

Build the Foundation First:

Why Data is What You Must Get Right Before Anything Else



IMPORTANT NOTICE

This document is being made available to you because your firm is a member of the Cutter Research consortia and has entered into a Member Agreement (the "Member Agreement") with Cutter Associates, LLC., Cutter Associates Europe Ltd or Cutter Associates Canada (collectively, "Cutter").

By accessing this document, you are confirming that you are authorized by your firm to access this document, that you are not an employee of a Prohibited Recipient (as defined below), and that you will comply with all of the terms and conditions of the Member Agreement.

In particular, you agree that:

- This document is the proprietary property of Cutter and contains information that is proprietary and confidential to Cutter and to the vendors reviewed and discussed herein (the "Vendors"). Disclosure of the information contained herein could cause irreparable harm to Cutter or the Vendors.
- You will safeguard this information with the same care as your firm affords its own confidential information. You will not provide access to this information to anyone who is not an employee of your firm and will not distribute this information outside your firm. Furthermore, you will not provide access to this information to any employee of any subsidiary, unit, department, or division of your firm that is engaged in any aspect of the software business involving third parties, including, without limitation, selling or otherwise providing software to third parties, assisting third parties in the selection or implementation of software, or providing investment related software information to third parties (collectively, the "Prohibited Recipients"). Excluded from this provision is information already known to your firm, information developed independently, information in the public domain, or information rightfully received from a third party.
- No part of this document may be reproduced or copied, in whole or in part, in any form, or by any means — manual, graphical, electronic, or mechanical, including photocopying, recording, taping, scanning, or information retrieval systems — without the express, prior, written permission of Cutter.
- Cutter has provided the information in this document in good faith. Many factors outside the control or current knowledge of Cutter affect the needs and project plans of members. Responsibility lies with the member to evaluate the information herein for applicability to the member's needs and circumstances. Cutter does not make any representation as to the accuracy, correctness, completeness, reliability, currency, or usefulness of the information in this document.
- In no event will Cutter be liable for any damages — direct, indirect, consequential, or incidental — arising from the use of this document.

Contents

Introduction	4
---------------------	----------

The Platform is Not the First Decision	5
---	----------

The data problem that is always there	5
---------------------------------------	---

The IBOR and ABOR question	6
----------------------------	---

What to do before selecting a platform	7
--	---

What a Well-Run Transformation Looks Like	9
--	----------

Change management is not a workstream. It is “the work.”	10
--	----

Scope as a discipline, not a constraint	10
---	----

What the Foundation Unlocks	12
------------------------------------	-----------

Real-time books of record	12
---------------------------	----

AI that works	12
---------------	----

The ability to grow without structural reinvention	13
--	----

Conclusion: The Sequence Is the Strategy	14
---	-----------

INTRODUCTION

There is a moment in almost every major investment management technology transformation when the program discovers what it should have found before it started. It arrives during data migration - six to twelve months in, after the solution has been selected, contracts signed, and internal teams committed. The discovery is not a technical failure. It is an organizational one, whereby the inherited data environment was never capable of delivering the outcomes the platform was purchased to achieve.



Position records spread across systems that were never reconciled to a single source.



An Investment Book of Record and an Accounting Book of Record that agree at month-end (because someone runs a manual reconciliation every night) but diverge at every other point in the trading day.



Pricing sources that are officially authoritative but practically overridden by individual desks.



Reference data maintained by more than one team, with no formal relationship between their versions.

None of these are unusual. They show up in most organizations through an accumulation of years of practical commitments made under real constraints. And while the system functions, it cannot withstand the scrutiny that a transformation brings to it.

Ridgeline's view is that data migration should not be treated as a one-time cleanup exercise at the back end of an implementation. It is a foundationally strategic opportunity for firms to decide whether they will preserve fragmented operating habits or standardize around a governed, repeatable data model.

To ensure a successful transformation, there are a series of data decisions that need to precede any platform investment. This paper explores what changes, across the full lifecycle of a transformation, when those decisions have been made.

The Platform is Not the First Decision

The data problem that is always there

The emergence of enterprise AI has accelerated a reckoning. For many investment managers, the first serious conversation about data architecture is no longer being driven by a system replacement or regulatory initiative, but by a simpler question from executive leadership:

Are we ready to deploy AI?


The answer depends far less on model selection than on data readiness. AI systems amplify the quality of the information they are given. Where data is fragmented, inconsistently defined, or governed differently across functions, AI does not resolve those inconsistencies; it operationalizes them at scale. An investment manager cannot expect an intelligent assistant, automated workflow, or predictive model to produce trustworthy outcomes when **the organization itself has not established consensus on a source of truth.**

These implications have re-shaped data evaluation. What was once viewed as an operational infrastructure decision has become a tactical prerequisite for enterprise AI. Firms increasingly recognize that they cannot modernize their technology stack, or realize the value of artificial intelligence, without first establishing a governed, unified, and trusted data foundation.

While AI didn't create the need for better data architecture, it has made the cost of avoiding it impossible to ignore.

Every investment management firm that has grown organically across multiple systems has, beneath its operating surface, a version of the same problem. Data that is abundant but not agreed upon. Positions that reconcile at month-end but diverge intraday. Reference data maintained by more than one team with no formal relationship between their versions. Pricing sources that are officially authoritative and practically overridden. Manual processes that compensate for gaps so reliably that the gaps themselves have stopped being visible.

None of this is negligence. It is the accumulated result of practical decisions made under real constraints: a risk team that built its own data feed because the official source introduced unacceptable latency or an operations team running a daily reconciliation because two systems were never brought into alignment. Each workaround addressed a genuine gap. Collectively, they form a data environment that can function under stable conditions but fractures under transformation. The moment the target state requires a single authoritative source, the organization is forced to answer a governance question it has deferred for years.



Every manual reconciliation process is evidence of a data governance decision that was never made. When the platform arrives, all those deferred decisions arrive with it.

The IBOR and ABOR question

Nowhere do these unresolved governance decisions become more operationally expensive than in the relationship between the Investment Book of Record (IBOR) and the Accounting Book of Record (ABOR). These two views of the portfolio are often produced separately, maintained by different teams, and reconciled at end-of-day. The process has become so embedded in the operating model that its true cost - in headcount, latency, and the constraints it places on future change - is no longer recognizable.

That reconciliation cycle is not a neutral operational cost. A real-time IBOR - one reflecting actual portfolio positions at any point during the trading day, not positions as of the last reconciliation - is the prerequisite for pre-trade compliance, intraday risk management, and any workflow that needs to operate on live data. If IBOR and ABOR are separately maintained systems, real-time is a qualified claim. It means recent or since-the-last-batch-ran, not truly real time. That qualification has consequences that grow more expensive as firms pursue AI-assisted operations and intraday risk capabilities.

From Ridgeline's point of view, the path to real-time operations isn't through reconciling faster, as that's addressing a symptom and not the problem. Firms need to eliminate the structural separation that made reconciliation necessary in the first place. By building a program on a single underlying data model, IBOR and ABOR become different lenses on the same continuously updated record rather than outputs of separate processes to be reconciled.

Before that architectural benefit can materialize, a number of governance decisions must be made. Whose transaction record is authoritative when systems disagree? Who resolves a pricing discrepancy? Which entity's version of a position prevails when an allocation spans legal entities? These are not technical questions. They are organizational ones, and the time to answer them is before the implementation clock starts.

What to do before selecting a platform

To move through transformation cleanly, the following should be treated as prerequisites, not parallel workstreams:



An honest data inventory

Not what systems say you have, but what every team trusts and acts on. For each critical domain (positions, prices, reference data, corporate actions, client records), ask yourself, "Is there a single version every stakeholder would accept as authoritative?" If the answer is, "No," that is a governance decision to be made before vendor selection, not during migration.




Explicit outcome ownership

Designating who is accountable for data accuracy as an outcome, not just for running the process that produces it. When something breaks post-go-live, process ownership provides no continuity. Outcome ownership does.



A migration scope grounded in reality

Sized against the actual state of the data environment, not optimistic assumptions about what the implementation team will find. Data migration is consistently the most underestimated workstream in any transformation. Its true scope can only be known after the inventory has been done honestly.



Ridgeline believes the implementation approach matters as much as the target architecture. Customizing around every exception slows firms down and adds unnecessary layers of complexity. Firms should focus on standardizing data definitions early, validating against a clear target model, and using repeatable transformation logic throughout the program. In practice, that means designing the implementation around patterns, staging and validation, and clear sign-off on authoritative data rather than bespoke, one-off conversion work.



Migration quality is often determined long before load day. Scope clarity, source-system understanding, and disciplined validation are what create confidence at cutover.

When firms clean up legacy groups, duplicate records, and unnecessary custom fields before migration, they are not just preparing data. They are simplifying their future operating model.


One firm significantly reduced validation churn by assigning a small internal working group to own source extracts end to end. Instead of sending different versions from different teams, they established a single repeatable export process, preserved full data precision, and reran the same files each cycle. That discipline helped the project team focus on true data issues rather than file-to-file inconsistencies.

What a Well-Run Transformation Looks Like

When the foundational work has been done, the character of an implementation changes materially. The program still carries complexity - any transformation of a live investment operation does - but it carries it differently. The team is not making governance decisions under time pressure. The migration scope is not expanding as previously undiscovered data problems surface. The go-live date is not hostage to a reconciliation that will not close.

Two things consistently determine whether an implementation moves cleanly or grinds:

1 How well the organization has prepared its data foundation.



2 How seriously has it treated change management as organizational design rather than a communications exercise.



The second is almost always underestimated, because it is less visible than the technology work and difficult to scope into a project plan.



Clean data does more than accelerate implementation. It reduces downstream reconciliation effort, reporting exceptions, and user confusion after go-live.

The fastest migrations usually start with the clearest source data, the most repeatable extracts, and the fewest manual touchpoints.

Change management is not a workstream. It is “the work.”

The most common mistake in an operational transformation is treating change management as a support function for the technology deployment. The output of this is typically a series of training sessions and communications scheduled around go-live. This fundamentally misreads what is being asked of the organization. A unified platform does not just replace a system. It restructures who is responsible for data, who has visibility into positions, who can act on a compliance alert, and who owns an exception when it crosses team boundaries. Each of those changes requires an organizational decision, and the decisions need to come before the program goes live, not after it surfaces the gaps.

In investment operations specifically, oftentimes the first week after go-live is diagnostic. Every unresolved question about data ownership, every process that was manually compensating for a structural gap, every informal agreement between the front and middle office about how a discrepancy gets handled - all of it surfaces at once, under the pressure of live trading. The organizations that absorb that pressure cleanly are those that have already answered the structural questions: which workflows belong centrally and which remain distributed, who adjudicates a disputed position record, and what the escalation path is when something breaks across team lines. The ones that have not are the ones whose go-live becomes a governance crisis wearing a technology costume.

One of the clearest markers of success is when customers approach migration as an operating-model exercise, not just a data handoff. In the strongest projects, firms designate decision-makers for data definitions, validation sign-off, and exception handling early, which shortens turnaround time and reduces ambiguity throughout the implementation.


Scope as a discipline, not a constraint

Scope is where programs are won or lost, and the pressure to expand it during delivery is relentless. Every team has requirements that did not make the initial specification. Every edge case generates a customization request. Every integration point has nuances the vendor did not anticipate.

To keep timelines and relationships intact through implementation, organizations must treat scope as a strategic discipline: clear about what the first phase must deliver, deliberate about what gets deferred rather than squeezed in, and honest about the difference between requirements necessary for go-live and requirements that are merely comfortable to have.

This discipline is easier to maintain when the data foundation has been established first. When governance is clear and migration scope is known, the implementation team can make scope decisions against a stable baseline rather than a moving one. The program is not simultaneously trying to transform the technology and resolve the data problems the technology exposed.

The programs that succeed are not those with fewer problems. They are those where the organization knows what it is solving for at every stage - and has the governance in place to make decisions when the unexpected arrives.



In one successful migration, the customer resisted the urge to recreate every legacy group and portfolio structure exactly as it existed. They reviewed which groupings were still actively used, retired outdated structures, and simplified exceptions in advance. That reduced noise during testing and made the future-state operating model easier for end users to understand.

What the Foundation Unlocks

Getting the data foundation right is not just about avoiding failure. The more consequential question is what becomes possible once it is done - capabilities most firms treat as aspirational that become operational when the underlying conditions support them.

Real-time books of record

The most immediate operational change for firms that implement on a unified data model is the elimination of the end-of-day reconciliation cycle between IBOR and ABOR. When both are native expressions of the same continuously updated record, position accuracy is a property of the architecture rather than the output of a daily process. That enables the headcount that previously ran that cycle to be redirected toward work that requires judgment. Firms are now able to prioritize exception investigation, client service, and analysis that was always deferred because the reconciliation had to run first.

Pre-trade compliance and intraday risk monitoring, previously dependent on the last batch having run, become real-time. Portfolio managers and risk teams work from the same position data simultaneously, without the latency that made real-time a qualified claim. This structural change provides a tremendous improvement to what investment operations can support, which is critical for firms expanding into new instrument types, multi-asset mandates, or derivatives.

AI that works

AI-assisted operations - in reconciliation, post-trade processing, client reporting, and compliance monitoring - are only as reliable as the data they operate on. This is the constraint that makes the sequence matter so much. Firms that invest in AI capabilities before establishing a clean, governed data foundation tend to find that automation accelerates their existing problems rather than resolving them. While the workflows may run faster, errors propagate further and exceptions multiply rather than shrink.

Firms with a sound data foundation have a very different experience. Their AI applications reduce manual intervention, surface exceptions worth investigating rather than noise worth ignoring, and enable them to redirect skilled operations staff from data maintenance to judgment work. That outcome is not a function of which AI tools a firm deploys. It is a function of what the data environment makes reliable enough to automate.

Ridgeline's view is that AI can't transform how your firm functions if it is layered on top of operational fragmentation. Data, workflows, and ownership first have to be brought into the same governed system. This provides AI with the necessary context and domain awareness needed to make automation trustworthy.

The ability to grow without structural reinvention

The most underappreciated benefit of a sound data foundation is what it does to the organization's relationship with change. In fragmented environments, every material shift - a new asset class, a regulatory change, a new counterparty structure - demands disproportionate effort. Because the environment was never designed to absorb change, each new requirement arrives as a custom engineering problem rather than a straightforward extension of what is in place.

Firms that have invested in a unified, governed data foundation absorb change differently. The new asset class connects to the existing data model rather than requiring a parallel one. The new regulatory requirement plugs into governance structures that are already functioning. The organization's capacity for change scales with its ambition rather than fighting it. This is particularly relevant for investment managers who are actively expanding their instrument coverage - firms adding swaps, private credit, or multi-jurisdiction structures to portfolios that were previously simpler find that the quality of the data foundation underneath them determines whether that expansion is operationally manageable or operationally consuming.

The Sequence Is the Strategy

Most technology programs in investment management are evaluated on delivery: did it go live on time, on budget, with the features specified. That is a reasonable measure of project execution and a poor measure of whether anything changed. To extract the most from major technology investments, firms need to ask the harder questions before the implementation begins and build the conditions under which the technology can do what it was purchased to do.

The firms that close that gap share a common discipline. They treat the data foundation as a first-order decision, not a technical prerequisite. They design governance before selecting systems. They resource change management with the same seriousness as the technical deployment. And they measure success not by go-live date but by whether the operation that emerges can do things the prior one could not.

The platform decision still matters. Architecture, data model design, upgrade continuity, and the breadth of instrument coverage all influence a firm's ability to adapt as markets evolve and investment strategies grow more complex. These are meaningful differentiators as well as selecting the right partner has a lasting, immutable impact. But platform quality is a ceiling, not a floor. The foundation is the floor. Build the right foundation first. Technology accelerates what already exists yet it cannot compensate for what doesn't.



Build the foundation first.
The platform determines what
is possible. The foundation
determines what happens.

BUILD THE FOUNDATION FIRST:
WHY DATA IS WHAT YOU MUST GET
RIGHT BEFORE ANYTHING ELSE

