

“Complex” or “Simple” Event Processing by David Luckham

Some people, I’m told, get scared when they hear the word “*complex*”, as in “Complex Event Processing” (CEP). They want to hear “*simple*” event processing — or so some IT marketing people tell me. The reason, I’m told, is a feeling that software systems have gotten too complicated and non-programmers can’t use them anymore. Well, let’s find out what makes an event, “*complex*”.

Start with a basic question, *is life simple?* Most people, if they think about that question, will truthfully answer, “*no*”. Events happen in life that are neither simple in *how* they happen, nor in the *effects* they have. We all know that. We face complex events every day of our lives. Take the 12-26-2004 Indonesian Tsunami as an example. NOAA has very sophisticated simulations to explain how that event happened. And its effects are continuing as we speak.

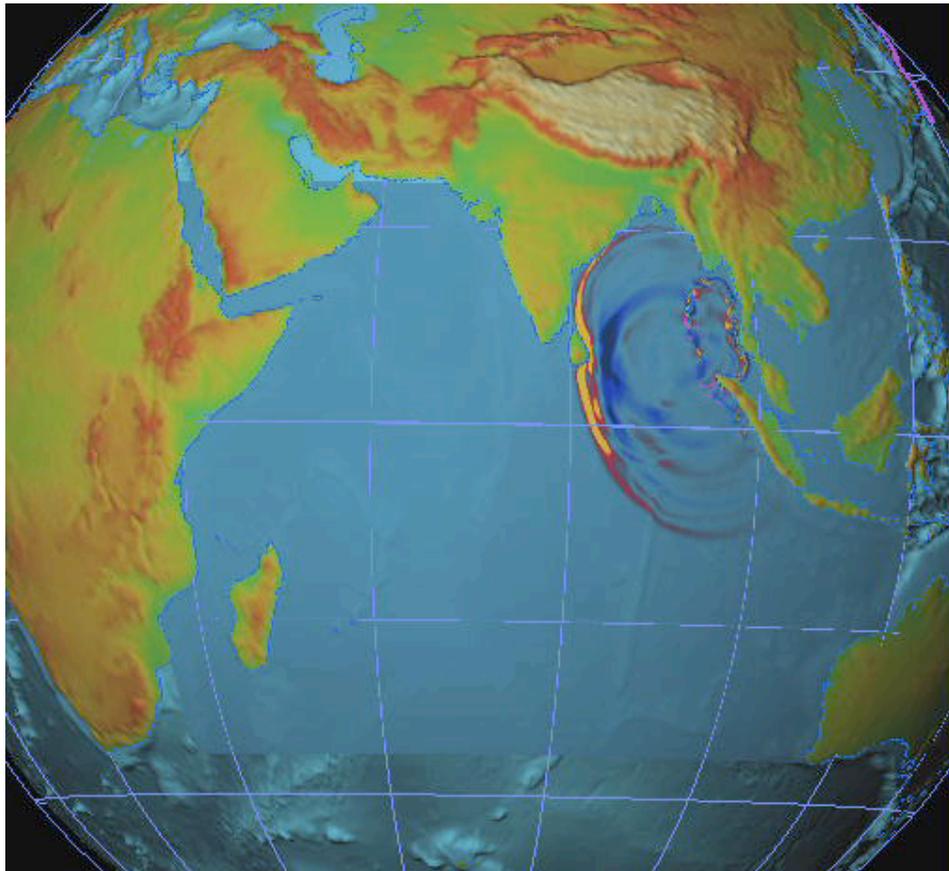


Figure 1: NOAA simulation of the 12-26-04 Indonesian Tsunami

If you want a technology that can deal with *life*, or in your case, with all of the events in the business and IT infrastructures upon which your enterprises depend, that technology will have to deal with complex events. You can't talk your way around that fact by using some other word, like "simple" or "composite"!

CEP is a foundational technology for detecting and managing the events that happen in event driven enterprises. One of the first objectives of CEP is to enable an understanding of the events in the enterprise. Only when we understand what is happening, or going to happen, can we plan and take action.

Here are some of the basic concepts aimed at achieving understanding:

1. Events in any enterprise can be organized into hierarchies, called *event hierarchies*.

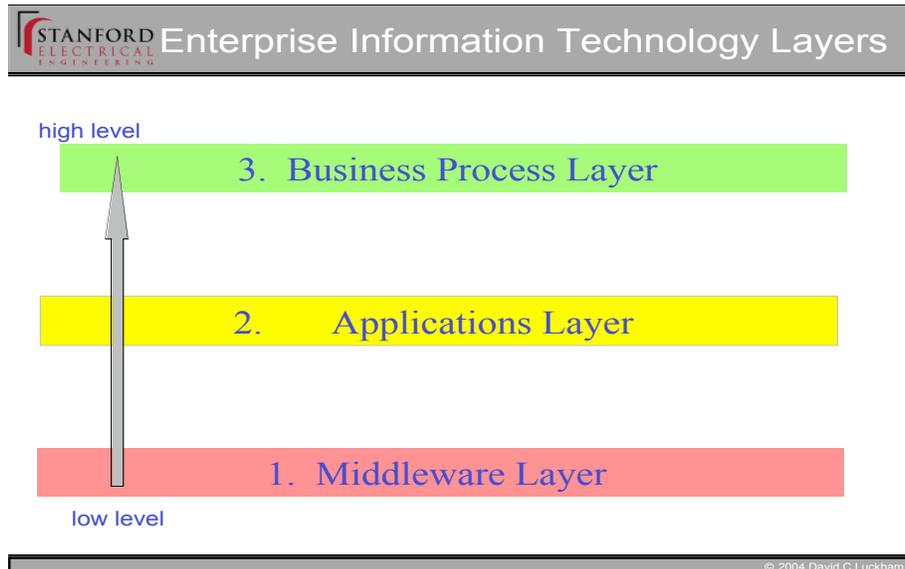


Figure 2: typical layers in an enterprise IT infrastructure

Figure 1 shows an example in which the event layers are (1) middleware events such as formatted messages with various subject headers, (2) application events, i.e., events resulting from use of applications such as database insertions, email, etc., and (3) business process events such steps in a sales transaction. Of course, there are other layers too, but let's focus on just these three. The idea behind classifying events into levels in a hierarchy is to achieve understanding of what's happening in the enterprise. One can focus first on the high level events associated with the business operations. Those events are closest to management and decision-making, and their significance for the

enterprise is most easily understood. To explain how they happened one uses their dependencies on lower level events. That brings us to a second concept.

2. Events are related to one another. Common relationships are *cause*, *time*, and *aggregation*.

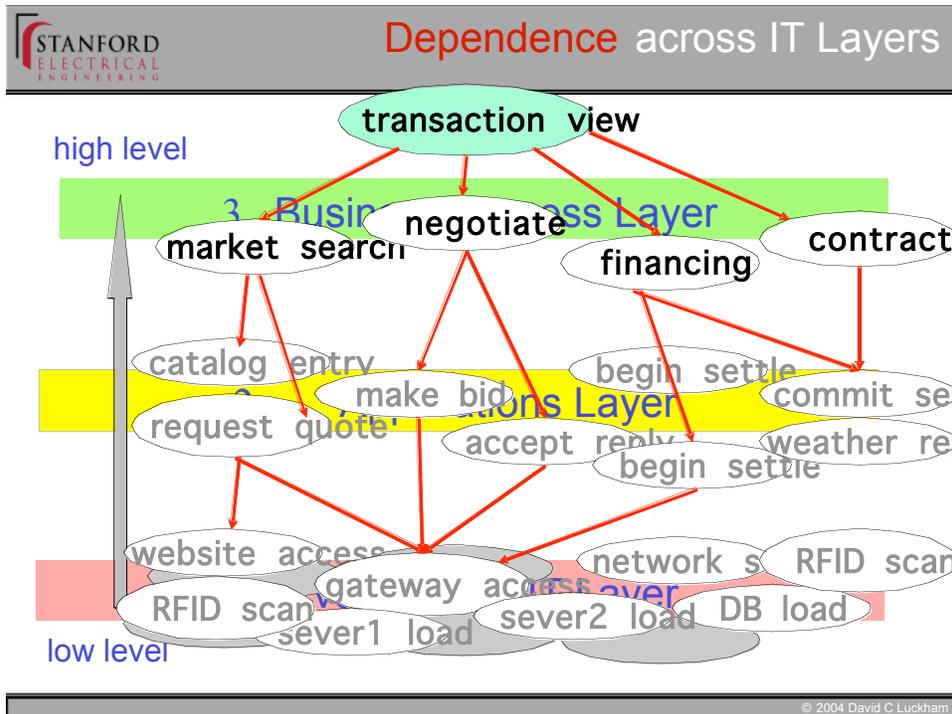


Figure 3: Events at different levels and dependencies between them.

Figure 3 shows some high level business process events like *market search* and *negotiate*. They could be created either by a business workflow engine or manually. They signify steps in various business processes, and they depend upon lower level application events in order to complete those steps. So, the process step events depend upon various events happening at the applications level. And so on. Figure 3 shows some of the dependencies as red arrows. The lower level events *have to happen* in order for the higher level events to happen. In fact, if the lower level events don't happen, say because there's a middleware error, the higher level events don't happen either. And the transaction hangs up.

Consequently, when those lower level events do happen, they **cause** the higher level event that depends upon them to happen also. This is shown in figure 4 by the blue arrows. A higher level event is called *complex* because it is caused by a pattern of lower level events.

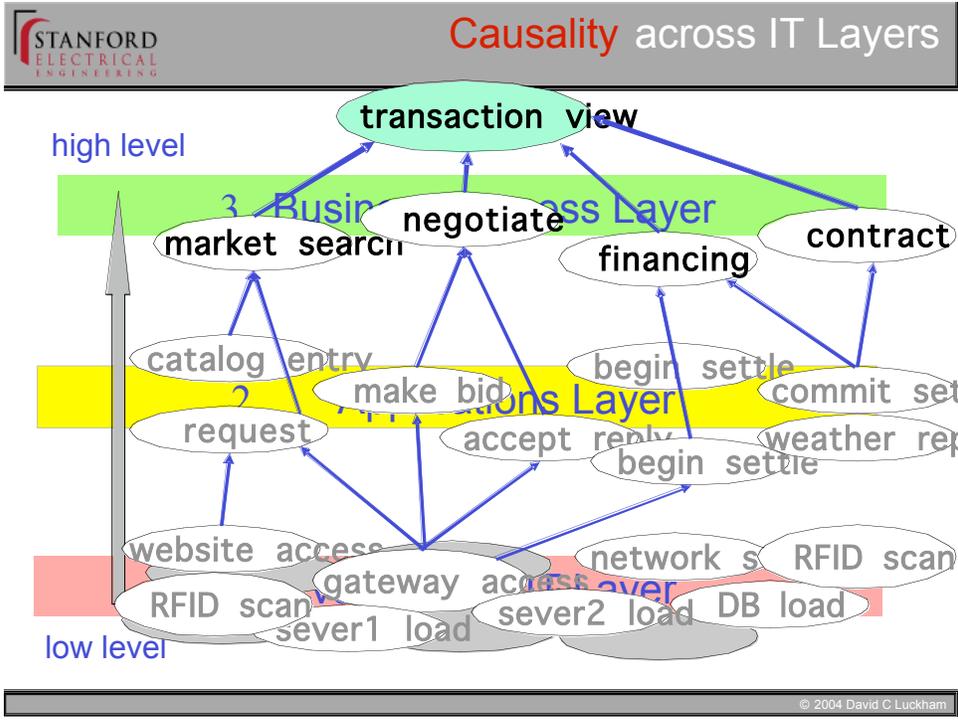


Figure 4: lower level events cause higher level events

The highest level event in figure 4 is a summary (or view) of the progress of a transaction. Usually that event won't happen at all unless there's a special tool that tracks process events and creates views. The tool uses CEP techniques to detect patterns of process events and create a higher level view event that contains a summary of the process steps thus far. A transaction view is an event that is caused by steps in one or more business processes.

A view of the progress of a business transaction is quite a complex event, depending upon lots of process steps. It is an *aggregation* of process events over a time period. Funnily enough, although it is complex in the sense that it is aggregated from many other events, it is easier to understand than the cloud of lower level process events. It abstracts essential data from those lower level events, and omits unnecessary details.

One of the benefits of organizing events into hierarchies is *understanding*. That brings us to another concept in CEP:

- 3. Different personnel need different views of the events in the enterprise, each view related to the role of the person.

For example, a CFO might want a view of the business transactions that are in progress. This knowledge may influence his financial planning. A process

architect or business manager on the other hand, might want a more detailed view of the steps in each transaction and also their dependencies on lower level events in order to analyze the impact of application or middleware glitches on the running processes.

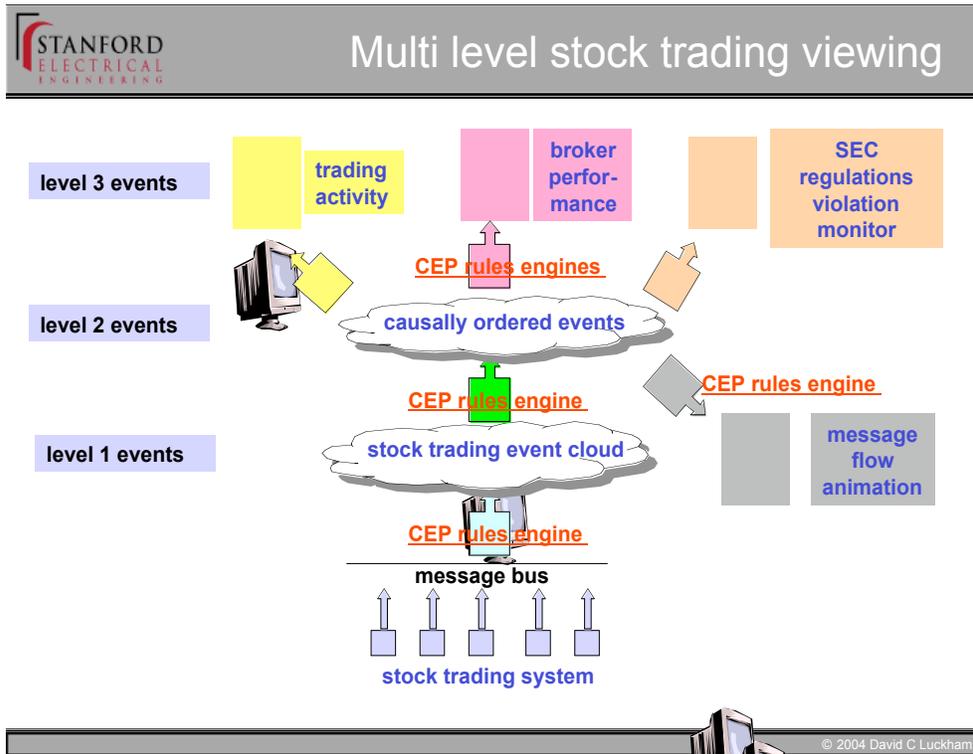


Figure 5: Separate real time views of events in a stock trading system

Figure 5 depicts role-oriented, real time viewing in a stock trading system. This is a highly event-driven global collaboration between multiple enterprises and individuals over various distributed IT media. Trading events, which may include anything from ticker tape messages to interest rate changes by the Fed, OPEC oil prices, etc., are processed off the messaging infrastructure. Event processing is done by a system of event processing agents (EPAs) — shown as colored boxes with arrows — distributed over the infrastructure. These contain event pattern recognition and processing rules about which we'll say no more here. Some EPAs filter events, and others aggregate patterns of lower level events into higher level events. And other EPAs process higher level events. High level events provide specialized views of the trading activity and are fed graphically to specialist role players.

One view might monitor the performance of a brokerage institution — e.g., how timely is its execution of stop loss orders. Another view might detect patterns of

events indicating suspicious situations where SEC regulations such as not trading ahead on a customer's order, may be violated.

So, CEP can be used to organize the clouds of events in the infrastructures of our enterprises into hierarchies. If you don't do that you just have a cloud of events in which you cannot see the business significance of anything. Indeed CEP is also used to create new events, like transaction views, that infer information from other events. The higher level events are complex in the sense that they are aggregated from patterns of lower level events. They contain data from the patterns that is needed to make decisions and perform various management roles. *Complex* in how they happen, yes. But also *simpler* to understand.

"Event hierarchy" is an application concept. Part of CEP is about the technology to organize event clouds into hierarchies in real time. This technology involves precisely defining patterns of events, detecting instances of patterns, and modeling causal, timing and aggregation relationships between events. Event patterns are used to create higher level events from the cloud of new events that is continuously appearing on the IT infrastructure. And the relationships between events are used in defining event patterns precisely. Event relationships are also used to reduce the search in drill-down explanations of how high level events happen, and why other events didn't happen. Once we, or our automated control processes, understand what is happening, the next step is to take appropriate action. CEP helps with that too. More on the topics in this last paragraph another time.