

Leveraging Cutting-edge Market Data Technologies For Greater Efficiencies and Lower Costs

Following is a transcript of the presentation and discussion that took place during the webcast on Nov. 10, 2009. David Brukman, VP technology at Interactive Data made a brief presentation and moderated the discussion on some of the latest processing and networking technologies, their benefits for market data consumers, and the issues involved with implementing them. The panelists were Stuart Breslow, CEO/CIO Townsend Analytics/RealTick; Falke Bruinsma, Chief Architect, WebSphere Front Office, IBM, and Kevin McPartland, Senior Analyst at the Tabb Group.

David Brukman I want to start with defining the challenge. Market data rates have continued to grow. OPRA (Options Price Reporting Authority), the largest source of data in the market data world, has been increasing at an annualized rate of 112% per year over the last six months. New York Stock Exchange feeds have been increasing at an annualized rate of 167% over the last six months

At the same time volumes have continued to skyrocket, tolerance for latency in our industry has continued to shrink. That means some of the more traditional computer science techniques for coping with huge volumes like pipelining are no longer valid because they increase latency in order to cope with throughput.

The next slide shows some of the tools we've been using both traditionally and most recently to cope with these huge volume increases. First of all, there have been advances in network technologies. In order to deliver the data we already have developments in 10 Gigabit Ethernet that is becoming commoditized now with 40 and 100 Gigabit Ethernet on the horizon. There have been developments in data center extensions. Infiniband has been an interesting technology. Various versions of wavelength division multiplexing, including all the way to ultra dense wavelength division multiplexing, have made it possible to put more and more data through the fiber. Finally, Remote Direct Memory Access (RDMA) for both Infiniband and Ethernet has been making inroads in reducing latency.

On the processing front we have some interesting developments as well. We have more power but new algorithms are required in order to cope with it. We have developments in general purpose processors including for instance in ones from Intel and AMD, things like multicore processors, more support for hyperthreading, ultra wide- memory bandwidth; we have developments in parallel co-processors that started off in many cases from the graphics world that are making inroads into general purpose processing including cell broadband engine, GPU and others; and finally we have probably the most customizable methods, such as field programmable gate arrays, both the actual tools and sometimes the cards and appliances are available from the vendors in their entirety.

However, while all of those techniques help cope with the data, they also present frequent challenges, and certain challenges are inherent actually in the problems themselves. First

of all, Moore's Law has been extremely helpful to the processing capability, and it states roughly that the number of transistors on a chip double every two years. From my earlier slide you saw that market data volumes actually double about every one year. So they are actually outpacing the rate of Moore's Law. The straightforward answer to that is to throw more hardware at it but we're beginning to run into issues with power and space. There's also another limitation on parallel processing which is called Amdahl's Law. In essence it says that the serialized portion of the problem is one that is going to dominate as the number of processors increase. The graph shows the linear increase that would be there if the problems had been completely parallel, and the lower line actually shows that as you add more processors, the incremental speedup actually decreases until eventually it flattens out. That's a problem that we all face with parallel processing architectures.

In summary, the market data volumes are rising and firms are asked to process these volumes in a shorter period of time. We went from milliseconds to microseconds and in some cases nanoseconds. We face limits on power and space, our ever increasing footprint is running into limitations on cost and actual building capacity. We need to redesign our algorithms to take advantage of non-sequential processing, and also we have organizations that have diverse needs of data, all the way to from extremely low latency tick by tick data to something that is actually consolidated and slows down because the applications can't cope and don't need extremely high volumes of data. So new processing and networking technologies can help us solve these problems, but the answers are often non-trivial.

What are the main factors contributing to the costs of effectively handling market data today?

Falke Bruinsma: Obviously the hardware costs which indirectly contribute to power consumption and space allocation and certainly there are the software costs that come with that. Additionally and more importantly is the maintenance of such a system. Keeping it up and running and maintaining it in the ever-changing market with exchanges consolidating and feeds changing, the upkeep of such a system is probably the vast majority of the cost.

Stuart Breslow: With respect to the costs, it's no longer effective to scale linearly to handle the exponential increase in market data rates. New approaches really need to be taken. What we also find is as the data rates increase so do the telecommunications needs, so the telecommunications expenses, in addition to the overall sheer processing, are probably one of the key elements to the overall costs increasing. Couple that with the fact that we do have more markets coming on the scene, both in the U.S. and Europe with respect to equities, options, futures. All this fragmentation is actually contributing to the need to have more telecom and then if you take into account the need for redundancy you then have multiple feeds into multiple sites. If you want inter-site and intra-site redundancy, one telecommunications line actually results in four telecommunications lines for that maximum redundancy. All that really adds to costs.

Kevin McPartland: I agree with Stuart. Horizontally scaling in many ways is just not the best approach anymore. Data center space is expensive, power is expensive and there just needs to be a more efficient way. So firms are looking for ways to vertically scale, whether that means buying faster processors, in some cases using hardware acceleration in order to have each box do a little bit more than it did before, and then looking at the bandwidth side. Tabb Group is in the middle of a study about sell-side infrastructures where we spoke to about 25 sell-side firms about their infrastructures and their biggest concerns and issues about their infrastructures and when asked how they were coping with increasing needs for bandwidth, a simple answer but everybody's response was "well we're buying more of it." So you're absolutely right there's a strong demand for the telcos, not just on local metro areas but on long distance routes between major market centers, and finding the best paths and big enough pipes to handle the data you need is an ongoing challenge.

What are some of the key areas of focus to ensure that a company is efficiently handling data?

KM: There are a number of things to look at. Prioritization is important—what data needs to be low latency? There are ways to segregate data flow for less latency sensitive and more latency sensitive information so data that's going to a market data screen for example, going to a point and click user, maybe to an operations professional who needs real time but not sub-millisecond data—there are ways to segregate your network to do that more efficiently. One of the other key themes from our sell-side infrastructure study was virtualization and that virtualization is becoming very popular amongst the sell side. However there is an inherent latency cost to using virtualization technology so I think in the next couple of years we'll see the providers of virtualization technology as well as the financial services firms themselves looking for ways to use virtualization technology to create a more virtual environment that allows them to efficiently use their hardware without adding so much latency that it gets in the way of trading strategies.

SB: Kevin definitely points out a couple of key areas to look at. If your focus is on making sure you have the capacity to handle every tick and that you do focus on delivering ultra low latency where it's applicable, you really need to go back and look at the way your solutions are architected. You talked a little bit about Amdahl's Law and it really comes into play here. Making sure that your processing is as parallel as possible is paramount and most developers are not necessarily taught in computer science programs to think in parallel. It really does require a new method of thinking; the old brute force method doesn't work any more. Another area that has tripped up a lot of firms is with respect to the middleware. As the market data rates explode your middleware--we all have distributed computing and it's all tied together by middleware--so you need to optimize the middleware to keep pace with the exploding rates. If by chance you have a single slow consumer or a slow computer processing some of that data, it can kick off these multicast storms which will just complicate and disrupt the overall processing.

FB: If we're looking at efficient handling of data we're looking at two parts. One part is to make sure that we're not necessarily paying for duplicate data but that you can process

a single data stream and distribute it to multiple types of consumers. Obviously the low latency/high volume applications will be the top-tier applications consuming data right off—for example the feedhandlers, whereas display applications or less demanding applications could also source from that exact same source. So by re-using the same ticker plant for many different applications and different business units we both curtail the bandwidth requirements that we discussed earlier and make efficient use of the feeds coming in. As Stuart just said about the middleware aspect of it—again quality of service comes into play there. If we're looking at a messaging API or messaging product, like IBM's LLM, we're able to deliver messages both at high capacity with ultra-low latency using a variety of network fabrics. The messages are delivered to those high demanding applications, at the top tier, using multicast traffic. That same data stream is also used to deliver messages to, for example, desktop applications and LLM is able to efficiently handle those negative NAKs that Stuart mentioned. Certainly NAK storms are with LLM for example a thing of the past and are well under control using suspension and expulsion of consuming applications.

How can companies best plan to have capacity for the continuing growth in data rates?

SB: It comes down to just having a good discipline about measuring and monitoring. You really need to know where your headroom is with respect to your bandwidth as well as your own processing capabilities, and then regularly measure to see how close you're coming to your headroom, and then decide on when it's time to further scale. We use some technologies in order to capture the peak data rates at each tier so that we know exactly how close we're coming to our overall utilization of our overall headroom. And then what you need to do is come up with statistical forecasting techniques. The Financial Information Forum (FIF) comes out with their overall projection and it's used as one element to the overall forecasting, but I think everyone needs to take it upon themselves to forecast with respect to their own capabilities and their own processing infrastructure.

DB: Very good point Stuart. From my experience the derivative and equity exchanges that often drive the overall processing workloads are not all very good at predicting their own volumes so you really have to look at their predictions with an additional level of analysis. Some of them have exceeded their published thresholds and others came very short of their published ceilings.

KM: I agree with Stuart that monitoring is critical here. From talking to the sell side about their equities infrastructure utilization rates they keep them around 30%. So that means 70% of capacity is essentially sitting idle during the trading day in the event of any market spikes. I think everybody was scared a year ago when we saw some of the crazy volatility numbers and market data volumes around the time of the Lehman Brothers bankruptcy. But that utilization level comes at a cost. There's a multimillion dollar cost to keeping that level of excess capacity but obviously as we're seeing it's necessary. So firms are going to be looking for ways to keep that excess capacity without spending so many millions to have those servers sitting there waiting in case something goes wrong, and this is both on the bandwidth side in terms of having enough fiber coming in to your

data center in your network and also on the compute power side that you have enough CPU power to process the data as it flows in through the big pipes.

FB: A lot of that comes down to communication—having the different groups and entities communicate about what capacity demands are going to be. Certainly FIFO is an input. In addition, there are business groups that continue to develop different algorithms with different input requirements and may place a different load on a ticker plant from a consuming perspective. So really it's up to an operations team to juggle all the aspects of it, both from an ingress and egress, and if a vendor like IBM for example is contracted to help monitor all that, to help share that information and make sure that capacity planning is done well in advance. So in that regard I think it's very much about communication, making sure that all the people are exchanging information about what's coming up, what's expected, both from an ingress and egress, and then plan accordingly to make sure the capacity is there. And that may go into some of these technical advances as far as what's possible to make the ticker plant as efficient as possible to be able to handle those loads.

DB: I would add that people in our industry such as Interactive Data have to also have a very robust process in place because this is the kind of thing that cannot be done once, cannot be done occasionally, it has to be done continuously with well-defined procedures.

What are some of the emerging technologies that should be explored to help companies stay ahead of the market data tidal wave?

FB: Lots of things are going on anywhere in the stack, from operating systems as far as real-time Linux and real-time operating systems vs. regular operating systems. There's a lot of development being done on the networking side, the network fabrics, you mentioned earlier Infiniband and 10 Gigabit Ethernet and of course the associated protocols that can be run on top of that. And then certainly the line gets blurred between those typically separate systems, the server and the network. The network technology becomes part of the server with more advanced network cards. TCP Offload may have been a precursor to that but that's certainly advancing into some of the feedhandler functionality and placing that onto or closer to the network card to have it do some of the processing in line if you will. And certainly that's something that is evolving over time and is being implemented in production today.

SB: Falke pointed out a couple of key ones. I'll just add in a couple of additional items. Within RealTick, we specialize in distributing data into a lot of interconnected clients all the way around the world, so what we've put some focus on--in addition to just processing it correctly through our own infrastructure--is how do we manage the wide area fanout as effectively as possible. We've actually looked at some specialized content routers to help with that. And then the last one is as you have exceptionally high data rates traversing through your distributed system and then being disseminated out, you can be much more bandwidth friendly if you focus on compression, so we've looked at some specialized devices to help with overall compression.

DB: I must say the last one in particular is something we at Interactive Data have explored as well including hardware-based compression methods for network bandwidth savings.

KM: In our study we asked what people thought is the most disruptive technology in financial services right now. When we asked the question I think we expected to hear about new chips and different widgets that do amazing things that we never heard of but the top response was actually the availability of bandwidth. That it's widespread—you can see that in your house with your home connection, how much faster it's gotten over the last two or three years. That same phenomenon has hit the enterprise as well. There's a lot of extra bandwidth available after the dotcom bubble burst in 2001 – 2002. So much of it was laid that then was given up because many of those companies went out of business. So prices are actually going down and there is some competition for key routes as I mentioned earlier between key cities but bandwidth really is available and prices are dropping. And the fact that it's out there and companies can get their hands on it is really something that's going to change the game and allow market data rates to continue to grow because people are able to consume them.

What are some other things to examine to ensure that companies' overall market data management is as efficient as possible?

KM: Some of this comes down to what I said before. It's about priorities and making sure you're doing what your trading strategies require, so if it's a traditional asset manager with a long-only strategy, then cutting microseconds off the time it takes to pull the datafeeds in is not necessarily critical, nor will there be a return on investment in that. If you're a high-frequency trading firm on the other hand that is doing market making or some stat-arb strategy then that is your business and it's worth spending nearly every penny to make sure you're faster than the competition. So it's not just technology for technology's sake—that is never the answer, especially not in financial services. It's really about understanding what your needs actually are and having a process in place to manage those flows and manage those needs.

FB: I agree with Kevin's point. Avoidance of duplication of datafeeds is certainly one that can be key. If data is already ingested into a ticker plant for high performance applications, there's no reason that data can't be forwarded on to less demanding applications downstream. So being able to route that through an efficient messaging infrastructure to all types of consuming applications will bring certainly the efficiency up and the overall cost of market data infrastructure down.

Are there any new technologies or tools that are helping with market data administration?

SB: There certainly are. You know, since RealTick has been around for quite some time we typically mostly use homegrown tools and specialty purpose tools for things like latency monitoring, data quality, sequence gap checking, crossed and locked markets, things like that. But there are emerging on the scene--some of them not so recent, some a

couple of years ago--some of them that are new that help with a much finer grain bandwidth measurement as an example, or do things like packet level data capture for phenomenal post troubleshooting, and now there are a couple of relatively new services that are on the scene that actually help with latency monitoring. So you're starting to see the emergence of a lot more available third-party tools and solutions to help with the overall market data management challenge.

DB: I will also add that in addition to the quality-type tools that Stuart mentioned there are also increasing changes in how the permissions are structured and contractual obligations evolve so it behooves people to manage market data—to invest in flexible strategies, whether in-house or outsourced.

FB: Certainly from an entitlement perspective it's critical to ensure that you only pay for the data that you're using. As you take in direct feeds vs. a consolidated feed it's certainly critical to keep track of that so we're not paying for all consuming applications for that data stream and can accurately keep track of them at no or very little overhead. So the entitlement processing goes hand in hand with what I mentioned earlier about the duplication of datafeeds to keep costs under control.

DB: I will add from our point of view we see increasing interest from our customers in more flexible permissioning and entitlement systems as well as giving them more control and some of the systems are becoming more popular from Interactive Data.

How are the technology changes affecting the hosting versus local processing of market data?

KM: That's a great question. A lot of it comes down to co-location as well--is it worth spending the premium to put your servers in the same data center as the exchange, and again as I said earlier it's important to understand what your trading strategy is. How critical is it to cut those milliseconds and to have complete control over all of those independent datafeeds to be sure you're getting them as quickly and accurately as possible? It's also a little bit about the scale of the firm. Is there a return on investment to spend money to create an infrastructure to handle multiple direct datafeeds or is it more efficient to look to a provider that has expertise and specializes and essentially spends all of their time figuring out how to solve these problems? In many cases that's the better option for those firms, rather than trying to hire or find all of that expertise in-house. But this in a lot of ways is the classic buy versus build decision that's always a complicated one for IT managers. I think for market data the decision comes down to your trading strategy and the return on investment for creating the infrastructure yourself or looking to an outside provider.

FB: From a hosting versus local perspective, it's key to match the requirements with location. If you need ultra low latency and you need to be able to co-locate with the exchange and the market data there, cross connect over a 1 gig or 10 gig cross-connect and get the absolute lowest latency. You want to make sure you have the flexibility and technology in place to do that by leveraging the same ticker plant that perhaps is

processing the data there, having it forwarded to your local hosting facility as well for downstream applications is a cost-effective way and having the best of both worlds. So rather than making the decision either/or, you can basically do both. You do the processing at the co-location site for the high-demand application, and then use middleware technology for messaging to pass it over the wide area network to your local facility for further downstream applications that have perhaps lower volume and can accept higher latencies and have those processed at the local data center. So given the right technologies, it's not an either/or option, you can actually do both.

SB: The technology changes that have evolved over time—it's no longer a one-size-fits-all, it comes down to how you need to consume market data. Depending on the nature of the business, the technology changes have evolved to have a solution that satisfies whatever your business needs are.

One of the things we discussed earlier was that in order to take advantage of parallel processing you actually have to in many cases redesign your program or algorithm. How do you train or hire developers for parallel processing?

SB: What we did when we started to look at this is just a couple of key things. One of them was we have a staff of very talented developers but most of them have not really been taught to think in truly parallel fashion or look for the parallelization opportunities. So we brought in one of the major microprocessor companies to conduct a customized training program we forced all of the developers to go through. That was a huge win for us. And then in terms of adding in new developers, we started targeting more the computer engineering programs instead of the pure computer science programs and we partnered with some of the local Chicago-based universities that specialized in computer engineering.

KM: We talked about this a little bit earlier but the human mind doesn't inherently think in parallel. It starts back in universities. When I was in school, computer science, learning how to program, we learned about threads and all but understanding parallel processing wasn't a common theme. There weren't multicore chips back then. I think schools are starting to get there and see that there is a demand in the marketplace and I'm sure Intel, AMD and the others are happy to supply the major universities around the world with their technology to help people learn how to use it and be most effective with it and I'm sure Visual Studio and any other development platforms are working on tools to make the process even easier. So almost to help programs parallelize themselves. So it's going to be a process but we'll get to a point in a couple of years where there will be a critical mass of expertise generally and on Wall St. where this won't be as much of a problem any more it will be common knowledge as to how to do this.

DB: My guess is because of the preponderance of von Neumann architecture trained engineers and furthermore business algorithms we'll still see a lot of it in a few years and just like with other techniques you mentioned earlier you'd have probably very targeted applications where you do very highly parallel processing for areas that are critical and still leave the sequential processing for most of the problems.

FB: In the case of IBM it's not a training or hiring exercise it's more applying the established experience on, for example, the IBM research team who have been working on many parallel problems for a very long time. We have software products that do massive parallel processing, we have supercomputers that people have worked on so really at this point what's being done is taking some of that experience and applying it to the financial industry. With IBM and the Websphere Front Office product suite, what we're doing is taking some of that experience, applying it to the market data and seeing how we can solve this particular challenge in this particular industry. So it's really an application of existing experience to this particular market.

What technology advances are being applied to High Frequency Trading?

FB: I think we've touched on many of those. I think one of the key ones is the convergence if you will of distribution and messaging with the applications and the processing of market data. Certainly typically messaging and the network was looked at independently from the application and really the line between those two is blurring at this point. We're seeing a blend of the processing—for example the line arbitration or line processing of feedhandlers being pushed closer to the network infrastructure while still keeping the business logic and the applications so it can be flexible and adapt to particular needs. So I think that's a key part that has taken place. And having the technology in place to adapt to those changes, while that may be one Gigabit Ethernet first, then going to 10 Gigabit Ethernet or Infiniband and 40 Gigabit Ethernet down the line, it's critical that that underlying infrastructure that the application might be built on is flexible enough to adapt to these new technologies so the application doesn't need to change.

KM: First it's important to point out that not all high frequency trading is necessarily low latency. It sounds a little counterintuitive but there are certain strategies where they're looking to profit from certain market inefficiencies in, say, options or futures where the prices don't tick that fast, so a couple milliseconds of latency will not impact the strategy's success. And on the flip side, not all low latency trading is necessarily high frequency. If you have say 10 orders throughout the day and they need to get done immediately then there's a need for cutting latency as much as you can, although that's not a high frequency strategy. So I think that's important to point out, even when we're talking about technology. One thing that I think is interesting that I've heard recently on the high frequency trading side involves messaging middleware. Inter-process communication is certainly a big one on the high frequency trading side. Even trying to cut out the middleware, making things closer than they were before and keeping everything on the same box and talking whether it be through memory virtualization or just the process of talking directly with each other through some proprietary communication medium created by the firms. I think we're at that level now where we're trying to figure out where a couple of microseconds per second could come from the high frequency, low latency guys. It's fascinating to see how fast and sophisticated they've gotten.

DB: Some of it is really good, old fashioned software optimization and we've found for instance things like keeping things in memory, buffer copy elimination – the kinds of things that can kill you even within one processor if you keep copying things around. Interrupt distribution is something we found very valuable. The more processors you have, the more careful you have to be to avoid bottlenecks in any single processor or processor core, so distributing the interrupt load is really a key part of high performance computing in this space even if you are avoiding the middleware.

KM: It's interesting--programming has sort of come full circle. When I was learning Java was very cool for the Web but it made everybody a little bit lazy you don't need to do garbage collection, you don't need to declare a lot of things – it does everything for you but there's a latency penalty for that. And of course Java has gotten better and faster, so we're seeing, as you said, some traditional software optimization methods and you need to understand memory allocation and interrupt handling and all those other things.

SB: Those are also very good points. I think that optimally, we talked about inter-process communication, merging everything into a single process obviously is the most effective but I would say it comes with its own challenges. We're talking about just general advances. One of the ones that we've all grown to accept as pretty normal now but it is an advance in the last year or two is the whole concept of co-location. Several years ago it was cutting edge and now it's become a little bit more mainstream and it's certainly a business process advance. And then the other advances that I think are paramount to high frequency trading is that the trading algorithms are now starting to move a little bit more towards a silicon-based type solution. You know FPGAs are getting a lot more mainstream, the toolkits are a lot more readily available; GPUs are being used a lot more regularly so we're starting to see a bigger adoption just of the pure microprocessor or processor-based calculations rather than doing it in code in an operating system.

Is technology driving the business strategy, or business strategy the technology?

KM: That's becoming a more and more common question. I think the easy answer is it's a little bit of both. It depends on what business you're talking about. If you're talking to one of the top high frequency trading firms, then I'm pretty sure they have ideas and concepts and they're waiting for the technology to catch up so they can implement the trading ideas that they have. On the flip side, if you're talking to maybe a retail broker or we talked before about a traditional asset manager, I think in those cases they see technology come into the market and think about ways they can use that to make things more efficient. If you look back to when equity trading in the U.S. started to go electronic, I think that was a case of people sitting around a table saying, we have this great technology, we have this new thing called the Internet, what can we do with it, how can we make trading more efficient and that's brought us where we are today so it definitely goes both ways.

DB: My own thinking is that although it is always both there are probably periods when one is driving the other more often. I think during transition periods like now technology probably drives business a bit more than the other way around. I do feel that trading

strategy the last couple of years has been undergoing quite a bit of transition to use some of the tools that traditionally have been used only in supercomputing—HPC and so forth.

FB: I think technology and business leapfrog each other. I think it goes in waves. Business may lead technology and then technology may lead the business. Certainly the current developments on messaging and infrastructure as far as network technologies is another example where vast amounts of bandwidth and extremely low latency are possible between machines and brings up interesting possibilities as far as using that data.

In developing new tools, should firms buy or build?

FB: One of the things you have to consider when you're building these technologies is that it requires a lot of expertise. I think it's clear from the discussion this far that the days of just writing a feedhandler using software only are long gone. It requires deep technical experience in not just software development—parallel or not parallel—but also experience in networking, drivers. You mentioned interrupts and being able to handle those correctly and the right CPUs. It requires experience in the networking aspect as far as being able to use 10 Gig Ethernet and even RDMA native verbs, for example. So I think the ante has been upped significantly in developing a competitive solution and in order to stay current with those evolving technologies it becomes very difficult to build those in-house I think and stay competitive in that regard. It just requires too much detailed expertise in too many different areas that are not the core competency of most of our typical customers.

DB: I would have to echo those remarks. In my experience we do exactly that--we have to look for the best technology solution which tends to require experts in the field but also we have to continuously evaluate areas where we do have to do things in house because outsourcing them doesn't give us the competitive advantage so you have to keep looking for those balance points.

Where do you see functional programming entering into this area (F#, OCAML, SCALA, LISP, etc.)?

FB: From an application perspective, there are different types of applications that would be using market data and I think I mentioned earlier that you want to make sure the single ticker plant can feed all of those applications. A functional and programming language like C# or F# or maybe one of those examples and those can be consumed to the API that those ticker plants provide.

DB: My experience is that many of these languages are good for specific types of problems but there is a vast majority of business applications right now so hungry for processing power that they seem to be going the other way around – echoing Kevin's remarks earlier high-level languages that do many things for you actually at a bit of a disadvantage right now, perhaps we'll see another wave in the future.

FB: Analytic applications are being developed to deal with current market data streams and those are in a sense very high level programming languages in which you can define

rules to reduce stream processing, for example, without actually doing any coding per se. So if you look at, for example, the InfoSphere product from IBM where we can consume market data coming in and using a proprietary language, that is a functional language you can define business rules and have the processing of them parallel across a distributed system. So that's a step beyond the typical functional programming language and more of the business logic can be programmed in that way. So I certainly think there are significant developments going on there and I think that's an area for example where the technology is driving the business. Technology is making those possibilities available and I think the business is catching up, figuring out how they can make the best use of that.

SB: I have one more point to add, which is really just a practical aspect to this. As we talked about this landscape continuing to evolve, there's a practical aspect to it. Everything that we write needs to be changed and evolved, and when you start to get into some of the more esoteric programming languages the availability of resources starts to wane a little bit and makes it a little bit more challenging to continue to evolve. If you stick to a mainstream programming language like a C++, C# or Java, the resources are much more readily available.

DB: Thank you very much Stuart that's quite a valid point and I think that probably reflects the reality quite a bit. And that's our last question, thank you very much panelists.