Thought Leadership Series
Transcript

#003: High Speed Market Data (Frank Piasecki – ACTIV Financial)

Candyce Edelen: Welcome to this installment of TrendSpotters. I'm Candyce Edelen and our topic for today is "High Speed Market Data." My guest is Frank Piasecki. He is the President and Co-founder of ACTIV Financial Systems. Frank has 22 years of senior management experience in the market data industry.

Before ACTIV he was at AT Financial for 11 years and there he spearheaded their institutional sales programs and their global customer service group. Frank also currently serves as a board member at Clearwater Management Company, which is a privately held asset management company. And recently Frank was recognized as a member of the Market Data Hall of Fame by Waters Technology. Frank, thank you for joining us, we're excited to have you on Trend Spotters.

Frank Piasecki: Thank you Candyce. I look forward to our discussion.

Candyce: Me too.

In 2009 and 2010 there was a lot of focus on reducing market data latency to the lowest possible level. Clearly, this is very important to the ultra low latency and high frequency trading shops. But really, to compete in today's electronic markets even the more traditional players have had to reduce latency.

This is particularly important in their market data infrastructure. So Frank, let's get started by talking about the key attributes that companies of any size or any strategy really need to be considering in defining their market data needs.

Frank: Well, it's an excellent summary of where we are and I would point to the age old problem that most folks start out with, which is “what is the required content.” If you've got a work flow that you've got information needs to satisfy the coverage and regional terms and in asset classes is a primal element in deciding and setting a strategy for any market data solution. You move very quickly as you point out now to how quickly that data arrives at your facility and into your applications and important to that is deciding the inventory of what you want.

Is it an eyeball user community you want to serve? Are you dumping the data into systems? Are you driving trading systems? Those all speak to the nature and type of interfaces and the delivery architecture of your market data content.

© 2011 PropelGrowth, LLC

So if you're buying content how fast does it have to be and what kind of interaction functionality is necessary? And you sum up all that into defining the application that you have and where it physically is located. This has become critical because in many respects, low latency strategies are satisfied very much by being inches or feet away from the matching engine.

So if you actually need market data, driving applications that are co-located disparately around the world at many trading venues, you've got to say so. And that's got to be the determination. So these all add up to an inventory of what you need where, how fast.

Candyce: OK. Thank you. So it's really about what are you trying to do with it and what do you need. So the push to enable high frequency trading has created this constant push to reduce latency to the lowest levels. Even to the point where people are starting to talk about zero latency which I don't think is really all that possible.

But to achieve the levels we see now, specialized platforms and infrastructure has been created across the market. What do you see that users can expect from these more specialized platforms?

Frank: Well, firstly that any approach, any technical approach has to run into the rules of physics, right? There's really truly a limit and you're exactly right, zero latency is not exactly... it's theoretically achievable. I don't think anyone will pay for it. So you've got limits of how fast data will move through fiber gear, you've got the costs to achieve these.

So if there's a law of diminishing return you better know about it, and you better know how your strategy will play against it. If it's constantly and necessarily the fact that you need to be fast and first, then clearly you can spend that money and it'll be justified. At some point there's a real business decision to be made there.

I would say that the most over looked as ACTIV deploys or as other vendors deploy these specialized technologies, you've got to understand that lots of other things are in your application stack that are likely improvable, not just market data. So it's being aware of the context of latency games, where they can be had and for how much.

Candyce: So we see the high frequency players really demanding that ultra low latency and being able to invest in the kind of infrastructure and costs that can get them there. But the broader more traditional players in the market don't need that ultra low latency for competitive differentiation. So what are some of the issues that these traditional players are facing and how different are they?

© 2011 PropelGrowth, LLC

Frank: I think those folks constantly have data breadth, asset class coverage needs that have been growing as investable assets across the world have been growing, both in instrument count and in physical location. The challenge with a lot of that and the growth of derivative instruments has been really data navigation. So you've got to get the underlying instruments wherever they are, which is a big aggregation challenge, which is a traditional market data function. And then there's a lot of linking data to relevant keys so that you can use it appropriately.

If you're talking about certain asset classes, reference data becomes extremely important, much more so than really latency. So going global and going across asset classes is really... I think that traditional groups' continual challenge. Always has been and will probably remain so. Cost in the face of higher message rates, so these folks are not maybe caring about the fastest access to data, but they are impacted by high frequency strategies that create a lot of data.

So the quote to trade ratio or trade to quote ratio, whichever way you want to look at it, has a very big impact. Because they have to build systems that handle the information flow in these liquidity venues. That means they have to buy more bandwidth, bigger switches, larger computers and all that capacity is very important. So their challenge, even though they may not need low latency, they're challenged by the fact that the low latency guys are creating more data and they need cheaper systems, or strategies to keep that impact as little as possible.

So filtering technology, conflation technology, those actually are real survival tools that they need to rely on that the technology stack needs to address.

Candyce: That affects any kind of player in this space, right? We're not just talking about the biggest firms or the most global, but all of the players are running into those issues around bandwidth and being able to handle the throughput that's necessary?

Frank: I think so. I think most people would be surprised at what size and scale firms operate globally now in the capital markets. Very modest shops now do trading between regions, Europe and Asia, North America and South America, there's a lot of feeding between derivative instruments that just happen in different regions and across, physically, many, many, many different areas. So there is an element to the scale of globalness reaching down into much, much, smaller consumers.

I would say that infrastructure built up generally by the industry to offer access to speedier sets of data also means that it's more broadly impacting the quality and speed of data delivered to people who don't necessarily need it.

© 2011 PropelGrowth, LLC

Candyce: Sure. Sure. Which helps everybody in the end. Now, what about these large global banks and institutions where they've got to consume and aggregate that data like you were talking about from that broad range of sources. But they've also got to distribute it to a broader range of consumers in real time. How are their changes different than some of the smaller and perhaps more nimble players?

Frank: I think you hit it straight on the head. Larger institutions have broader requirements that not only mean that they have to source an aggregated set of data with particular latency characteristics into an application or a set of applications. But typically they've got to source internally generated data and they also have to collect that and distribute it to a very wide group of consumer applications that may have nothing to do with each other and have very different views on information.

So the challenges of the larger banks and institutions globally have been very different, but they also have high frequency strategies that have been embedded or that they service. So it's not that they don't care about low latency, it's more that they care about that and other data aggregation challenges that may not have latency as a primary requirement. The ultra specialization of use cases within the enterprise really mean that you've got to have broad platforms, or you have to run lots of different solution sets.

And I think organizations are always looking to get rid of the complexity around market data and if you can develop different strategies where the dependency for different use cases is solved by single or unified platforms – that's extremely attractive.

Candyce: So Frank, the global markets are becoming more and more interdependent and this seems to be affecting companies regardless of how global or how large they are, because the market data infrastructure in the different geographical areas seems to be varying widely. What are the most challenging issues you're seeing, let's say facing European companies right now, or companies that are trading in Europe, where they're dealing with new issues that are being created by method and lack of data aggregation?

Frank: I'd say probably in Europe the number one issue we're hearing about from customers is the consolidated tape issue. Getting a unified look across competing liquidity venues, dealing with various exchange rate regimes, making very disparate trading rules, or trade condition normalization established, is a very serious challenge to those folks. So there is in Europe a lot of generic, how do we make data more widely available so that we can get a fair and balanced view of the market, but allow trade competition between venues to help us attack the costs?

© 2011 PropelGrowth, LLC

That's been a challenging issue and will continue to be one. I think finally in Europe where lots of different MTFs cropped up recently and attracted electronic trading strategies, we saw data rate growth, which we've seen in North America for some time. But now it's starting to show up there in earnest.

So people are literally starting to deal with, wow, that's a message rate growth that I didn't anticipate. It's forcing me to buy different equipment and more bandwidth. And when does this stop? And that usually presages the need for a pretty sophisticated solution.

Lastly I'd say that challenge has been symbology. There's a proprietary symbology that's deeply embedded into the European trading system, much more so than anywhere else in the world. Of course that's the Reuter instrument code issue and I know the regulatory authorities are looking at that. And I think that remains something to be solved.

Candyce: I saw an article, I think it was last week, that was talking about the struggles with aggregating all of the data in Europe and suggesting that standardization that's so desperately needed is going to take a very long time. What's your crystal ball say?

Frank: I'd agree. I think there are pretty profound regulatory differences that remain. I think that there's a cost of switching. Standardization is always very good, but it is going to take some time and planning. And I think you have to pick which issue you're specifically talking about. But in terms of real-time exchange-traded data, there's a lot of variety out there amongst many venues. Now, it's being aided and abetted by exchange consolidation, if that does occur at pretty large scales as we've heard recently with LSE and TMX and of course the Deutsche Börse/NYSE. So it may be solved somewhat through market events that we're seeing now. But in the end it's going to take some time.

Candyce: Now Asia's got very different problems. Part of their problem, my understanding is, that they're more geographically diverse. They also have very different regulatory bodies, different levels of automation. What kinds of things are you seeing as the biggest issues firms are dealing with in Asia?

Frank: Well, precisely. The sheer distance between venues is just enormous. So the cost of communications, the language barrier, the different regulatory requirements of operating businesses in different regions or different countries, the complexity there is big.

It's not a market that has had a lot of emphasis on transparency, necessarily. There has not been the kind of electronic trading facilitation that's been pursued in other regions. But they're picking up and they're picking up fast.

© 2011 PropelGrowth, LLC

You've seen with Tokyo, with Arrowhead and the opening of the many PTS' there recently. That's created an information challenge that market data folks have to address. So the scale, the opening and the relevance of these folks for electronic trading arbitrage, now makes these problems somewhat important to address. And that's going to hit basic infrastructure cost and aggregation cost.

So I think this is going to be interesting all around, but they're really starting it. It's early on the curve, but I think like in most things you'll see the Asian region pick up these subject matters, these capabilities and move very quickly.

Candyce: So looking back at Europe and Asia, what are some of the things that firms are doing to address these issues?

Frank: First and foremost they're trying to determine which venues are attractive for what kind of strategy. So, if you're talking high frequency, there's not a whole lot in Asia. There's many more in Europe that have opened up. And what's very clear is that you need to pick your strategies around solutions with partners who are there or going there.

And by that, I mean you have a physical element of technology that has to be implemented. There's infrastructure and architectures that need to be deployed. And do you have a relationship or internal capability with folks who've done this before?

As we've been doing this over the many years now, it's clear that there's a significant amount of local knowledge that has to be brought to bear and then a significant amount of infrastructure that you have to deploy. And then you've got to run a good data service.

And that requires knowledge of the different instrument types and sourcing non real-time data to make that relevant and accurate. And then developing quality control capabilities across the world, in so many different languages. Again, so many different regulatory regimes, it is complex. So you really need internal knowledge and partners who really have expertise in this.

Candyce: Right. That makes all the difference.

Let's jump now to a very popular topic, hardware acceleration. It seems that this concept and FPGA is even on the tongues of CEOs of banks now—nobody knew what it was a few years ago. Are you seeing a lot of firms adopting it?

Frank: Well, everyone's talking about it. And I think it's typical of sophisticated technology solutions. What's very amusing about the emergence of this technology is that most people take the face value attribute of lower latency.

© 2011 PropelGrowth, LLC

And whether you're talking about FPGA strategy or other hardware acceleration strategies, you're kind of in an ASIC world. You're doing specialized programming on specialized devices to move data through a stack quicker.

And general purpose CPUs with standard operating systems worked very well for market data for a very long time. And in fact, will continue to do so. The problem is that there are points in this data-flow that truly benefit from some engineering effort in alternative technology.

So we've found, operating a global infrastructure where our enemy is message rate growth, and of course that's what the market's going to do and continue to do as long as computers are involved, that hardware acceleration has been more than anything, a footprint reduction. So there's an efficiency of building data services. And if you're using less horsepower, less space, less electricity, you've got a more simple, less complex computing environment, then you're going to be a low-cost producer of such service.

I think not well understood, is that FPGAs and the like offer extreme advantages in reducing footprint. And that, I think, as an operating entity that's got to pay bills and be profitable to do R&D to get into the next technology stack, we need to make sure that we're making money doing what we do. And we do that by leveraging hardware acceleration for its operating efficiency on top of its latency benefit.

Candyce: That's an excellent point.

Frank: You have to understand what happens inside an FPGA. Are you doing normalization? Are you creating states? Are you facilitating access to BBO information? What's going on inside that box? And knowing that helps you discern the differences between the technology solutions.

So it's one thing to say, hey, we're using an FPGA. What you're actually doing on it can be profoundly different. It's very simple to do a table look-up in FPGA technology, which is extremely relevant to things like pre-trade risk which is very important for direct access now. For no matter how fast your market data is, if you can't trade it because you're prime won't let that trade go through, you've really got a problem.

So we are seeing FPGAs being levered for those very simple capabilities, but they're quite powerful as well. So I would sort of cover lastly on the fascination with hardware acceleration: generally these technologies are only as good as the surrounding architecture.

© 2011 PropelGrowth, LLC

So broadly, market data as a topic needs a lot of things. It might have expanded or contracted over the years. But ultimately you need a co-lo hosting capability, you need networks globally, you need gateways into trading venues. This is all on top of the technology footprint, the ticker plant, feed handlers, and middleware underneath all of that.

Then you need data management support for quality service against the data quality. You need field technicians, phone folks. And then you've got to do that around the clock.

So point solution technology is great and that it's being led by a bunch of innovative small institutions, that's fine. But you've got to remember, if it's going to be useful, it's got to be sitting inside of a much broader set of services. I think that's critical. It's great to be excited about a technology—how is it actually deployed and supported?

Candyce: So we're starting to run out of time. But before we wrap, I'd like to hear your predictions for where you see the focus changing in 2011 and 2012. Can you quickly kind of talk about what hot button trends you predict?

Frank: I think a lot of folks are going to be looking to get out of infrastructure cost spiral, which is that they step into the low latency game with one or two markets, realize and develop strategies, successful ones that then need other latency values on other markets. And then they find themselves in the mini ticker plant infrastructure management. And then they go, well wait a minute, I can't do this globally. I don't have enough qualified people and I can't keep them.

So I think it's going to be a big effort to kind of find technology solutions, whether they are managed solutions or they rely on managed infrastructure, or you go to your traditional large scale vendor for these solutions. That really remains to be seen. But we believe that that's really the issue.

Mostly SAS, ASP, Cloud, whatever you want to call these, all in demand services will play a larger role. I think this year we'll see people kind of trying to virtualize access to information services, certainly for market data.

We call this centricity, but there's going to be a network centricity to solutions that has really not been there in the past with a lot of different solutions that have been deployed.

So delivery, an all encompassing platform that can address many issues, these will be important to that design.

© 2011 PropelGrowth, LLC

Candyce: Well it looks like we are out of time. But Frank, thank you again for being my guest. This has been a very interesting conversation and I have really enjoyed it.

Frank: Well thank you, Candyce. I enjoyed it as well. There's so much more to talk about. But I think this is an exciting time in this industry. I've never seen so much change happen so quickly, and I think it's a great time to be addressing these challenges.

Candyce: I agree. My guest today has been Frank Piasecki, President and co-founder of ACTIV Financial Systems. If you have questions for Frank, you can join the TrendSpotters group and ask him there. You'll get instructions about that in a moment.

Or you can email Frank at frank@activfinancial.com. That's frank@activfinancial.com.

I'm Candyce Edelen, CEO of Propel Growth. And I look forward to joining you again next time on TrendSpotters. Have a great day!

We invite you to join our discussion of this topic at the TrendSpotters Thought Leadership Forum. Join the conversation at www.trendspotters.propelgrowth.com. TrendSpotters is produced by PropelGrowth, helping companies create competitive advantage in the capital markets. To contact Candyce, email her at trendspotters@propelgrowth.com. To hear more of the Trendspotters series, visit www.propelgrowth.com/trendspotters. To find out more about PropelGrowth, visit us on the web at www.propelgrowth.com.

©2011, PropelGrowth LLC. Contents of this transcript and original recording are the sole property of PropelGrowth and may not be copied or repurposed without our express written consent. TrendSpotters is licensed under a Creative Commons attribution/non-commercial/no-derivative works 3.0 license.

PropelGrowth is a trademark of PropelGrowth, LLC. Progress, Reuters, Currenex, Hotspot, CME and any other trademarks contained herein are the property of their respective owners.