

CEP02 – Daniel Chait & Ross Hamilton, Lab49 (pt.2)

Announcer: Voices in Business presents the Complex Events Podcast, sponsored by BEA Systems. Show No. 2.

[music]

Announcer: Hello and welcome back to the Complex Events Podcast. Today, we feature the second half of our interview with Daniel Chait and Ross Hamilton of Lab49. Daniel and Ross continue their conversation with Vicki Zesses, discussing the adoption of Complex Event Processing [CEP], how buy side demands drove its use to solve the needs for algorithmic trading in 2006 and 2007. And what they expect will be broad-based adoption across the enterprise in 2008, particularly in the areas of market data, real-time pricing, fixed income and enterprise-wide credit risk management.

Vicki Zesses: Daniel, you spoke earlier about how we got here, to the point of CEP. As you're working with various financial institutions, what are you seeing, how far along is the market in adoption? Are there places where it could be used and it's not yet being adopted that you see value?

Daniel Chait: Absolutely. I mean, that's a great question. If you look at especially where we are right now, I think, it's an incredibly interesting time for CEP technology. If you look back, let's say, to 2006, I would say that the vast majority of practitioners within financial services, IT, had not heard of any kind of CEP engine whatsoever, or if they had, it was probably within the context of a specific point solution that they were working with.

Ross Hamilton: And probably something like algorithmic trading, which is the classic use case.

Daniel: I think, certainly there were kind of leading areas, like algo trading, that really came to this stuff first. In 2007, I think for the first time, it kind of bubbles up to the popular consciousness. You started seeing a number of different forums and industry conferences dedicated to specifically CEP, a number of magazines and publications ran special issues around it, a number of products were launched in this space for the first time.

You mentioned earlier, when BEA announced their product in 2007 in the real-time event-driven applications, that was I think, a signal to the industry, both to the other vendors as well as to the market, that this was something that needed to be taken seriously, because for the first time you had one of the top global firms entering this space. So, 2007 really was a very transitional year. It was the year where it really came to the forefront, a lot of people were talking about it, a number of people started implementing it.

But, by the end of 2007 you still had, I would say, the vast majority of use cases where it might be used, still untapped. So, you still had the existing systems that had been built over the years or the existing ways of doing things being practiced broadly.

And I think, if you look forward to 2008, certainly when I've talked to our customers and see the projects we have on-going at Lab49, it's clear that the adoption is really ramping up. And I think 2008 is going to be the year of broad-based CEP adoption. I think, certainly from a standpoint of the customers that we've seen, I would be pretty confident saying that we've done our last custom event-driven system that doesn't use a CEP engine.

Vicki: And let me ask, are you talking about this just in algo trading, because as you mentioned Ross you've seen it be applied in the algo trading environment. Or, are you saying 2008 is when it will be applied in other areas?

Ross: We actually have already seen areas where some of the more innovative companies are starting to realize that this technology can be applied to different use cases. To Dan's point, I do see those other use cases starting to be figured out and adopted in 2008. But, specifically, I think, we should talk about some of those.

And algo trading absolutely is the classic use cases, and it's set up well. It's all about low latency, fast response times, placing orders, executing trades in electronic marketplaces. However, we see that as only part of the overall application architecture. There's other things you actually have to do to make that happen. You need to be able to deal with vast volumes of market data, ever increasing numbers of market data at low latencies that you have to scrub, cleanse, normalize that market data, to feed those algo systems. You have to be able to do pre-trade analytics. You have to be able to do real-time risk management. Both market and credit risk management is becoming more cross-asset and it becomes a challenge across different books, desks, lines of business. Right?

Credit risk, as we all know, is a heavy issue for people; right? And with the vast volumes of activity going on, the sensitivity to that's only going to increase. We've even seen horizontal use cases where you can use these event-driven technologies to actually monitor your automated trading environment. Look at the latencies and bottlenecks from the second or the microsecond that you get a market data tick all the way through to actually settling the trade. You can see where your performance bottlenecks actually are.

Daniel: That's actually an interesting case because historically a lot of the CEP technology was created in a kind of signal processing or distributed monitoring context within academia. So, a number of projects were started at several universities around this idea of putting out a number of sensors throughout a system, let's say it's a factory or a fleet of trucks, which would each generate data at a fairly high rate. To build a system to collect all that data and process it in real-time to produce a real-time kind of picture of what all these sensors were out there actually doing, and it's interesting to see it come full circle is that sort of triggered some use cases within finance and now, people are actually using this to monitor both infrastructure and applications within finance.

Ross: Why don't you talk a little bit more, Daniel, about the market data side of things, because we've definitely seen certainly in the buy side, a lot of people figuring out that that's actually one of the critical points in the overall trade flow.

Daniel: Absolutely. Market data processing in general is becoming a major use case for CEP engines. We talked earlier about the kind of trends that have been driving the market. One of those trends is a massive increase in the number of market data sources as well as in the volume of market data from each of those sources. So, electronic crossing networks, ATSS, the classic exchanges, are all producing higher and higher volume. At the same time, the number of those venues has been proliferating over the years.

Ross: I think, another interesting point to note is that if you actually look at what some of these hedge funds and other proprietary trading desks are doing in the algo space is, they're placing these orders, these trades and as they see the market move while these orders are being filled, they actually are starting to modify or cancel and replace these trades and orders. And that just increases the number of trades in the marketplace by orders of magnitude. So, all of that activity, the really, really dynamic, fast, responsive trading activities driving these volumes up, in addition to just the sources of liquidity.

Daniel: Yeah, absolutely. So, the number of sources of data are proliferating, and the volume of data from each of those sources has been skyrocketing. So, what you're dealt with is this massive flow of market data that you have to somehow cope with.

Now, what are the problems that can create for you? Well, one is just simply the ability to architect systems to actually drink from that fire hose, right? You've got thousands or tens of thousands of messages per second that can be handled, and to write a system from scratch that's capable of doing that fast enough without memory leaks, with a high availability and disaster recovery is difficult. Doing it on a platform that provides those kinds of infrastructure for you solves some of those problems.

But secondly, you have issues of error detection and correction. So, what if a number comes from an exchange that is either wildly incorrect, like let's say the price has been at 99 for several ticks and then you get a price at nine or 999 for one or two ticks? And so, building kind of validation systems in a CEP engine to look at the stream and take a sort of sampling or do some pattern recognition around that stream of data, is a very, very classic and becoming a very popular use case of market data.

And then a third, that I'll just point out, is, often for the same kind of data there may be several sources. So, for equities prices we all know there are a number of different venues on which equities trade. Sometimes, your connection to one or more of those venues goes down. So, what you may have and what firms are doing, is they'll set up some rules around where their preferred source is for a given ticker symbol; but, if the source drops out or if we have a connectivity issue, let's pull in the number from the backup.

So, it's very, very easy to write some simple logic in your CEP engine to connect two sources and run them through some logic that says, "If this source hasn't ticked in so

long, let's consider it dead and pull in from our backup source." So, people are using it for some of the kind of backup and recovery solutions as well.

Ross: One of the things we actually didn't really talk about in terms of what the core needs are for the CEP technologies and what their drivers are, is the complex part of the name "Complex Event Processing." It's not just about processing these market data events very, very quickly and responding to them.

These technologies actually offer you the ability to look at very complex relationships across different symbols, different instruments, across different venues. So, the classic case would be doing basket trading or correlation analysis or trend analysis and looking for outliers or deviations from what you perceive to be a trend or a valid state of a marketplace. And that's really where some of the power of CEP technologies come in, in their ability for you to create those types of pattern recognition.

Daniel: Yeah, I think that's a great point. Two things have been happening, one is, as you mentioned, firms are becoming much more sophisticated in the way that they analyze the market and look for trading opportunities. And the other is that the actual instruments that people are trading are becoming themselves much more sophisticated.

So, one of the key areas of innovation within the financial markets over the last several years has been the creation of ever more complex derivatives. We've all heard of CDOs and some of the effects that those have had on the credit markets.

Ross: Just a couple of effects.

[laughter]

Daniel: Just a couple of small effects. Some folks maybe perhaps wish they hadn't been invented, but, they do serve a useful purpose in terms of structuring some of the risk and exposures that their customers can have. And that trend is actually only increasing.

There's a sort of commoditization cycle where a firm will come up with a new kind of derivatives contract, that contract will gain in popularity, eventually it will become standardized. At which point it can move from being traded over the counter or sort of on a custom basis to being traded on an exchange. We've seen that happen time and again, it's only going to increase in popularity.

So, those kinds of complex derivatives, structured products, mortgage backed securities, CDOs, synthetic derivatives of various kinds, hybrids, you name it. Those have all created challenges for firms in how do they actually model those instruments. They can get incredibly complicated, they can require lots of interesting kinds of simulation just to determine the prices for them.

And what you're seeing is that as the volumes increase in terms of the trading volumes of these instruments, firms need to rerun those pricing analytics, those models, ever more

frequently. So, it's a very, very natural fit for the kind of CEP model.

You may have an instrument like a mortgage backed security, the price of which and several of the analytics are sensitive to some underlying interest rate. So, what a firm will do is they'll tie into that underlying interest rate feed, say the 30 year treasury or the LIBOR rate and as those rates tick, on every tick or several times per second, will generate - it's event driven in a sense - will generate a series of downstream repricing jobs that the CEP engine will kick off. And that logic can get quite complex. It can do lots of Monte Carlo simulations, it can call out to external analytic libraries, and you want that to be this kind of cascading effect, where when this rate changes, it triggers some downstream calculations.

Ross: I think, a really good example of that, and we've actually talked to a number of customers about a specific scenario. For example, pricing off the run treasuries. The liquidity in the on the run treasury markets is pretty high, and everything's pretty much been driven off those market ticks across all those instruments. And it's a challenge pricing all the off the run treasuries in a very, very fast way, right? And I think, CEP technologies could definitely be an interesting solution in solving some of those problems, right?

Daniel: Yeah, again, it's one of the kind of classic if you look at the complex side of Complex Event Processing, where typically firms will price an off the run treasury in terms of some on the run treasury price and then some function of that - whether it's a simple spread or a more complicated curve based or multiple curve based model of what the difference is between that off the run and the corresponding on the run.

So, what you'll have is again this cascading effect. You'll watch this high frequency on the run trading market, and you'll then cause a cascading downstream ripple of calculations and repricing of your off the run portfolio.

Ross: Another great example would be, again in fixed income, where you have convertible bonds. So, where the bond prices and the optionality that's embedded are a function of the equity associated with the company's bond, right? So, now you've linked up the fixed income and equity market, right? So, being able to price that stuff is tough, apart from fact that embedded optionality itself is actually a complex process to initiate and compute.

Daniel: Absolutely.

Vicki: I mean, these are all very complex things, all related to pricing to quantitative factors.

Daniel: Yeah.

Vicki: What about credit risk? Ross, you mentioned earlier that this is ripe for application, too.

Ross: Well, the credit risk I was referring to has all been... Credit risk, fundamentally, is your exposure to counter-parties and whether they're going to default and whatnot. And there's all aspects of credit risk, and there's all types of tools and derivatives, instruments, to help hedge against that credit risk. So, CDSs and CDOs and all that type of stuff.

But, one of the points I was making earlier was actually more to do with the control systems that are actually required to know what your exposure is to those counter-parties. So, if you can imagine a range of desks inside a global market's business - fixed income, equities, FX, you name it. They're all trading, quite often, with the same counter-party, and they're all running up exposures against the same counter-party. So, at some point, there's a net exposure to that counter-party.

Now, in the real-time markets, how do you aggregate all that up, in real-time, and know what your actual exposure is to that counter-party such that you can actually limit or control the trading activity itself?

And this is something that we've talked about a lot. Trading systems are becoming very, very advanced, but, we do need to look at how the control side of systems can balance off and actually proactively feed back into the trading systems. Very complicated problem. Big enterprise-wide problem. It's very, very non-trivial, right? [laughs]

Vicki: You mentioned earlier about hedge funds and some of the things that they're doing applying CEP. Do you see a difference between the buy side and the sell side in their adoption of CEP?

Ross: Well, I think, the buy side has actually been more innovative and has been looking to adopt these technologies, because they actually see very quickly how they can take advantage of the technologies and pick off opportunities in the marketplace. So, they've certainly been a lot more nimble.

But, that said, there's been absolutely pockets of innovation in a lot of the top sell side participants. Prop desks, equity desks have certainly been driving it. But, you're starting to see a lot of pull-through now across other desks. And I think, you are definitely seeing a lot of the enterprise-wide architects and groups starting to look more seriously at this as being a platform for the enterprise rather than just a number of point solutions.

Daniel: Yeah. Simply due to the way that hedge funds are defined, they are smaller than the big global investment banks are, and so that tends to lead to them being a little bit more nimble and a little bit more agile in terms of their ability to evaluate and adopt new technologies. So, I think, you have seen that.

At the same time, the big players have big systems and big challenges and lots to do. So, they're looking at some of the bigger, more complex installations of this kind of technology. So, I think you are seeing kind of two parallel tracks there.

Ross: Right.

Vicki: And Daniel, you mentioned the maturing of CEP, from 2006 to 2007 to where we are today, 2008. So, overall, who do you see coming to the party now, as far as the various uses? Who do you think will come in 2008... and maybe, even going forward, maybe coming in 2009?

Daniel: I don't think it's an understatement to say everyone, or nearly everyone, in the sense that it's the same kind of picture, that you would say... again, getting back to the application server, "Who's going to be the customer of these application servers?" And initially, it was e-commerce, right? It was building a shopping cart. But, eventually, it came to the point where you're going to build and deploy an enterprise Java application without an app server? It's not done, right?

And I think, if you look at the broad sweep of applications that are built within financial services, it's going to coalesce around this model, I think, for a number of use cases, where it's just everything is moving to event-driven and real-time, and it's just going to start to become increasingly common.

Ross: And if you're not in the market, you're going to lose out, right? If your competitors will be able to do this type of complex processing in a much more performant way, how do you keep up with the Joneses?

Vicki: So, would you say that happening inside the major financial institutions right now, both buy side and sell side, is this very conversation that everyone needs to be at the table and be participating?

Daniel: Absolutely. Absolutely. I would say, right now, the vast majority, if not every customer that we have - within both the hedge funds, within the institutional asset managers, within the market providers and liquidity sources and networks, as well as in the sell-side financial institutions - are all considering, in one way or another, CEP applications.

Vicki: So, it's not just the trading desk over here in this corner?

Ross: No...

Daniel: Not anymore. Not anymore.

Vicki: All the reason for us having this conversation today. Excellent, thank you.

Ross: Great. Well, I think, this has been an interesting conversation, certainly in terms of introducing CEP and where we come from and, potentially, where we're going. I think, what we are going to do in some of our other podcasts is explore a little bit more some other angles. We're going to be talking to Brad Bailey in our next podcast about some of the research and industry analysis that he's done. And hopefully we'll be talking to some other people from the client-side, get their perspective on actual adoption in their companies.

Daniel: Great.

Vicki: All right. Thank you so much, Ross and Daniel.

Daniel: Thanks a lot. This has been fun.

Announcer: That was the second half of an interview with Daniel Chait, co-founder and managing director of Lab49, and Ross Hamilton, the company's director of client engagements. Both Ross and Daniel will be appearing as co-hosts in future episodes of this series. Our next episode will feature Brad Bailey, research analyst with Aite Group.

For more information on Lab49 and BEA Systems, for details of how to subscribe to this podcast, or to leave comments and feedback, go along to our website at www.complexeventspodcast.com. There's also links to Lab49, where you'll find the feature, 'Inside the Lab,' containing some excellent demos, research and analysis, and a link to our sponsor, BEA Systems, where you can learn about their solutions for Complex Event Processing.

The Complex Events Podcast was brought to you today by Voices in Business, helping our sponsors achieve thought leadership in their business sector. For more details of what we offer, go along to www.voicesinbusiness.com. Thanks for listening. Good bye.

[music]