Risk Analysis on the Fly:
Fast Markets, Complex Portfolios

TABB GROUP

E. Paul Rowady, Jr. / Adam Sussman | V07:0012 | May 2009 | www.tabbgroup.com

Vision

In response to unprecedented market stresses in the past year, highly complex financial organizations – such as large multi-strategy hedge funds and bulge-bracket banks – need to enhance their data management capabilities, principally for purposes of on-demand portfolio risk analysis. This facility with data management can, in large part, be established by extending a unified and high-performance technical architecture to the entire enterprise.

This unified technical architecture needs to be designed for the improvement of data quality, dramatically higher data volumes, and much greater computational speed whether the underlying market conditions, strategy or portfolio requirements warrant it today or not. The core components of this enhanced architecture already exist within those business units focusing on high-frequency algorithmic trading and sophisticated quantitative research.

The methods employed for data storage – of both streaming and historical data – are central to achieving enhanced enterprise data management (EDM) capabilities, particularly for the exploding volumes of time series data. Better ways of providing the terabytes of trading data to various business functions can drastically improve the efficiency of the organization. Portfolio risk analysis, quantitative research and systematic trading strategies can all benefit from database innovations.

Regardless of how tightly regulated the financial services industry becomes over the next several years, investors and trading partners will demand a new level of transparency and risk management from their money managers and counterparties. Financial institutions will need to support an on-demand reporting infrastructure, capable of reporting across asset classes, portfolios or business lines, as a standard cost of doing business. Enhanced EDM capabilities are a critical requirement for achieving on-demand portfolio risk analysis of complex portfolios as well as multiple portfolios of varying degrees of complexity. Indeed, the difficulty of multi-portfolio risk analytics is effectively established by the most complex instruments held across the firm, usually of the structured or OTC variety.

While high-frequency trading strategies based on exchange-traded products have set the bar for being able to process, analyze and respond to market data, the requirement for more timely risk management has quickly spread across the industry. Now that “complexity” and “financial innovation” are four-letter words, firms must be able to tame that complexity by breaking down instruments into byte-sized datasets that can be quickly and confidently valued.

The pursuit of an on-demand infrastructure – one that responds as fast as changes occur in the relevant underlying market(s) and predicts such changes with increasing regularity – is the best strategy for taming complexity because it typically forces a new level of data and process automation (as well as a “data culture”). If your trading organization can respond on-demand, then it is likely it can respond as fast as will ever be necessary.

Table of Contents

Introduction

If the economic crisis of the past year has taught us anything, it is that more-active measurement and management of risk in complex portfolios and across complex trading organizations is a potential source of profound differentiation. We need not look any further than the algorithmic trading community for proof of what is now possible in taming complexity risks.

Despite the loss or disembowelment of some of the most venerable financial institutions and hedge funds, a number of strategies performed well throughout 2008. One common theme among these success stories involved highly-automated and high-frequency trading, predominantly in exchange-traded securities, such as equities, options and futures (see Exhibit 1). In fact, of the top 10 best-performing hedge fund strategies of 2008, six depend upon some mix of the kind of on-the-fly analytics discussed in this report. While the performance of these strategies may not fare as well in the years to come, there is no doubt that the technologies and requirements needed for them will enable those firms to succeed in other strategies and therefore pull ahead of their competition.

Exhibit 1
Top 10 Best-Performing Hedge Fund Strategies of 2008

Fund Strategy	# Funds	2008 Return
Short Bias	31	30.61%
Mortgages	29	20.40%
CTA/Managed Futures	623	10.97%
Other Arbitrage	16	9.44%
Short-term Trading	13	6.49%
Statistical Arbitrage	21	4.18%
Options Strategies	86	0.79%
Global Macro	250	-1.75%
Merger/Risk Arbitrage	25	-2.89%
Market Neutral Equity	244	-3.42%

Green indicates high degree of automation.

Source: HedgeFund.Net

The algorithmic trading community has been at the forefront of technical advances in financial markets for well over a decade. The "quant" technology stack has some of the best gadgetry in the business. High-performance analytics tools perched on top of an infrastructure designed to perform within the grasp of the speed of light are becoming the standard for electronic trading environments. In fact, many of the competitive advantages in high-frequency trading have now become competitive necessities. As a result, there is a zero-latency arms race being waged as firms buy, build and configure the hardware and software necessary to shave the last few microseconds from a trade cycle. Even urban legends about cutting the last 1/8-inch of slack out of network cables have begun to circulate. The landscape has matured to the point where the barriers to entry are substantial, though there are still some opportunities on the fringes using computational art. Clearly, the game has evolved from the milliseconds scale to the microseconds scale, and while we once talked about market microstructure, we now must consider market and systems nanostructure.

Managers of systematic, high-turnover strategies have led the demand for all of the tools necessary to support their research, trading and risk analysis. The trick is to configure the hardware to deliver the maximum speed for the maximum level of capacity (and at the highest fault tolerance). After that, the next trick is to configure a pattern recognition and decision support platform with the highest level of consistency and adaptability requiring the minimum level of human intervention.

The wisdom of high-frequency trading strategies begins with the exploitation of the most liquid, standardized and transparent markets. While it is by virtue of these market attributes that an abundance of clean, granular data is yielded, successful strategies ultimately hinge on an incredibly advanced ability to manage time series data. From the vast store of time series data, a spectrum of pattern-recognition techniques are deployed that ultimately result in trading signals. (In fact, today, a trading signal and an order are indistinguishable.) As one might expect, there are numerous tasks that must be done *simultaneously* in order to construct the optimal real-time data management infrastructure. The following shed some light on the issues and challenges faced by this modern day, high-performance "data engine:"

- **Updates:** Keeping up with the amount of streaming data;
- **Access:** Allowing multiple queries on the same dataset(s) while not slowing down the retrieval of the data;
- **Logic:** Performing analysis and making trade-related decisions;
- **Redundancy:** Backing the data up in real-time; and
- **Dynamic Storage:** Quickly and cleverly moving data from memory to disk based on business function.

It used to be that there was sufficient time – or lack of an alternative – to upload data into analytics tools in order to make risk-transfer decisions. Now, given the critical time sensitivities of the on-demand trading enterprise, particularly in the areas of trading and risk analysis (as well as in research and even compliance), performing time series analytics *inside* the data engine is the only way. Without this functionality, "on-demand-ness" or "real-time-ness" is simply not possible.

Coupled with the advanced speed of automated order matching, human inefficiencies and biases are all but eliminated from the trading strategy. In essence, it is these market attributes that provide the very foundation for automation. Automation provides the foundation for speed. Speed provides the foundation for agility. And agility – combined with intelligence – is how you survive and thrive in uncertain times, which are most of the time, and definitely like right now.

Meanwhile, at the other end of the spectrum, the recent failings of fixed income, structured products and over-the-counter (OTC) derivatives businesses have fallen into the crosshairs of public opinion, although perhaps for reasons that the mainstream media have yet to fully appreciate. Unlike the quant community, the secrets of their headline-grabbing failures are actually not secrets all. The reasons are quite clear, and the first reason, oddly enough, is their success.

Since 2000, many structured products businesses grew rapidly and profits gushed – a timely tonic for dwindling commission revenue, post-dot-com deal fatigue and other decaying sources of alpha. Like a go-go start-up, policies, technical infrastructure and systems capabilities may have lagged behind the growth,

thereby fueling operational, market, credit and liquidity risks that we all have had the great misfortune to witness, and which are common to young businesses. Moreover, it turns out that structured products and OTC derivatives businesses have the kind of characteristics that lead to high degrees of complexity. For instance, in these businesses, there are few standards, very little transparency, illiquid markets, numerous data challenges, and no central clearing counterparties; this is a scenario that makes the accumulation of time series data for purposes of research and risk analysis problematic.

This is almost the exact opposite of the attributes of high-frequency strategies. In fairness, automating the work- and data-flows of a structured products business is really, really difficult. So complex, in fact, that there may have been little choice, given the circumstances. Pursue the business without the fully-baked, omniscient infrastructure, or stay on the sidelines and cede the winnings to the competition.

Complexity Defined

In computational complexity theory, the complexity of problems – that is, the difficulty of solving them – is defined by the time it takes for a computer program, usually an algorithm, to solve it.

Participants in financial markets are generally accustomed to complexity and the process of solving problems and providing clarity where little or none had previously existed. After all, this lack of clarity or transparency is often where market inefficiencies live, and therefore, where un-harvested profits typically can be found. So it should come as no surprise that financial innovation and engineering are engaged in the process of discovering market inefficiencies, harvesting "alpha" and developing products and markets that serve unmet needs in new, unfamiliar and often challenging ways.

Today, we can think of each financial asset or product class along a perceived spectrum of complexity, from the relatively simple (e.g., publicly-traded common equities) to the extremely complex (e.g., structured products, like ABSs or the nearly unfathomable poster child for perceived complexity, "CDOs Squared").

For financial products, TABB Group believes there are at least six primary attributes that explain most, if not all, of the potential complexity of a given financial product. These factors include:

- **Computation:** The number of variables that need to be calculated in order to arrive at a reasonably accurate asset price;
- **Standardization:** The similarity of attributes for products in a given group or asset class;
- **Liquidity:** The frequency by which a transaction price can be observed;
- **Data:** The amount and specifications of data, including granularity, cleanliness and/or depth of historical archive, necessary to accurately estimate a current asset price;

- **Technology:** The amount and nature of equipment – hardware and software – and personnel expertise necessary to reliably conduct research, trading and risk management in this product category; and
- **Counterparties:** The level of counterparty risk; the presence of a centralized clearing counterparty.

Each asset class can be judged somewhat objectively according to its complexity across each of these six attributes. The total complexity for each asset class would be the sum of its complexity scores. Interestingly, there is also a relationship among these six attributes. For example, as a product becomes more standardized, it becomes possible to distribute more data. As more data becomes available, more technology is needed to process it (see Exhibit 2).

Exhibit 2
Complexity Factors by Asset Class
(On a Scale from 1-5, 1= Low, 5=High)

Complexity Factors	Equities	Futures / Options	Fixed Income	OTC Derivatives	Structured Products
Computation	1	3	2	5	5
Standards	1	1	2	4	5
Liquidity	2	3	4	5	5
Data	1	2	3	4	5
Technology	1	2	2	5	5
Counterparty Risk	1	1	5	5	5
Total Complexity Score	7	12	18	28	30

Source: TABB Group

Now consider the additional complexity inherent in a particular trading or investment strategy. Other than the increased complexity caused by the combinatorial effects of the aforementioned asset classes, it turns out that the need for speed, intelligence and agility are the primary antidotes to complexity. Just like cholesterol, there appears to be "good complexity" and "bad complexity." Taming bad complexity is what we think it means to perform various tasks "on the fly." Given an infinite amount of time, anything can be accomplished and thereby dilute the perception of complexity.
The objective, however, is to accomplish more inside the boundaries of maximum intelligence in the least time, consistently, while the landscape beneath your feet is shifting.

Based on the foregoing discussion, our search for the outer boundaries of complexity – and "complexity risk" – in financial markets would end at the point where all complexity factors exist in their highest forms and at the greatest scale. For example, imagine an entity with multiple portfolios at multiple turnover frequencies across multiple combinations of asset classes – many of which are illiquid – in multiple regions at large scales for both exchange-based and exotic instruments. On the sell side, this scenario describes the largest, integrated financial intermediaries, such as global bulge-bracket banks and insurance companies. According to recent financial reports, the products that are the most difficult to value, deemed Level 3 by the Financial Accounting Standards Board's Ruling No. 157, represent 5% of reporting firms' balance sheets, while products of intermediate difficulty make up 94% of the balance sheet (see Exhibit 3). On the buy side, this can only mean the largest multi-strategy hedge funds and asset managers, sometimes known as "near banks."

Exhibit 3
Bank Holdings by FAS 157 Definitions

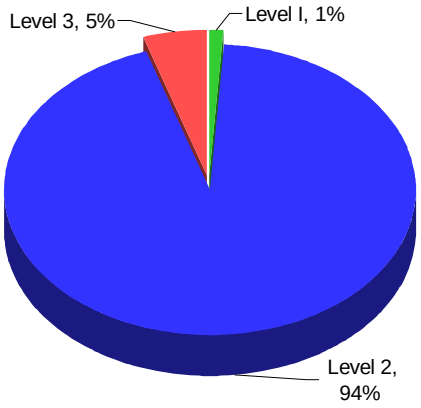

Source: Company Financials, TABB Group estimate

The illustrations to describe complexity in today's financial markets seem almost endless. Consider, for example, the complexity of a mortgage-backed securities business. First, to do any asset-backed securities (ABS) strategy well, you should have transparency to loan-level data. For many players, you can stop right there because they don't have it. For those who do have loan-level data and the transparency to detect the linkages, the historical dataset is comprised of over 1 billion records – easily terabytes of data. Optimally, your ABS strategy should also be able to match loan-level data with credit data and changes in credit histories. If this dataset exists at all, it's another footprint measured in terabytes. Once the granular data is identified, algorithms to estimate and predict prepayment, default and other risks need to be developed.

While it is true that there is no high-frequency trading strategy on Level 3 assets, it is equally true that the ability of a financial institution to properly assess its risk at any given time is dictated by the weakest link in its risk analytics arsenal. If a multi-strategy fund has all of the right technology for its high-frequency equities fund, but lacks the data and modeling techniques to properly value its exposure in CDOs, from a risk analysis standpoint its enterprise risk assessment could be dangerously out-of-date, or worse, just plain wrong. Indeed, the quant meltdown of August 2007 is often laid at the doorstep of a certain multi-strategy hedge fund that did not aggregate its natural gas exposure on the same risk platform as other strategies. (Or at least it did not update this risk exposure at the same frequency as its other exposures.) Indeed, one of the requirements that could result from regulators' focus on systemic risk will be the ability to calculate on-demand exposure and other risk measurements.

In this report, we differentiate between real-time and on-demand with regard to the update frequency of various processes. Real-time often conjures up the perception that a given process must be recalibrated every few seconds. This is often unnecessary since the underlying markets may not update at that rate, but instead, perhaps once every few minutes or hours or sometimes days (in the case of the most illiquid markets). Therefore, in general, a process should be updated as often as the underlying market. For this reason, we emphasize the use of the term, "on-demand" for its greater accuracy in explaining requirements. On-demand also implies that all of the data is aggregated and ready to be analyzed, as opposed to a bunch of ad-hoc queries being run on disparate databases and then loaded into an application, or likely a spreadsheet, for analysis.

Risk Analysis on the Fly

Enterprise Risk Management (ERM) can trace its roots back to the 1970s when Swedish risk manager Gustav Hamilton first introduced the concept of the "risk management circle," an exercise to identify and consider the interrelationships among all of the organization's risks. If the devastating market events of the past several months are any indication, it appears that many lessons have yet to be learned and that ERM still may have a long way to go.

ERM was supposed to break down the silos, let the data flow (except in the case of regulatory barriers), and usher in an era of greater knowledge about firm-wide market, credit and liquidity risks. The reality, as we now painfully know, is that the promise of ERM has yet to be fully realized. The good news is that the concepts of ERM are still fundamentally valid. Achieving greater clarity about portfolio or firm-wide exposures is a positive step. The bad news is that in the aftermath of a meltdown, budgets are slashed, personnel jettisoned, and lines of authority obfuscated. Triage is the order of the day, and it will take some time and effort to chip away at and tame the complexity that has been created.

The two primary obstacles to achieving an optimal ERM platform are breaking down the silos of asset classes and business lines, and the related problem of sharing data across the institution. Tackling these obstacles is the litmus test to determine the health and status of ERM, and the ultimate prerequisite for performing on-demand risk analysis. As the tenets of ERM become more firmly ensconced in the requirements of regulators, investors and counterparties, financial institutions managing portfolios containing complex securities or multiple, large-scale portfolios will need to make further headway on these two fronts.

Break Down the Silos

The pillars of the Knowledge Enterprise are people, process, technology and data. Develop best-of-breed in each of these categories, establish and maintain maximum harmony among them and your organization is likely to occupy a leadership position in the business.

First step: Break down the silos. The only thing that belongs in a silo is corn and, possibly, lawyers. In all seriousness, there are valid reasons to "wall businesses off" for regulatory reasons, but everything else should be set up for mobility. Sure, people spend most of their time in a vertical or horizontal orientation for purposes of an organization chart and command-and-control operations; but there is even a good argument to suggest that rotating personnel into different seats over time would pay dividends.

Technology certainly should not be constrained to a silo if it can serve a broader audience. If it works, it should be extended to the entire enterprise, wherever applicable. Most importantly, data should not be constrained to silos, either. Any dataset that can be used to improve the effectiveness of research, trading and risk management needs to be set up for broader access with consideration for speed,

intelligence and agility. This is particularly true for all types of time series data, from market data to risk factors to economic and other fundamental data.

None of this advice should be new. And there is plenty of evidence to suggest that many trading organizations have not only implemented this strategy, but have made knowledge management a key tenet of their culture. You can bet that the firms with the best track records have figured this out. Another way to think of it is that the best long-term track records are directly correlated with the ability to manage data.

Data on the Fly

At the end of the day, people, process and technology in a trading organization are all employed to manage data. Discovering, capturing, storing, accessing, analyzing and manipulating data are the lifeblood of all trading enterprises. Doing these things really well – and now, really fast – is the lifeblood of the on-demand trading enterprise. It seems strange to think of it this way, but it's true. Finding unique sources of it, identifying novel patterns in it, and making decisions with it (sometimes faster than the rest of the market, sometimes just at the right time), are all aspects of managing data that ultimately create value for stakeholders.

Data that is useful for trading is typically broken down into three categories:

- **Market Data:** Includes both streaming and historical, and is comprised of long strings of prices (trades and quotes) and volumes;
- **Reference Data:** Denotes all types of metadata to describe the terms of securities and the companies that issue them, and is most often static; and
- **Fundamental Data:** Financial, economic or calendar data, comes in a broad spectrum of shapes, sizes and frequencies and can usually be represented as a time series, like market data. (Risk factor time series are derived from both fundamental and market data.)

Streaming market data has received much of the recent attention because of its unique complexities, particularly in relation to volumes and velocity. For example, in 1993 the New York Stock Exchange averaged fewer than one million trades and quotes per day. Since then, decimalization, steadily decreasing average trade sizes (due to automation), and market fragmentation caused in large part by Reg NMS in the US and MiFID in the UK, have played a role in the explosion of equity and equity-related market data volumes. In comparison, the 2008 peak trades and quotes stream exceeded one billion records in a single day.

Recent observations show that average daily message rates across all live US equity and options market data feeds is between 400,000 and 600,000 messages per second (mps). In terms of speed, the historic peak in market data rates, as measured since April 1, 2008 by Market Data Peaks (www.marketdatapeaks.com) was 1.56 million messages per second (mps) on December 5, 2008, more than tripling the record just three years earlier (see Exhibit 4). Recent automation through various electronic trading platforms has also had a dramatic impact on market data volumes and velocities for futures and even fixed income.

Exhibit 4
Aggregated One-Minute Peak MPS Rates
Arca, CTS, CQS, OPRA, NQDS

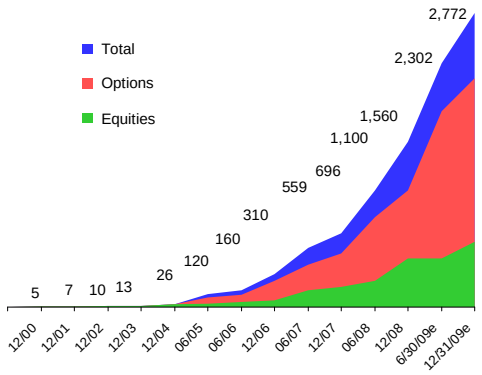

Source: SIAC, OPRA, NASDAQ, FIF, TABB Group Est.

The current expectation is that market data volumes and message rates will continue to expand across all of these asset classes. Indeed, even though market data spending will remain flat over the next two or three years, because the increasing velocity of messaging rates and the recognition of the value in processing and responding to those messages is critical to successful trading strategies, financial firms will increase their spend on direct market data feeds at a 17.7% CAGR (see Exhibit 5) between now and 2012.

Exhibit 5
Global Market Data Spend (in $MM)

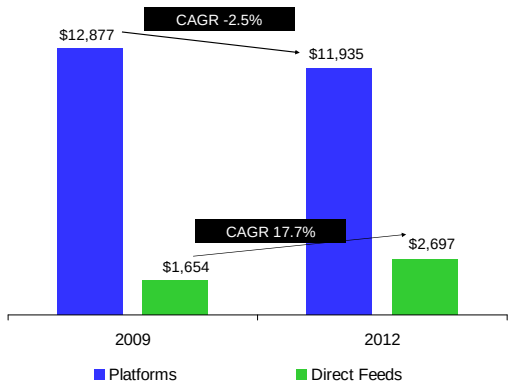

Source: SIAC, OPRA, NASDAQ, FIF, TABB Group Estimates

If all that weren't enough, streaming data must ultimately be placed in storage to become part of the historical archive used to support all manner of quantitative research, risk analysis and more often of late, compliance. In terms of historical data, all of this activity – ticks and quotes data – currently translates into roughly 50-400 gigabytes a day in global uncompressed market data that must be carefully collected and stored (see Exhibit 6). To add a finer point here, many market participants opt to collect all of this data themselves to insure the highest quality possible throughout the research and trading process. Furthermore, extend this analysis beyond exchange-traded instruments and the data volumes have

Exhibit 6
Storage Requirements for US Equity Options

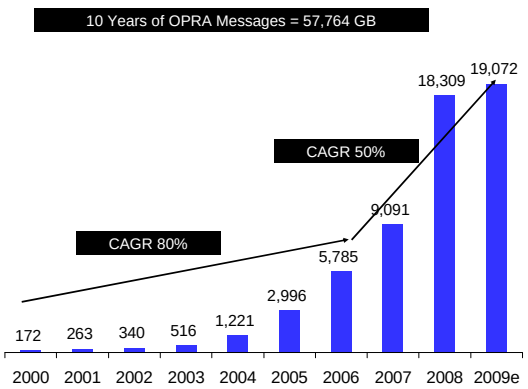

Source: OPRA and TABB Group Analysis

the potential to become even more impressive. Finally, consider that data is being retained for longer periods of time because it can improve the accuracy of analyses and because compliance may be forcing longer retention.

Over the course of a year, all of this new data can be measured in terabytes, or trillions of bytes. All in, including market, reference and fundamental data, the typical enterprise data store for large, complex trading organizations can total hundreds of terabytes of data. Luckily, the current economic environment has conditioned us to think in terms of trillions!

The data challenges are particularly exacerbated in illiquid markets, such as structured products and OTC derivatives. Since there are often few standards here, there is rarely a centralized trading mechanism, which of course means there is little automation. Much of the data management functions here rely on highly manual processes which then lead to numerous errors and inefficiencies. And since market prices often cannot be observed in illiquid markets, elaborate modeling is typically required to value these securities. In these cases, modeling can be effective, but only given broad archives of comparative data (from similar securities) and deep archives of underlying fundamental data, such as loan-level data in the case of asset-backed securities.

Adaptive risk analysis is judged by its ability to rapidly reflect changes in market conditions. This rapid reaction and corrective prediction requires an advanced (and relentless) ability to manage and manipulate data, as in a data engine. It is elemental. In other words, the only way to establish on-demand risk analysis – Risk on the Fly! – is to first establish data on the fly.

Data Engine: A Requirement for Take-off

The level of speed, intelligence and agility that a trading enterprise must have will be determined by the needs of the underlying trading strategy or strategies it supports. Needs are also determined by the specific processes within the overall strategy workflow, e.g., research, trading, risk analysis, compliance, etc. Although this report has alluded to some of the most extreme cases of complex asset classes and sophisticated high-frequency strategies, the truth is that most firms have yet to attain an optimal state of technology for many of the most common scenarios.

Core to achieving this optimal state is making sure that the data most impactful to the enterprise is able to reach any touch point within that enterprise. It is not just how quickly an infrastructure receives the data that determines its ability to react to that data. In order to make sure silos can be broken down and data flows are omni-directional, the storage facilities for the gigabytes of streaming data and terabytes of historical data that are generated each day must be able to quickly share that data throughout the trading infrastructure (see Exhibit 7).

Exhibit 7
Trading Infrastructure

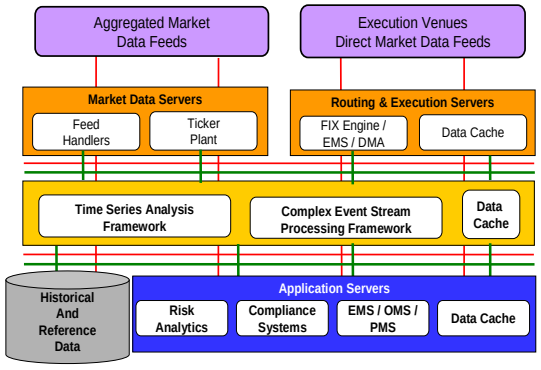

Source: TABB Group

One of the most common uses for this data is to conduct quantitative research, where there is a critical need for relatively fast – but not quite low latency – access to a vast historical archive of market and fundamental time series data. And since quants are always looking for new correlations and indicators among old and new datasets, quant research is a primary beneficiary of the "no silos" mandate. Quant research is highly iterative, using various pattern recognition techniques on nearly infinite combinations of datasets to identify possible sources of alpha. This process is always time-consuming and fraught with numerous blind alleys that typically yield little redeeming value other than the knowledge of what not to do next time. The challenges here are grossly augmented when time series data is unscrubbed and dotted with errors.

In the past, quants would spend much of their time collecting, scrubbing and otherwise nursing data for testing. Fortunately, market data capture and processing routines have advanced to the point where dirty data is no longer nearly as much of a problem and researchers can now focus much more time on actually performing tests. However, data quality and cleanliness problems of the recent past in equities and related markets ominously foreshadow the challenges yet to be overcome in many of the developing markets for complex securities.

Once the data is clean, the process of rigorous walk-forward backtesting can begin in earnest. In the old days, say back in the mid- to late-90s, many backtests were performed overnight due to computational complexity and the sheer volume of data necessary to perform the test. Of course, back then, quant research was comparatively rare. Imagine waiting for hours to complete a backtest. Then, imagine the aggregate costs in terms of time and lost opportunity across a large trading enterprise caused by all that waiting around. Although we are unaware of that ever being quantified, it is likely to represent many person-years costing millions of dollars.

Beyond quantitative research is the actual implementation of a strategy. The most often talked-about quantitative strategies are of the systematic, a.k.a. algorithmic, variety. Because the decisions have a direct impact on P&L and the timing is more important, the data management objectives are different. In algorithmic trading, streaming market data is flying from the end of the proverbial fire hose at increasing rates of speed. The challenge is not only to capture all of it – without missing a single drop – but keeping it all IN ORDER! Then, to add insult to injury, the trading platform must be able to respond to this streaming data – allowing time for computational processing – in increments of milliseconds and microseconds. If that weren't enough, this massive cache of data must eventually be transferred from memory to disk with sufficient cleverness and efficiency so as to allow for simultaneous performance of queries on the very same data. Talk about high-performance multi-tasking!

These days, the focus has shifted somewhat from quantitative research and systematic trading strategies to portfolio and enterprise risk analysis. Yet again, the data management challenges for portfolio risk analysis – and certainly anything approaching on-demand portfolio risk analysis – are also unique. In this case, it is time series data representing numerous derived risk factors which are periodically loaded into memory to allow scenario analysis to be performed at the highest speeds. The challenge here is for the data management platform to allow high throughput from disk to memory, particularly for scenarios that require frequent updating of the time series data.

With these diverse data management needs within strategy groups and across the trading enterprise, it would appear to be a significant technical challenge to devise a unified enterprise-class data engine to properly balance both the varied response rates and the memory-to-disk issues while simultaneously providing consistent scaling and complex provisioning at reasonable costs. Fortunately, the quant technology stack can provide some guidance here.

Database Innovation

Given the aforementioned requirements for the on-demand trading enterprise, the database has undergone a dramatic transformation over the past few years in response to significant demand for speed, scalability and interoperability at lower costs. Indeed, alongside messaging middleware, the database forms the critical path to the information sharing and the enablement of on-demand risk analytics.

The key transformative database features are as follows:

- **Columnar Design**: Storing unique metadata types in columns instead of several metadata types in rows to address throughput speeds, storage capacities and total cost of ownership (TCO);

- **Compression and Encoding:** Methods to minimize the storage requirements – in terms of both space and cost – of raw data while maintaining full query access to the data (i.e., no decompression required) and while maximizing throughput speeds;

- **Fast Loading:** Methods to transfer data from memory to disk (and back) at high speeds while maintaining continuous access to all data (rather than prohibiting access to data while it is being written to disk); and

- **Embedded Processing Logic**: Methods to perform complex analysis and decision support at ultra-low latencies while data resides in memory.

Traditional relational databases store data in rows, hence the name "row store." These rows, often known as records, are made up of various fields. In order to perform a query on the database, each of the relevant rows – and each of the fields therein – must be accessed to extract the relevant data. The time taken to read irrelevant and redundant fields is highly inefficient and can lead to severe throughput bottlenecks in the market analytics platform – like slow backtests or infrequent updates to risk factors. Alternatively, there are many kinds of data that can be stored in a column orientation, thereby allowing the same field types – like date or security identifier – to be stored in the same column. This revised database design is known as a "column store." When performing queries on a column store, only those columns with the relevant fields are accessed, thereby eliminating some of the inherent inefficiencies experienced when querying a row store. Time series data is particularly well-suited for storage in a column store due to its standardized structure and minimal metadata types.

For maximum efficiency impact, data compression goes hand in hand with a columnar database design. In a column store, the values within each column have the same field type and typically a lot of repetition. A timestamp field is a good example of this, where few numbers actually change from one record to the next. This repetitive scenario, coupled with the homogeneity of field types found in column stores, lends itself to much higher compression ratios than in a row store, since redundant data can be encoded to a smaller byte size.

At the end of the day, data compression is powerful for its effects on the total cost of data warehouse ownership since less overall data requires a smaller hardware footprint. In the US equity markets, trade data spread across 10,000 instruments on all exchanges and liquidity pools at roughly 100 million trades per day can amount to 30-50 gigabytes of raw, uncompressed data. Over the course of 252

trading days a year, this data can amount to more than 1 terabyte (TB). Using standard opaque compression techniques can yield a compression ratio of 2:1. A comparable column store can yield compression ratios of 5:1 on trade and quote (TAQ) data and 10:1 on trade audit trails (according to one vendor). This means that data compression could reduce the amount of storage by roughly 9 terabytes per year for US equities alone. For larger organizations trading global equities and derivatives, the storage space savings would be well over 100 terabytes.

In addition to compression and encoding, data loading can have a dramatic impact on performance, particularly in trading platforms that have a continuous need to digest and use streaming data. Today's database infrastructure must be capable of transferring data from memory to disk as quickly – or nearly as quickly – as the speed of the underlying data. This necessity poses challenges because the physical limitations of disk drives – due to disk rotation speeds – prevent storage that is in harmony with the speed of the data. One technique has been to write data only to the outer 30% of the hard disk because these regions can be accessed faster relative to inner disk regions. Unfortunately, this technique can only go so far.

In-memory storage and methods to efficiently and cleverly transfer data from memory to disk have grown to address this challenge with solid success. In some cases, high-performance databases can load continuously, while the database is being queried, thereby allowing for on-demand analysis of the absolute most current time series data.

The key to storing mass quantities of high velocity data is to cache it in main memory (RAM) for processing later – which in many cases is just a few microseconds later. The alternative is to be on the wrong side of Moore's Law. While processing speeds continue to increase, hard-drive disk rotation speeds have topped out at roughly 15,000 rpm. Therefore, databases that trade a lower dependence on disk drive access speeds for more reliance on processing power – through in-memory caching –benefit from the prevailing trend represented by Moore's Law. Caching data is critical to the process. Otherwise, data can ultimately be lost or the data engine chokes. Overall, there must be a highly intelligent balance struck between memory and disk.

Not only does the main memory cache serve as a buffer, but it also serves as an in-memory database with embedded processing logic that is able to aid in the performance of complex and ultra-low-latency analysis and decision support. This feature is absolutely essential for on-demand analytics. Without the ability to tame the most complex aspects of time series data management inside the data engine, the vision of the on-demand trading enterprise will remain elusive.

There are many tools configured to work in harmony that are required to do anything on the fly within a trading organization. Having the right database platform – a data engine capable of extremely high performance on both streaming and historical time series data (measured in microseconds and terabytes, respectively) – is essential to all of them.

Conclusion

The Algorithmic Age has ushered in a level of capability never before seen in human history. The accomplishments have been impressive. At the core of it all is the ability to manage the markets' most precious resource – data – with unprecedented speed, intelligence and agility. Capture, slice, dice, filter, sort, calculate, compare, analyze, decide and store it, all in less time than it takes a hummingbird to complete half a wing-stroke. Now, consider enough hummingbirds to cause an eclipse of the sun, and we begin to appreciate the volume, velocity and diversity of today's data management landscape and the speed necessary to capture and use it.

Yet, the rest of the industry has not necessarily benefitted from the unprecedented amount of data and technology available to the black box trader. On August 6, 2008, the Counterparty Risk Management Policy Group (CRMPG) published a report including a new set of guidelines and recommendations for major market participants. If complex trading organizations are the same entities that struggled during the recent market crisis, and there seems to be considerable evidence to support this conclusion, then this is precisely where a more actively-managed, on-demand risk analysis paradigm needs to be targeted.

Three of the CRMPG's recommendations are particularly relevant to our discussion and highlight the disparity that exists between the transparent and straightforward exchange-traded securities and the more-complex, opaque financial instruments, as well as the growing adoption of more sophisticated portfolios and multi-strategy entities:

- Make a reasonable effort to speed up the adoption of electronic platforms; presumably this includes both internal management and for trading.
- Increase usage of tools and methods that facilitate the standardization of instruments.
- Insure interoperability across the enterprise and throughout the industry by deploying third-party tools when the market has a suitable product.

These three recommendations are born out of the progress of more mature financial instruments, shaped by years of busts and booms, technological advancements and regulatory oversight.

The necessity to bring the same fluidity of data that exists within most US equity trading desks to the entire enterprise will require a painstaking amount of work, industry-wide cooperation and innovation from third-party vendors.

- Technological silos must be viewed as no more than virtual, and necessary only to illustrate how data is organized and disseminated according to the regulatory and business requirements of teams and business units.

- Allocate the resources to nail down detailed workflows. This step is essential to establishing the behaviors necessary to foster a culture that cares about proper handling of the data required for world-class research and risk analysis.

- Unified technical and operational infrastructure must supply the plumbing and security so that data can flow with more speed, intelligence and agility throughout the enterprise.

- Extend the core data management components of the quant technology stack to the entire enterprise to minimize the complexity of portfolios with different types of instruments.

Accomplishing these goals means more than just gigabytes of memory, hundreds of ultra-fast, multi-core processors, terabytes of storage, and the latest fiber optics and 100 GbE network switches to connect it all. It also means deploying the most refined processing logic and a high-performance data store – dubbed herein as the data engine – to keep the monumental list of datasets concurrently organized and transparent for a vast and diverse array of end users. This is a requirement to both overcome the intense challenges of time series data management and for timely, adaptive responses to the markets.

In the end, complexity can only be perceived on a relative basis. In terms of taming modern market complexity risks, the algorithmic trading community has provided ample evidence that the bar has been raised. To borrow a phrase: Impossible is nothing.

About

TABB Group

TABB Group is a financial research and strategic advisory firm focused exclusively on capital markets. Founded in 2003 and based on the methodology of "first-person knowledge," TABB Group analyzes and quantifies the investing value chain from the fiduciary, investment manager, broker, exchange and custodian. Our goal is to help senior business leaders gain a truer understanding of financial issues and trends so they can grow their business. TABB Group members are regularly cited in the press and speak at industry conferences. For more information about TABB Group, go to www.tabbgroup.com.

The Authors

E. Paul Rowady, Jr.

Paul Rowady is a senior contributing analyst at TABB Group. Mr. Rowady became affiliated with the company in 2007 after serving for several years as an independent consultant focusing on the design, development, strategy and research for market intelligence systems. Prior to this, Mr. Rowady spent sixteen years with proprietary trading and hedge fund companies in leadership roles for research, risk and portfolio management related to derivatives and high-frequency trading. He earned a Master of Management from the J. L. Kellogg Graduate School of Management at Northwestern University and a BS in business administration from Valparaiso University.

Adam Sussman

Adam Sussman is director of research at TABB Group. Mr. Sussman joined the firm in 2004 as a senior analyst, after serving as a senior product manager responsible for order management systems, routing and next-generation trading tools focused on the equities and options markets at Ameritrade, Inc., a brokerage industry subsidiary of Ameritrade Holding Corporation. He earned a BA in philosophy and comparative literature at the University of Rhode Island.

At TABB Group, Adam has authored a number of reports, including "*Equity Swaps and OTC Options: A Buy-Side Perspective*"; "*International Perspective on Transaction Cost Analytics*"; "*Performance Anxiety: A Buy-Side Study on Benchmarks and the Investment Process*"; "*European Institutional Equity Trading 2007; "Modular Algorithms: The Growing Choice of Buy-Side Execution Strategies"; "Institutional Equity Trading 2006"; "Hedge Funds 2006; Outlook on Algorithms"; "Trading Under a Microscope: The Buy-Side Perspective on Transaction Cost Research";" Institutional Equity Trading 2005"*; and *Managing Risk in Real-Time Markets".*

TABB GROUP

www.tabbgroup.com

Westborough, MA
+ 1.508.836.2031

New York
+ 1.646.722.7800

London
+ 44 (0) 203 207 9397

