

A Short History of Complex Event Processing¹

Part 3: the formative years

by
David Luckham

Third of two articles on the recent history of complex event processing: the origins of CEP vendors and the struggles for survival of the early vendors.

In “A Short History, Part 2²” we left off our story of CEP at a period of “hard times” for the vendors of CEP products during the years 2000 to 2007. This was the first of the four stages of Complex Event Processing. It was the take-off stage for CEP. Figure 1 show the stages. The vendors were struggling to convince the




Figure 1: The four stages of CEP development

¹ © 2008 David C Luckham

² “A Short History of Complex Event Processing Part 2: the rise of CEP” by David Luckham, <http://complexevents.com/?p=397>

one market that understood event processing, the stock trading and related financial services market, that their products could make a profitable difference. And, at the same time these vendors were also wandering in the wilderness searching for potential new markets and trying to educate them about the benefits of event processing.

1. WHERE DID THE CEP VENDORS COME FROM?

First of all, there were the startups. CEP ventures started being formed about 1999. Most of the startups were formed by members of university research projects working in one of the four major areas of event processing: event-driven simulation, networking, active databases and event processing in middleware³. One of the early startups grew out of work on event-driven simulation while another grew from work on generalizations of middleware event processing. But the largest number came out of an active database background. This is the reason why many of the commercial languages for event processing are extensions of SQL. In some cases the CEP technology was actually developed and experimentally tested on the research project before the company was formed. But in most cases the companies were formed with little more than some ideas and a business plan.

The second source of CEP technology was the large IT suppliers. And there were three different pathways within these companies to CEP products. The first was research. There were research groups within these companies engaged on similar projects to those in universities. Some of the groups ended up producing prototype CEP products. In these cases, the decisions as to whether to let these prototypes go forward to productization were haphazard, probably because they were made by a management that didn't understand the technology or its potential. The second pathway was an explicit business decision to develop CEP products as add-ons to existing products in the areas of SOA and ESBs. These add-ons tended to emerge towards the end of the formative period, say 2007 and later. By this time management was playing catch-up with the startups! New add-ons are continuing to appear with increasing frequency as enhancements to product suites. The third pathway was acquisition. Some of the large IT suppliers simply bought CEP startups. In some cases one can find all three pathways being followed within the same company. Today, some of the large IT suppliers have three or four competing CEP offerings.

³ see "A Short History of Complex Event Processing, Part 1: Beginnings" by David Luckham, <http://complexevents.com/?p=321>

It is not our purpose here to name vendors or give categorizations since the landscape is still very fluid, to say the least. A comprehensive categorization of CEP vendors can be found in various Gartner presentations and publications, particularly those by Roy Schulte. Gartner does a very good job of keeping up to date. The latest Gartner classification⁴ lists about forty vendors, small and large, in different categories. The categories include

- “CEP engines”,
- “Integrated Development Environments (IDEs)”,
- “Embedded CEP platforms”,
- “SOA Infrastructure”, and
- “ultra-low latency messaging”.

I suspect not only the actors, but the categories themselves will change in the next year or two.

2. LIFE IN THE EARLY CEP DEVELOPERS, 2000 - 2006.

These were not easy times for small CEP startups. Some of the startups were acquired by the large companies. Others simply went out of business. Actually the failures were surprisingly few in number, although many of the small vendors were probably perilously close to folding from time to time.

At this point in time the BAM packaging for delivering CEP to customers by means of dashboards had been invented, or more accurately, recognized and named.⁵ Secondly, new markets were being pioneered by small vendors. And thirdly, the large IT suppliers were sitting on the sidelines watching and wondering – as we mentioned in part 2. Most of the pioneers were struggling, staying alive on venture investment and running at a loss most of the time. The sales of proof of concept applications kept them afloat and gave them hope. It also validated the belief of the investors.

These struggles on the part of the small vendors were hidden from an outside observer. A brave and optimistic face was always presented to a casual visitor. But there were indications of the real situation, for example the frequent turn over in senior personnel such as VPs of marketing. On the positive side, the sales of proof of concept trials in new potential markets were increasing in frequency. This activity was all part of the “educate the markets about CEP” activity that the

⁴ September 2008.

⁵ BAM, Business Activity Monitoring, so named by Gartner, 2001.

vendors were being forced to invest in. It wasn't easy. But I like to think that the publication of my book, "The Power of Events", in 2002 helped. It provided a roadmap to the technology development and a vision of how the technology could be applied. I know of some marketing departments that were buying 10 or 20 copies at a time and distributing them both internally and externally to potential customers. Today, at the beginning of 2009, we are about one third of the way along that roadmap!

The book had originally been intended to be the center piece of a marketing campaign by a start-up called ePatterns. The company was founded in November 1999 to commercialize the CEP technology that had been developed on the Stanford Rapide project. Unfortunately the business plan was overly ambitious and the CEO spent all the funding in the first year hiring a large organization before it had built a product. When he asked for the promised second round of venture financing he found that it was not forthcoming. The investors reneged during the dotcom implosion of 2001. I continued to write the book after the company folded. Today, the book is all that remains! If ePatterns had persisted through these times by a more judicious approach to company development and the expenditure of capital, it would be doing quite well today!

3. THE DOTCOM IMPLOSION

In the period leading up to mid 2000, money from many sources, e.g., pension funds, university endowments and private investors, was pouring into hi tech through the venture investment firms. All kinds of schemes and silly ideas were being marketed up and down Sandhill Avenue – or "Vulture Row" as it was often called. Hi tech startups were popping up everywhere. I remember chit chat between the cubicles, like "I can't believe the valuation of this or that company – it's worth a tenth of that!" Stories of overly ambitious valuations and large investments coupled with unrealistic goals were commonplace. But in fact, that's the way bubbles go, first up and up, and then, the flow of money dries up, and everything comes down with a bang!

The same companies that had outrageous valuations in 2000 suddenly disappeared in 2001 or 2002. They "folded their tents in the night, and like the Arabs, as silently stole away." The office entrance and window signs disappeared. And there were frequent large auctions of accumulations of repossessed office equipment. In fact, if you were an employee of a terminating company, you could hand in your company laptop and then go to an auction a

few weeks later and buy it back at a fifth of its value – it would probably still have your name on it! And people that one had met just a few weeks before while visiting these companies were out looking for jobs again.

While ePatterns was not the only CEP startup that folded or was acquired in the aftermath of the dotcom implosion, there were surprisingly few failures among CEP developers. However, the small startups that appear today to have survived through the hard times and are now on the road to CEP success were founded later in this period, around 2003. So while they had to raise funding in an unfavorable climate, they didn't have to survive the dotcom implosion.

Finally, it must be said that all manner of high tech companies failed at this time. Such stories as this were certainly not limited to CEP. They were very common across all technologies.

4. TECHNOLOGY CHAOS AND THE EPTS

The variety of backgrounds of the CEP vendors seems to be the only explanation for the correspondingly great diversity of technology offerings. It is true to say that no two CEP products are similar. The formalisms for defining event processing problems are all completely different. Even if two vendors use SQL variants as input, those variants have different ways of expressing the same problem, different features and execution semantics, and different answers will result. The user interfaces bear no resemblance to each other, other than they are all graphical. The engines all have different performance characteristics. And if you look for some commonality in the architecture of these products, you'll have to look very hard indeed! True, many of them are implemented in Java, but that's about it.

This is not a good situation. Typically, it is an indication of a market in an early stage of development. But it has led some potential customers of CEP tools to build their own in-house rather than rely on current offerings. And those customers have tended to be large enterprises. So the sooner things change and some commonality is introduced into the event processing arena, the better for all involved.

Towards this end the Event Processing Technical Society (EPTS)⁶ has been formed. EPTS is explicitly *not* a "standards organization". But it is devoted to developing foundations for CEP activities. For example, it is developing a

⁶ <http://www.ep-ts.com/>

glossary of common event processing terminology, perhaps a very first step that needs taking. And it has work in progress on *Interoperability Analysis* with the goal of defining a set of “mechanisms” that will allow interoperation between event processing systems from different vendors. Also there is work going on to write a survey of the existing languages. This is expected to involve a classification of language capabilities and may result in a position paper with possible alternatives for language standardization. Other EPTS work includes uses cases, benchmarks and reference architectures. This is all stuff that needs to be done now in order to take event processing beyond the stage of Creeping CEP to the next stage, a Key Information Technology.

The next article is on the Creeping CEP stage. Times get better for the CEP vendors in 2007 and research companies begin to track and estimate the CEP market.