

4/16/2009

03:45 PM



Penny Crosman

News

Connect Directly

[0 COMMENTS](#)[COMMENT](#)[NOW](#)[Tweet](#)[Share](#)[G+1](#)[0](#)

## Citi Taps CEP for Analyzing Equity Data

**Global financial giant uses Coral8 software to perform complex pattern matching on equity data in real time.**

Tech-savvy Citi executives have become avid users of [complex event processing](#) software.

According to Marc Adler, director and chief architect of equities at the firm, Citi finds CEP helpful for analyzing equity data, including order execution information, market data, and real-time greeks, volatility changes and news.

Although Citi has long had technologies that could process real-time data, before 2007 it lacked a way to easily correlate the data and perform complex pattern matching in real time, Adler relates. "That's what CEP offers," he says. "Before, you used to have to write these kinds of engines yourself. But the CEP vendors have made this a commodity."

Typically CEP vendors provide a series of input and output adapters that connect their CEP engines to external data sources. But, Adler notes, Citi doesn't use these because it has developed its own. "We use just the core part of the engine without a lot of decorations, because who knows when we'll have to change our CEP engine," he explains.

Nonetheless the CEP vendors do provide useful tools. Adler points to the visual development environments as well as debugging and monitoring support that CEP products provide as particularly helpful.

Citi began exploring CEP technology at the behest of a former global head of equities. According to Adler, this executive perceived a need for technology to capture and analyze the multiple events that stream through Citi's equities area in order to make business decisions and improve situational awareness. "He didn't say, 'Go and buy a CEP engine,' " Adler recalls. "He said, 'We need some way to process and analyze the multiple events that are constantly streaming through our equities flow.' "

The global financial services giant declines to elaborate on the specific business purpose for which it's using CEP software (beyond analyzing equities data), but Adler acknowledges that "CEP is generally useful for back-testing and real-time trading."

### The CEP Selection Process

Citi first began evaluating CEP technology in fall 2007. "We knew nothing about CEP, so we read all the literature and blogs on the subject," Adler says. "We went to an event processing summit in September, [and] we met with a number of the vendors for educational purposes and to see what was out there in the marketplace." In addition, Adler notes, he spoke with colleagues at other investment banks to see what they were using.

Next Citi invited Coral8, Aleri and StreamBase to conduct proofs of concept. As for the fourth major CEP contender -- Progress Apama -- "We wanted to evaluate their software, but we found that they had a heavyweight sales process. For instance, they wouldn't let us download the software," Adler explains. "To be fair," he adds, "John Bates, the CEO of Apama, didn't know about this, and he says that if he did, he would have given us the software."

Citi did not test the three available products for speed. "We didn't care that much at the time about latency or which engine was the fastest," Adler says. "We had a particular use case we wanted to implement, so we were looking at a lot of intangibles, such as ease of development, how hard it was to do what we wanted in the language, the visual development environment, and support for things like monitoring and debugging." Citi also looked at the companies' customer support and developer friendliness -- anything that would impact the development experience, Adler adds.

It was the development environment, according to Adler, that ruled out Aleri. "Aleri is like a V8 engine that's surrounded by a clunky body," Adler says. "Our impression was that they had very smart Bell Labs scientists who had a great engine but didn't fully appreciate what the development experience should be."

StreamBase, according to Adler, backpedaled on its performance promises and had a higher price than Coral8, Citi's ultimate choice. "Coral8 didn't have the best developer tools or documentation, but overall the company felt good to deal with," Adler relates. "It was developer-friendly and had the lowest price point."

In March 2009 [Coral8 and Aleri merged](#), retaining the Aleri name. According to company announcements, product pricing will stay the same. Executives have vowed to both continue to support existing CEP engines as well as create new, best-of-both-worlds products. According to Adler, part of the integration plan is to abandon Aleri's version of SQL in favor of Coral8's [Continuous Computation Language](#) to handle queries.

One Aleri feature that Adler would like to be able to use in his environment is its Splash procedural add-on language. "What a lot of these SQL-based CEP engines lack is the ability to move out of the SQL paradigm into a more procedural approach," Adler says, noting that developers often use Splash to obtain performance improvements. He adds that with the recent merger of Aleri and Coral8, he is hopeful that Splash language capabilities will be incorporated into the new Coral8 product.

### The Results

Citi's experience with Coral8 to date has been "good and bad," Adler reports. "Streaming SQL is a new paradigm; a SQL developer might be immediately productive with Coral8, but writing queries in such a way that they perform well can be challenging." For instance, in the Coral8 language, one error can slow overall message processing to a crawl, allowing the message queues to back up, Adler says.

"Without the in-depth help the Coral8 architects gave us, we wouldn't have been able to figure out how to get it to perform as well as it does now," he continues. "You could get yourself into trouble with the language very easily and not know why. There's not much transparency in Coral8 as far as what modules and queries are performing badly. I hope the merger will result in improvements in monitoring tools and traceability of query performance."

Any CEP product requires the support of high-end servers, Adler relates. "Our system runs on an eight-core machine, and we've improved our code a lot" to take advantage of all the cores, he says, explaining that multiple cores allow developers to assign a different process to each core. "But I would not run any of these engines on anything less than a four-core machine."

It's also important to be able to scale up to multiple machines that communicate with each other so you can perform load balancing among multiple Coral8 instances on multiple servers, Adler adds. "We find that is amazingly valuable," he says. "The more hardware you throw at it, the better, because of increasing message rates."

## The Future

Eventually Citi may expand its use of CEP to other parts of the firm, according to Adler. "We'd like to have CEP everywhere," he says. "We think there's a new paradigm coming where everybody will start to publish data into us. We'll take in more kinds of internal data, across asset classes and regions, and analyze it and provide intelligent alerting to people up and down the value chain." The firm, Adler notes, already has developed its own .NET-based visualization tool to enable people to easily read views of data.

One thing that is becoming more important in the broader world of CEP, according to Adler, is real-time online analytical processing, or **OLAP** -- technology that lets users easily and selectively extract and view data from different points of view. "This seems to be the way Wall Street firms use CEP, although that was not the original intention of the CEP vendors," he says.

"Using OLAP, you could have real-time aggregations across many dimensions that you show visually, and different people could subscribe to different aggregations," Adler continues. "The tricky part is managing the different entitlements and views among regions or trading desks so what gets displayed out of one OLAP cube to one user can be totally different from what's presented to another user. Imagine you're analyzing something for the entire firm, but you don't want the sales trading desk to see what the prop desk is doing and vice versa; you need to be able to give each individual trader a different view of that world. That's something that's very interesting to do in real time." Adler says he would like to see vendors start to provide more support for different entitlements and views of OLAP data.

A wild card in the future of CEP, Adler adds, is open source offerings such as **Esper**. Citi didn't choose this solution because it lacked full-time, dedicated support at evaluation time and it doesn't come with development tools, he says. But Merrill Lynch uses Esper in equity derivatives.

"With a high-profile project like we have, if one of our developers is stuck on a problem, we don't want to have to wait for one guy to answer our question; we'd rather call up a dedicated support

line and get our questions answered in hours, if not minutes," Adler says. "But that takes nothing away from [Esper's] technology."

[COMMENT](#) | [EMAIL THIS](#) | [PRINT](#) | [RSS](#)

#### **MORE INSIGHTS**

##### **Webcasts**

[Can You Trust Your Data?](#)

[Finding a Cloud Solution That Meets Your Security Requirements](#)

##### **More Webcasts**

##### **White Papers**

[Open Source Security Risk: Strategies, tactics, & tools to help reduce Your risk](#)

[Print Security: Everything You Didn't Think You Needed to Know](#)

##### **More White Papers**

##### **Reports**

[\[Secure Application Development\] New Best Practices](#)

[Ransomware Report](#)

##### **More Reports**

Copyright © 2016 UBM Electronics, A UBM company, All rights reserved. [Privacy Policy](#) | [Terms of Service](#)