

Why Are We Still Building Investment Performance Systems?

Mark R. David, CFA

Essex River Analytics

January, 2006

One of the first projects that Essex River Analytics undertook, back in 1996, was a performance attribution solution for a global equity manager. We began by searching for and interviewing several commercial vendors of attribution systems, and quickly concluded that none of them offered quite the sophistication that our client needed. Several were close, providing a basic Brinson-style attribution with variants, but the most common shortcomings, at that early date, included:

- Monthly frequency of calculation, with no daily capabilities. Most adopted a buy-and-hold method over the month, sometimes ameliorated with a Dietz or modified Dietz algorithm.
- Inflexible bucketing capabilities, limiting the analysis to pre-defined, static sector definitions. Some required holdings to be already aggregated before loading into the attribution.
- Large residuals relative to reported return as produced by the accounting system, with weak accounting for transactional flows, at best.
- Lack of robust multi-currency functionality.

Nonetheless, the future looked promising for the commercialization of performance attribution as an off-the-shelf product. The systems we looked at had not been on the market very long yet, and they were *close* to what we were looking for. Portfolio accounting systems were already robustly productized, and trading/order management systems were just becoming very competitive. Our forecast in 1996 was that - within five years - solutions in the investment performance domain would entail vendor search and integration projects only, never custom builds.

Cut to ten years later, and we're architecting and building more custom performance measurement and attribution solutions than ever before. Where did our prediction go awry? In this white paper, we will relate the observations and experiences that now lead us to the following conclusion:

The current demand for custom investment performance system development stems not from any failure to implement standard functionalities, but instead from rapid structural changes in the investment markets themselves. These changes have vastly outpaced, even transformed, the business requirements of a performance solution. The fundamental nature of this transformation makes it extremely expensive to modify the software that's already been developed in this space, thereby obliterating the capital investment and value of current in-house and off-the-shelf solutions and promising a strong future for custom solutions providers.

The Usual Suspects: Vendors *have* been keeping up with many of the requirements and developments that made news over the last decade.

The menu of commercially available solutions for performance measurement and attribution has improved dramatically in ten years, in terms of both number and quality of choices. There are today dozens of vendors competing in this arena, propounding a wide variety of methodologies supported by very robust systems. Many of these are well-integrated with accounting and trading systems from the same or partnered vendors, while others are stand-alone performance packages. Furthermore, the shortcomings we experienced in 1996 have all been addressed quite thoroughly:

- Almost all offerings offer both a monthly and daily calculation capability.
- Flexible bucket (or sector, or segment) hierarchy definition, conforming to any desired client analysis, is the standard today. Practically any attribute¹, static or dynamic, can be used to define security aggregates. These are aggregated by the vendor system, from basic position, trade, benchmark and security data interfaces.
- Though not a universal feature, several vendors offer a menu robust intra-period transaction processing capabilities that minimize residual vs. reported return at the portfolio level.
- Full multi-currency capabilities are so ubiquitous today as to be passé.

Another substantial development has been the introduction and widespread adoption of standards for performance presentation within the institutional investment management community. Beginning with the AIMR PPS, and closely followed by the GIPS, these standards immediately presented both a challenge and an opportunity to the vendors - one which most of them met with admirable speed and reliability².

Based on the weight of these significant advancements (and the capital expended to accomplish them), we can see that vendors have indeed been quite busy and can't be faulted for sloth or inattention. What, then, are the root causes of their inability to satisfy current business requirements in the investment performance domain?

Swaps in the Machine: The jump in the volume and liquidity of multi-legged derivatives has confounded the assumptions built into systems' architectures.

The last couple of years have experienced a phenomenal growth in swap contract volume. Investment managers use them routinely as a tool for risk management, while complex hedge fund strategies employ them as a means of creating exposure and leverage. Interest rate swaps, while available for some time, continue to grow in acceptability within institutional portfolios. Credit default swaps have undergone an enormous surge in popularity in the fixed income domain, and total return swaps are becoming more and more commonplace by the month. While swaps aren't in any way new, they're

¹ This statement is true so long as that attribute is the attribute of a *security*, a distinction that we will expand upon later on.

² It is worth noting that software can support GIPS standards, but is not in itself certifiably "GIPS compliant". It is the manager that is certified for compliancy with the standard.

increasing use is. What few swaps were used five years ago were often accounted for manually as an adjustment to the automated systems; now the business is demanding their integration into production accounting, trading and performance environments.

The problem with swaps, from the systems point of view, is their structure: they are essentially a tandem investment in (at least) two different securities – long on one side short on the other. Thus, the bucketing and aggregation process of performance measurement often wants to place each leg into a different bucket. Take a vanilla interest rate swap: if one wishes to present fixed rate returns separately from floating rate returns, then one leg of the swap falls into each of these buckets. Now imagine an OTC total return swap, trading the return of a custom basket of stocks for the S&P 500, and a portfolio manager who wants to see a breakdown of the return contribution by industry; it gets messy fast.

What it boils down to is a requirement to be able to look at a swap as a single instrument for trading and compliance purposes, but to split it into component legs for performance measurement and attribution. Very few, if any, existing performance systems can support this more complex model of a security; neither do they incorporate in any way the concept of leg as a sub-component of a security.

Caveat Creditor: Fixed Income Attribution has surged as an area of both innovation and investment.

The new millennium's difficulties experienced by traditional equity organizations opened a window of opportunity for their fellow practitioners in fixed income. Prodded by the stick of decreasing equity allocations and coaxed by the carrot of attractive and diverse international bond markets, fixed income became an increasingly important profit center for many investment firms. As a result, larger and larger portions of IT budgets have been allocated to fixed income technology.

However, the state of the science of fixed income performance attribution was not – and still is not – as well developed as that of equity. Historically, performance attribution papers had focused on a handful of sector-based, allocation/selection methodologies and variants that, while putatively general, were strongly biased toward traditional equity investment processes. New research into fixed income attribution burgeoned, and the proliferation of papers, methodologies, conferences and forums continues still.

Thus, while we see a surfeit of demand for investment in fixed income attribution technology, no industry-standard conventions for what these systems should *do* have yet gelled. This situation, while presenting a major conundrum for system vendors, categorically demands a custom solution.

Strategy Happens: The growth of the hedge fund industry has promoted a completely different way of looking at performance.

During this same period, the explosive growth of the hedge fund industry has spawned the kernel of a demand for an entirely new and different method of looking at performance. Furthermore, many of the major features implemented by vendors to satisfy the mature institutional fund market are simply not relevant to hedge fund managers:

- Performance reporting standards do not exist (not yet, anyway).
- Widely accepted benchmarks do not exist, and may not ever be relevant given the diversity of strategies lumped together as “hedge funds”.
- Performance measurement to date has focused on P&L and detailed accounting of the cost basis of positions, with little or no reference to holding period or time-weighted returns.

These dissimilarities have caused the development of a whole separate market for hedge fund software. Comprised of both independent vendors and large prime broker services offerings, these systems typically focus on the trading and accounting functions; sometimes a performance report is tacked on to the end of the process. Indeed, the last bullet point above seems to call into question the need for any sort of return-based performance measurement at all.

However, a different sort of performance analysis requirement has arisen amongst hedge fund managers, and in several more sophisticated institutional departments as well. Generally known as “Strategy-based”, this method attributes contribution to performance to active, manager-initiated strategies as opposed to pre-defined buckets defined by security attributes.

As a simple example of a strategy, imagine a manager shorting a basket of US stocks vs. the Nikkei 225 futures contract. In the case of a hedge fund, this strategy is probably implemented in absolute terms – i.e., with actual longs and shorts in some manager-defined ratio. In the institutional world, it’s more likely defined as exposure relative to a benchmark, and implemented through increases/decreases to current exposure. In either case, what the manager wants to know is, “How much did this strategy contribute to my fund’s return (absolute or excess) during this period?”

At first glance, this would seem to be something that any vendor system could accomplish with only a small amount of effort. All you would need to do is provide some interface or other means of placing these two securities together in the same user-defined bucket by themselves, and proceed as usual. But this approach dramatically under-estimates the differences between strategy-based performance measurement and the current (let’s call it “security-based”) method.

Currently available systems aggregate and report on performance into buckets. Which bucket a position falls into is determined exclusively by one or more attributes of the security held. But in our example, the two sides of the strategy share no security-intrinsic

attributes whatsoever. US vs. Japan, basket of equities vs. futures contract, USD vs. JPY; there isn't anything in common between the two except the fact that the manager is spreading them against each other.

As suggested above, we might create a special attribute for each security that specifies which strategy it's a part of. But suppose that, even though the manager is bullish on Japan over the US, they like the UK even better. So they initiate another strategy shorting the Nikkei vs. the FTSE 100. Now the Nikkei contract has to be able to belong to two different strategies. Worse than that, every system available today will net the two Nikkei strategy holdings into a single position, perhaps eliminating it entirely!

In order to support strategy-based performance, the data model for a system has to be able to do two things that no other system in place can do:

- Maintain multiple positions in the same security, on either side, without netting them together³.
- Identify the inclusion of a position in a strategy at the position level, not the security level.

The Shadow of the Past: Modifying legacy systems to fulfill these new requirements is too expensive for their owners to contemplate.

As we've pointed out earlier, the software currently providing performance measurement and attribution in a given firm may have any of a number of provenances:

- Vendor of institutional performance systems, perhaps integrated with accounting and trade/order management systems from the same vendor
- Vendor of hedge fund operations systems, with or without a performance module
- Off-site system provide as a service of a prime brokerage or custodian
- In-house developed system

Modifying any of these systems to support the functionalities presented in the preceding sections is an extremely expensive proposition. Some of the difficulties are universal; others are unique to particular types. In this section, we'll show why we believe that all of these systems owners will need expert assistance to meet the new demands of their clients, either to overhaul existing systems or build new ones from scratch.

Multi-legged securities

The difficulty the systems' architect encounters with this concept is that the fundamental bottom-level object of almost every performance system⁴ is the security. Most data models have no representation of a leg – an object, one or more of which belong to a security. Furthermore the expense of adding such a construct to existing software is, in

³ Note that this isn't achievable by making use of tax lots; these are generated as a result of trade execution, and do not necessarily correspond 1-to-1 with strategy adjustments. Several institutions have wasted significant resources trying to implement strategy-based performance through "trade-tagging" without fully understanding this issue.

⁴ Or accounting and trading system, for that matter.

general, truly enormous. It is much easier to add aggregate objects to the top of a relational structure – and the code that updates and accesses it - than to the bottom.

So in the short term, the system architect contemplates ways to support swaps without rebuilding everything from scratch. There are two obvious avenues down which to go: represent the swap as a single security, or as two or more securities that (somehow) get transacted simultaneously from one trade. The latter course is fraught with difficulties;

- Traders think you are nuts. To them, the swap is a single security, period. That's how it's quoted, that's how they trade it, and nobody downstream should think of it any other way. In prime broker operations, where the lion's share of IT investment power lies in organizational "silos" – fiefdoms organized around trading operations by security type - this may well be the final word on the matter.
- Many, if not most of the swaps are OTC, and the individual legs don't have CUSIPs or other identifiers, and aren't priced regularly from any current source.
- Compliance system rules will kick in when you start holding short leg positions in the portfolio, and will somehow have to be outwitted.
- It would be even more costly - and thus improbable – to modify the accounting and trading systems to recognize this new security/leg data model. Thus, any maintenance of the relationship is likely to be implemented as a function that occurs after security master scrubbing, pricing, etc. This rapidly becomes a recipe for operational nightmare.

In short order, the expedient solution becomes irresistible: represent the swap as a single security and forego the ability to separate its legs for the purposes of performance measurement and attribution.

Multiple positions in the same security

The difficulty implementing this requirement is painfully obvious to the data architect of any legacy system: Security ID is always in the primary key of the holdings table. To modify this assumption when building a new system is a non-trivial task in data design; to remove this assumption from an existing system would represent a colossal undertaking. Additionally, it would be nigh impossible to pull off without modifications of similar scale in the accounting, trading and compliance systems

Bucketing based on position attributes

Were one able to accomplish the functionality above, this might actually be a comparatively feasible undertaking. All of the bucketing algorithms would need to be modified to look first to the holding attribute and, if missing, inherit through to the security attribute. This wouldn't be a minor change, but more often than not would represent more of a major enhancement than an overhaul. Of course, we're still stuck with that little prerequisite problem...

Problems specific to prime brokers

All of the impediments detailed above apply equally to the technology behind the services provided by large prime brokers to their clients, particularly the hedge funds.

For these firms, however, the degree of difficulty is multiplied many-fold by factors unique to the nature and structure of their business:

- Enormous transaction volumes have often forced these systems to be built in-house. The full cost of revamping/rebuilding them must therefore be funded internally.
- In order to implement these changes, they have to multiply the level of effort by the number of silos with independently built and maintained accounting, trading and other interfaces, and then add the margin system to boot.
- Their historic focus on trading and cost-basis/P&L has left many prime brokers with accounting systems that lack a feature key to performance measurement: “as-of processing”, or the ability to correct position inventories on prior days based on corrections and cancel/rebook activity. Until you can accurately determine what you held yesterday, you haven’t got a hope of calculating today’s return.

Our Conclusion: The near-term future is very bright for the consultant who has expertise in the business requirements and technology of performance systems.

It is clear from our experience and observation that the demand for investment performance systems is not only growing very rapidly, but is fundamentally *not* able to be met by the current batch of available software. The manager demanding these capabilities is therefore presented with one of two rather bleak near-term prospects:

- If their performance analysis software comes from someone else (vendor, prime broker, custodian), then they won't get satisfaction any time soon. Maybe never.
- If they depend on in-house systems, then they may have some greater degree of control over the outcome. But it is, nonetheless, a mammoth undertaking and isn't likely to happen this year.

Both circumstances provide a prime opportunity for experts in this field. We will deliver value to prime brokers and custodians by:

- providing deep and accurate insight into the business requirements for performance analysis
- translating these into feasible data and systems architectures
- maximizing resource efficiency in the development phase of the resultant system

As for the ultimate clients - the institutional and hedge fund managers – the first link in the value chain will be to intelligently assess build vs. buy options, based on our familiarity both with client business requirements and potential available solutions. Subsequently, either a development or an integration project will be undertaken, in the context of which our experience and expertise will once again be invaluable.