordination between Legal, Finance, and Operations stays reactive, and a missed or misapplied term is found only when the affected LP raises it. That is both an LP relationship problem and a defensibility problem in front of auditors and regulators.

Where you start, and how it grows

The Jumpstart Agent

Obligations Register and Process Mapping Agent

You start with a single agent that parses side letters and LP agreements, extracts investor-specific terms across fees, MFN, reporting, and liquidity, structures them into a centralized register, and maps each obligation to the downstream processes it governs, including fee calculations, reporting cycles, and distribution waterfalls.

Week 1 Output: A structured obligations register with investor-specific terms, downstream process mapping, and ownership assignments.

The Full Agentic Cluster

From that baseline, the cluster expands to operationalize obligations enforcement across finance, reporting, and liquidity processes with ongoing monitoring.

Agent Group #1

Extraction and Structuring

Agent Group #2

Process Mapping and Ownership

Agent Group #3

Monitoring and Compliance

What you get

Every obligation in one place. Side letter and MFN terms live in one structured register with a named owner, so Legal, Finance, and Operations work from the same record of what each LP is owed instead of re-reading agreements each time.

Obligations checked before the period closes. Each obligation is checked against fees, reporting, liquidity, and disclosures before the period closes, and upcoming triggers are flagged ahead of the deadline rather than discovered after a miss.

Defensible in front of anyone. The fund can enter each period able to show LPs, auditors, and regulators exactly how obligations were met, backed by the register rather than memory.

How this cluster is governed

Automated obligation tracking against the register catches missed terms before they reach an LP statement or advisory committee review, eliminating the manual gaps that produce disputes.

Structured term extraction and pre-period checks catch inconsistent MFN and fee treatment before it reaches an LP statement.

Trigger-based alerts on reporting and distribution dates surface disclosure obligations before a deadline passes, not during an exam.



Regulatory and Compliance Operations

Sound Familiar?

“We find out we have a compliance issue when it’s flagged to us, not when it happens.”

What’s Breaking Down

The compliance program runs on spreadsheets, email, and manual checklists, so the CCO has no consolidated view of the firm's regulatory posture at any given moment. Filings like Form ADV, Form PF, 13F, and 13H are assembled by hand, marketing materials are reviewed case by case, and regulatory change is tracked informally. The common thread is that compliance evidence is rebuilt for each deadline rather than maintained on an ongoing basis, so gaps surface only when an examiner or a filing date forces them into view, and every SEC exam starts from scratch.

Where you start, and how it grows

The Jumpstart Agent

Regulatory Posture and Filing Baseline Agent

You start with a single agent that inventories current compliance program elements, populates and validates a baseline filing such as Form ADV or Form PF from fund data with cross-checks against prior submissions, and assesses marketing rule compliance posture across investor-facing materials to prioritize gaps for full cluster deployment.

Week 1 Output: A regulatory posture scorecard, a validated baseline filing draft, and a prioritized compliance program gap report with a remediation roadmap.

The Full Agentic Cluster

From that baseline, the cluster expands to operationalize the compliance function with testing, filing automation, and SEC exam-ready evidence.

Agent Group #1

Filings and Testing

Agent Group #2

Marketing and
Personal Trading

Agent Group #3

Change Management
and Evidence

What you get

Exam-ready between exams. Preparation becomes review rather than reconstruction, and the CCO gains a single, current view of the firm's regulatory posture instead of assembling it under deadline.

Faster, cleaner filings. Form ADV, PF, 13F, and 13H file faster and with fewer errors, each backed by an audit trail, and the Marketing Rule applies consistently across every investor-facing material and channel.

Regulatory change that lands on an owner. Regulatory changes are caught early and routed to an owner with an impact assessment attached, rather than tracked informally until a deadline forces attention.

How this cluster is governed

Scheduled testing and a standing evidence repository close the program gaps that turn a routine exam into an enforcement referral.

Filings populated from fund data and cross-checked against prior submissions cut the manual errors and missed dates on Form ADV, PF, 13F, and 13H.

Marketing Rule review of investor-facing materials flags non-compliant claims and performance figures before they are published, not after.



Treasury and Banking Operations

Sound Familiar?

"We do not know our true cash position until end of day."

What's
Breaking
Down

Cash sits in dozens of accounts across funds, SPVs, and blocker entities, and Treasury sees each one through a separate bank portal rather than a single position. Capital calls, distributions, management fee sweeps, and deal funding are authorized through manual, email-based approvals, which invites both error and fraud. As the entity count grows, this does not scale. The structural weakness is that Treasury activity and fund accounting run on separate tracks, so cash movements and the records that should support them drift apart, opening reconciliation gaps and segregation-of-duties exposure.

Where you start, and how it grows

The Jumpstart Agent

Consolidated Cash and Payment Control Baseline Agent

You start with a single agent that monitors bank accounts and cash positions across funds and entities, validates recent payment requests against budgets, contracts, and approval thresholds, and identifies dual-authorization gaps and unusual movement patterns to scope the full cluster.

Week 1 Output: A consolidated cash position dashboard, a payment control gap inventory, and a baseline fraud-risk pattern report.

The Full Agentic Cluster

From that baseline, the cluster expands to operationalize Treasury with automated controls, fraud detection, and accounting integration.

Agent Group #1

Visibility and
Position
Management

Agent Group #2

Payment Controls

Agent Group #3

Accounting Integration

What you get

One cash position, not many portals. Fund, SPV, and entity-level cash appears in a single position instead of many bank portals, and a daily cash position with exception detail reaches Treasury and Compliance.

Controls that catch fraud early. Payment authorization and segregation of duties move into enforced workflows off email, and unusual movements and dual-authorization gaps are flagged before a payment leaves the account.

Cash tied to the record behind it. Every cash movement ties to the fund event and accounting record behind it, and the model adds funds and entities without adding Treasury headcount in proportion.

How this cluster is governed

Enforced dual authorization and pattern-based detection close the manual approval gaps a single actor can exploit to move funds.

Standing segregation-of-duties monitoring and automated treasury-to-accounting reconciliation remove the blind spots that let misappropriation go unnoticed.

One consolidated cash position replaces the portal-by-portal view that hides shortfalls until a capital call or distribution is due.




Getting Started with CrossCountry

The first agent is running within weeks.

Firms typically start in one of two ways. Both lead to the same place:

- **Cluster-led:** Pick one place where the work is breaking down, prove value on your data, and expand from there.
- **Program-led:** Start with a leadership working session, align on the agentic roadmap on day one, and prioritize by material outcome.

Sample cluster-led approach



The people who built this guide have sat in your seat.

Every cluster in this guide was built by practitioners who have worked inside Asset Management operations, and by the engineers who build alongside them.

Closed NAV under audit pressure

Fund accountants who have produced and defended NAV when the auditors were in the room.

Rebuilt liquidity models on deadline

Finance leaders who have stood up capital and liquidity models when the clock was running.

Built governance from nothing

Compliance advisors who have designed and stood up governance programs from the ground up.

On every engagement, these practitioners work side by side with an engineer who turns their judgment into working agents on your data and your systems. That pairing is what you get when you build with CrossCountry.



Value in weeks

Proof on your own data before any full-cluster commitment.



Built to withstand scrutiny

Output that holds up to audit and LP diligence.



Human judgment in control

People stay in control of every decision.



Experts and engineers together

An engineer embeds alongside the practitioners to build and deploy on your data and systems.

Start with one cluster. Prove it on your own data. Build from there.

Ready to get started? Let's talk.

