

A Short History of Complex Event Processing¹

Part 2: the rise of CEP

by
David Luckham²

This second article follows on from part 1 on the history of complex event processing³

CEP is the logical and obvious next step in the development of event processing that were described in our first article.⁴ The explosion of event traffic over the past twenty years has created a new set of demands. The IT layers (i.e., the company's networks, middleware, enterprise service bus and websites) are humming with this traffic and people want to listen in. This created a demand to extract information from the event traffic. As we shall see, it is difficult to think of anything else that could have taken place in event processing at a particular point in time, around 2000, but CEP. The final question, yet to be answered of course, is about the future and what will eventually happen in event processing.

CEP consists of principles for processing clouds of events to extract information, together with technologies to implement those principles. I have written a separate overview of the main concepts of CEP.⁵

1. MOTIVATION AND OPPORTUNITIES

The motivation for complex event processing results from the growth of event processing in general. First there was the explosion in networking in the 1970's. This was followed in the 1980's by an exponential increase in the use of messaging at the management level in the day to day operations of modern enterprises. But at this time the only analysis of this traffic was in terms of network loads and flows. And then came the Internet!

¹ © 2007 David C Luckham

² I wish to thank Roy Schulte for helpful and detailed comments on this article.

³ See <http://complexevents.com/?p=321>

⁴ See figure 1 in article 1.

⁵ A Brief Overview of the Concepts of CEP

1.1 Event clouds

By the late 1980's communication by higher level events had become the basis for running enterprises everywhere – in business, in government, and in the military. Any large enterprise had linked its applications across the networks from office to office, sometimes around the globe. It now operated on top of what was referred to as “the IT layer”. Business and management level events – say trading orders, or planning schedules or just plain email – were entering its IT layer from all corners of the globe, from external sources as well as from its own internal offices. And in all the different formats used by the applications. Enterprises were essentially operating in a veritable *cloud* of application level events.

So now there was another problem. Not only did one need to keep the enterprise IT layer running smoothly, but one was faced with trying to understand what was happening in it. The event cloud obviously contained a lot of information – often called business intelligence – that would be useful in managing the enterprise. But the cloud did not come with any explanation, events simply arrived! Sometimes at rates of many thousands of business events per second. And attempts to extract this intelligence amounted to storing the events in data warehouses and later on trying to do data mining.

What kinds of information might one try to extract from the cloud of events flowing thorough the IT layer of an enterprise? Here are a few examples:

1. Are our business processes running correctly, and on time? This question might apply, for example, to supply chains or on-line retail websites or trading transactions.
2. Is our information at risk, is anyone trying to steal from us? Are our financial traders violating their permissions?⁶
3. Are our accounting processes complying with government regulations?
4. Is there an opportunity currently developing between different financial markets for our trading programs to make a profit?
5. Are our call centers servicing our customer requests in good time? How are our customers reacting?

There are many other questions like these. They range from managing the enterprise and detecting market opportunities as they happen, to protecting the

⁶ France to Fault Societe Generale's Controls in Report. See also http://en.wikipedia.org/wiki/J%C3%A9r%C3%B4me_Kerviel Carrick Mollenkamp and David Gauthier-Villars. Wall Street Journal. (Eastern Edition). New York, N.Y.:Feb 4, 2008. p. A.3

enterprise and ensuring that its processes conform to policies. Indeed, all kinds of possible uses of the information in the event cloud have arisen. And the emphasis is on “real-time” – getting the answers as the information crosses the IT layer. These are the opportunities that have led to CEP.

1.2 CEP versus Custom Coding.

Vendors of CEP products entering new markets such as financial trading in those early days around 2000 - 2003, would often complain “our biggest competition is from custom coded solutions”. What did that mean? You might well think that it meant some machine coded hackery. But in fact it simply meant that the competition was coded by the customer’s own IT department and not purchased from an outside source! Often a custom coded solution was indeed a program that was focused on a single problem and had an overly restricted range of applicability. Even so, it might apply principles of CEP, although the IT department may never have heard of CEP. And usually a custom coded solution worked pretty well because it had been working for some time before the CEP vendors entered the market. Also, the first generation CEP solutions were far from perfect, and most of them applied only a little CEP, as we shall discuss.

However, event processing was “taking off”. And when the customers wanted to extend their custom coded solutions to other event processing issues, even in the same problem domain, it usually turned out to be a labor intensive and expensive task. That led potential customers to re-evaluate the “build it or buy it” equation. So CEP began to make inroads into these markets. This battle is still going on today!

1.3 The Dawn

With the advent of networks came a host of network management tools⁷ whose job it was to track and trace events and display the results graphically. Some of these tools, such as HP Openview, CA UniCenter, IBM Tivoli NetView and BMC Patrol, could be considered as the early precursors of BAM (Business Activity Monitoring). They were monitoring events in the network, and sometimes trying to reconstruct events at more abstract levels. So by the mid 1990’s we can see the tip of the CEP iceberg! Whether the implementers of these tools had any explicit definition of CEP principles in mind may be open to doubt. But two things are worth mentioning.

⁷ see Article 1, figure 1

First, these network management tools made graphical dashboards the main access medium for the user. The dashboard was developed to supply the user with tables, graphs, and pictures of the monitored network traffic. It was established as the default interface for delivering insight to the user, or at least kidding the user that insight was being achieved! I'll come back to dashboards and implied insight later.

Secondly, an astounding fact. The major vendors of these early network management tools were slow to capitalize on the opportunity to extend their technology to business level events. It is clear that such extensions could have been made far in advance of other entrants into the business event processing market. I have asked many commentators and students of event processing why this is, and I have yet to find a satisfactory answer. Perhaps they were too focused on the existing network market and didn't really believe there was a market for higher level business event processing. But there can be no doubt that HP, IBM, CA and BMC all missed an opportunity. And for this, the small start-up vendors in the CEP field today must be eternally grateful. Of course some of these large companies, IBM for example, are in the business event management market now, along with Oracle, and doubtless the other big players will follow. But today this market is a battleground for market share, and it is unclear if the big players will still dominate. Their foresight and timing could have been better!

2. THE FOUR STAGES OF COMPLEX EVENT PROCESSING

There are four stages in the evolution of CEP tools, starting around the year 2000. There are no sharp distinguishing time points when one stage ends and the next begins. In fact each stage overlaps in time with the next. Figure 1 shows the approximate time intervals of these stages.

2.1 Simple CEP (1999 – 2008)

Event processing applications in these early days explicitly implemented only a few of the simplest CEP principles⁸. So, I refer to this as the era of simple CEP although the *data* processing (as opposed to *event* processing) was often very

⁸ A Brief Overview of the Concepts of CEP

sophisticated. In other words the events being processed might be complex, but only the simplest event processing techniques were used. Even so, the intended users were business analysts or software developers, trained professionals.

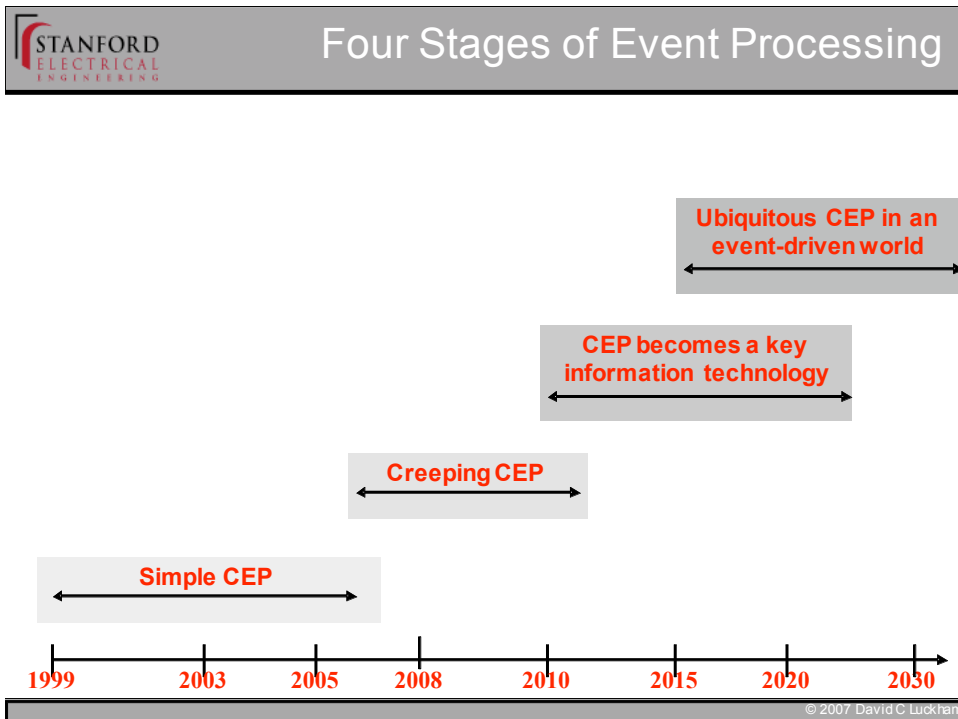


Figure 1: The four stages in the evolution of CEP applications.

For example many event monitors intended, say, for keeping tabs on business processes, allowed only one predefined type of events to be used to trigger alerts or ECA rules. Sometimes there was a limited facility for users (professionals) to specify new (not predefined) event types as triggers of ECA rules. In cases where the monitors allowed users to specify additional types of events, or event patterns and event triggered rules, they usually had to be coded in Java script or an extension of SQL. Also, the kinds of event patterns that could be used were simple Boolean combinations of events, or maybe some subset of regular expressions. Timing conditions could be added to events in a pattern. But nobody thought causal or independence relationships between events worth the effort.

Most of this era's tools came with a graphical output display called a dashboard for displaying statistics computed from events in real time. These statistics were called Key Performance Indicators (KPIs).

In 2001 Gartner coined the terminology "**BAM**" for Business Activity Monitoring to classify the market that these event monitors were targeting. This simple label had the effect of focusing attention on this new area and helping people to understand in general terms what it was all about, and what its goals were. True, acronyms have proliferated beyond all understanding in the business applications area, but this particular one was worth its weight in gold!

As the development of CEP tools progressed, graphical input tools were introduced to help users define event patterns. Graphics gives a user a way to easily compose an event pattern. Even so, at some level in the pattern definition a user often has to resort to writing program text of some kind.

The first established market for CEP was financial trading, specifically algorithmic trading. Here the event cloud contained market trading data from various sources, usually including several stock market feeds. Financial trading had a tradition of custom coded tools so CEP vendors found that event processing was already well understood. In fact, the biggest competitor to CEP products was custom coded solutions. Elsewhere vendors found that their most significant problem in penetrating new markets was educating the marketplace about event processing and its potential benefits.

2.1.1 Much ado about event streams

In financial trading the important components of the event cloud were usually stock market feeds reporting transactions and stock prices. It is easy to think of these feeds as streams of events – linear sequences of events ordered by time. One of the requirements for applying a CEP tool to such data was that the tool could handle the rate at which events arrived in these feeds, often thousands of events per second. Thus a name for a subset of CEP was coined – event streams processing (ESP). The intention was to infer that a tool that did ESP could handle high speed event feeds. Also some marketeers feared that the "complex" in CEP would arouse fears often associated by their customers with impenetrable and faulty software. So some vendors used the term ESP as a marketing ploy for a year or two, but they soon began to realize that they needed to explore other markets where the event clouds were, perhaps, more complex.

As time went on, ESP was subsumed back under CEP, and everyone was once again doing CEP, which in fact was the truth of the matter all along! Today “CEP” has become an accepted term in business journalism.

I have placed the early (2000 – 2004) financial trading applications under “simple CEP”. This may cause some wonder since the actual trading algorithms were often sophisticated and proprietary secrets. Well, although the trading algorithms were far from simple, and the events could be complex, the event processing involved was very simple. Consider a typical component of such event processing, the Volume Weighted Average Price (VWAP) computation (see Figure 2).

The figure shows a typical event feed as a sequence of stock trade events where each event contains data such as the stock symbol, number of shares of that stock traded, price and time. As events arrive, they are processed over a time window which may be a few minutes, or perhaps an hour. The VWAP computation shows two feeds being processed simultaneously. First a filtering operation is applied to split the two feeds into a set of streams, one for each stock symbol. This is certainly an event processing operation. These streams of single symbol events are sent to two functions. One computes total dollar amounts over a time window, while the other just adds up the total volume of the trades in that window. The results (shown by different colored arrows) are a stream of dollar amounts and a stream of share volumes for each stock symbol. These two streams are fed into a third computation that divides the dollar amounts by the share volumes. The result is a stream of complex events giving the VWAP for each symbol. Only simple techniques of event processing are involved: filtering stock feeds by symbol, and moving events between computations. The complexity of the algorithm is captured in the functions that are applied to the events.

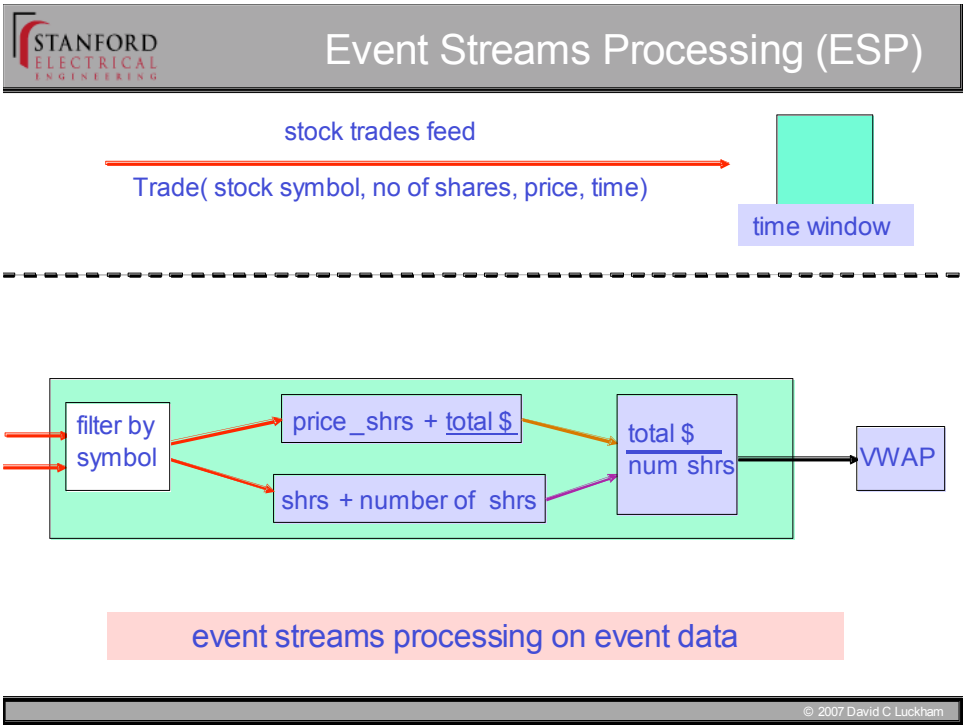


Figure 2: event streams processing to compute VWAP

Actual financial trading applications will involve networks of such like computations in which VWAP is just one node, but the event processing is still very simple.

2.1.2 Hard times

The years between 2002 and 2007 were difficult times for most small CEP vendors. It was a time when business analysts might talk about the need to “gain traction” in markets. Some small vendors were assimilated by larger companies, and others failed. But most continued to pioneer new markets while running at a loss. Belief, the stuff of innovation and enterprise, kept them going. Meanwhile the large players were sitting on the sidelines developing event processing additions to their product suites and trying to decide if there was a business event processing market.

The third of our two articles on a short history of CEP will deal with the “gold rush era” in development of CEP, the new potential markets, use cases, and future technical progress.