

Holistic Event Processing Future Applications of CEP by David Luckham

In the past couple of years, 2007 -08, complex event processing has achieved what the business analysts call “market traction”. The vendors are doing well nowadays. To quote from one of them recently,

“It has indeed been a good year. We’ve moved past the stage of spending most our time with “tire kickers” and now have almost more business than we can handle. This creates a new set of problems of course, but they’re the right problems to have”

Forrester¹ assessed the 2007 CEP market as about \$100 million in sales, and as “a market in the early stage of development”. Forrester expects the growth of the CEP market to accelerate driven by “a broad array of CEP applications across a wide range of industries”, among them the “*green revolution*” – an example they gave was applications of CEP to energy consumption in buildings.

Interestingly, Forrester also conducted an awareness survey over 1017 large companies. They found that 33% of them were already using CEP/EDA, and 38% were aware of, but not yet using, CEP/EDA. So, if all of this assessment is right, the vendors can look forward to a bright next few years. And they won’t have to keep on educating the marketplace about events and CEP either!

¹“ CEP Adoption Is Broader, Deeper, And More Business-Driven Than IT May Expect” by Charles Brett, Jan 2008

Now here's a question. How would you characterize the current generation of CEP applications?

Well, let's consider what we have. There are applications in markets such as algorithmic trading and financial services, patient flow monitoring in hospitals, routing and crew scheduling in trucking, monitoring service level agreements in call centers, consumer behavior in on-line retailing, and airline baggage handling, just to name a few. Most of these applications are processing multiple event feeds, maybe hundreds of them, and different types of events. And they're keeping track of state and timing too. The financial services applications are often computing very complicated statistics and averages over the data in events. These are not "simple applications." But they tend to apply to related sets of problems within a small number of types of events.

Notice, a small number of different **types** of events does not imply a small number of events. There can be a lot of events of one type, for example, say stock market quotes. I call the current generation of CEP applications "*small event type space applications*". A "type space" is simply a set of types of something.

Take an example, airline baggage handling. Typically the baggage handling system processes event feeds from sensors, either barcode or RFID readers. The sensors are placed at various points along the baggage transport system through a terminal from passenger check-in to the loading gates. The system also has a database (or state) containing information about the gates at which the planes are located, and their flight information, airline connections, etc. The system tries to detect as early as possible when a bag might be heading in a wrong direction. It should alert the baggage handling operations as soon as possible. Certainly, the system should stop a bag before it gets loaded onto a plane that isn't going where the bag is supposed to go, or cannot make a connection that would get the

bag to its destination. This isn't such a simple application as might seem at first sight, especially if one of the goals is to keep false alarms to the minimum, and another is to permit maximal use of an airline's carrying capacity. Bags might not always travel with the passenger, but should be there when the passenger arrives, wherever that may be! But my point is, the space of types of events being processed and output is quite small.

Now, consider all the other problems to which an airline, or group of collaborating airlines, might want to apply CEP. Figure 1 is a depiction of a small part what is involved. Here the system that is running the airline's operations is a dynamic EPN, that is, a network of event processing engines, some of them on board the planes, some in the airports, maintenance hangers etc. This system is processing a wide variety of types of event feeds, from the FAA, satellites, terminals and airports, global weather, passenger, crew and aircraft data, etc. The output types of events are equally varied. This application of CEP has the goal of keeping the airline running while minimizing costs and customer dissatisfaction!

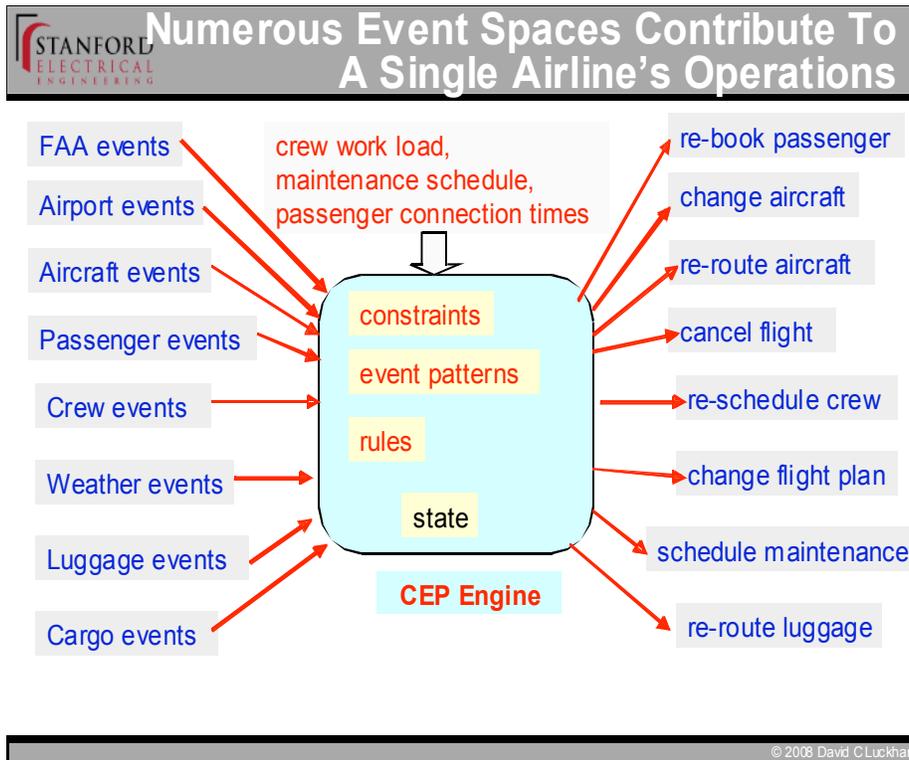


Figure 1: Future unified airline operations using CEP

An application like the one depicted in figure 1 can only be developed by an evolutionary process. First, smaller systems that solve parts of the whole operations problem, such as baggage handling, or aircraft maintenance, are developed. And later on they are unified gradually to form a total operations management system.

But why would anyone want to unify a number of specialized event processing systems? One motivation would be that there are optimizations that can be made in the overall operations of the airline. Obviously there are also a lot of risks! It would happen gradually over many years. And a lot of things have to conspire along the way to enable it to happen:

1. comprehensive format standards in EP,
2. technical innovations in EP in the special problem domains,

3. development of an event processing infrastructure that provides CEP capabilities over large event type spaces,
4. political will, e.g., collaboration between all parties concerned.

When a system gets to the stage where it is unifying a number of smaller event processing applications that deal with problems in different type spaces, and attempting to leverage the inter-relationships between them, I call it a **holistic event processing** application. A holistic event processing application is a large event type space system. The term comes from the Greek, holos, meaning all, entire, total. The system is attempting to achieve what Aristotle described as, "the whole is more than the sum of its parts."

If airline event processing applications ever develop from the small event type space systems being built today, to holistic operations-wide support systems it will take a long time, my guess is 10 years from now.

The world is full of small event type space event processing systems that might eventually evolve into global holistic systems. Often there are pressures towards unifying specialized systems such as economies of scale. Indeed, holistic systems will happen simply because at some point, when enough smaller systems are available, they become an obvious next step. But I believe that item 3 above is critical.

Here's an example of an application that has actually "gone holistic" but needs a holistic event processing infrastructure in order to develop to its full potential – the US Geological Survey's Natural Hazards Support System (USGS NHSS). See Figure 2.

Our world is wired! There are sensors in satellites airplanes and balloons, at the bottom of the oceans, in our forests and rivers and geological faults, polar icecaps, etc. They belong to separate monitoring systems, for tsunamis, atmospheric studies, earthquakes, and of course weather forecasting.

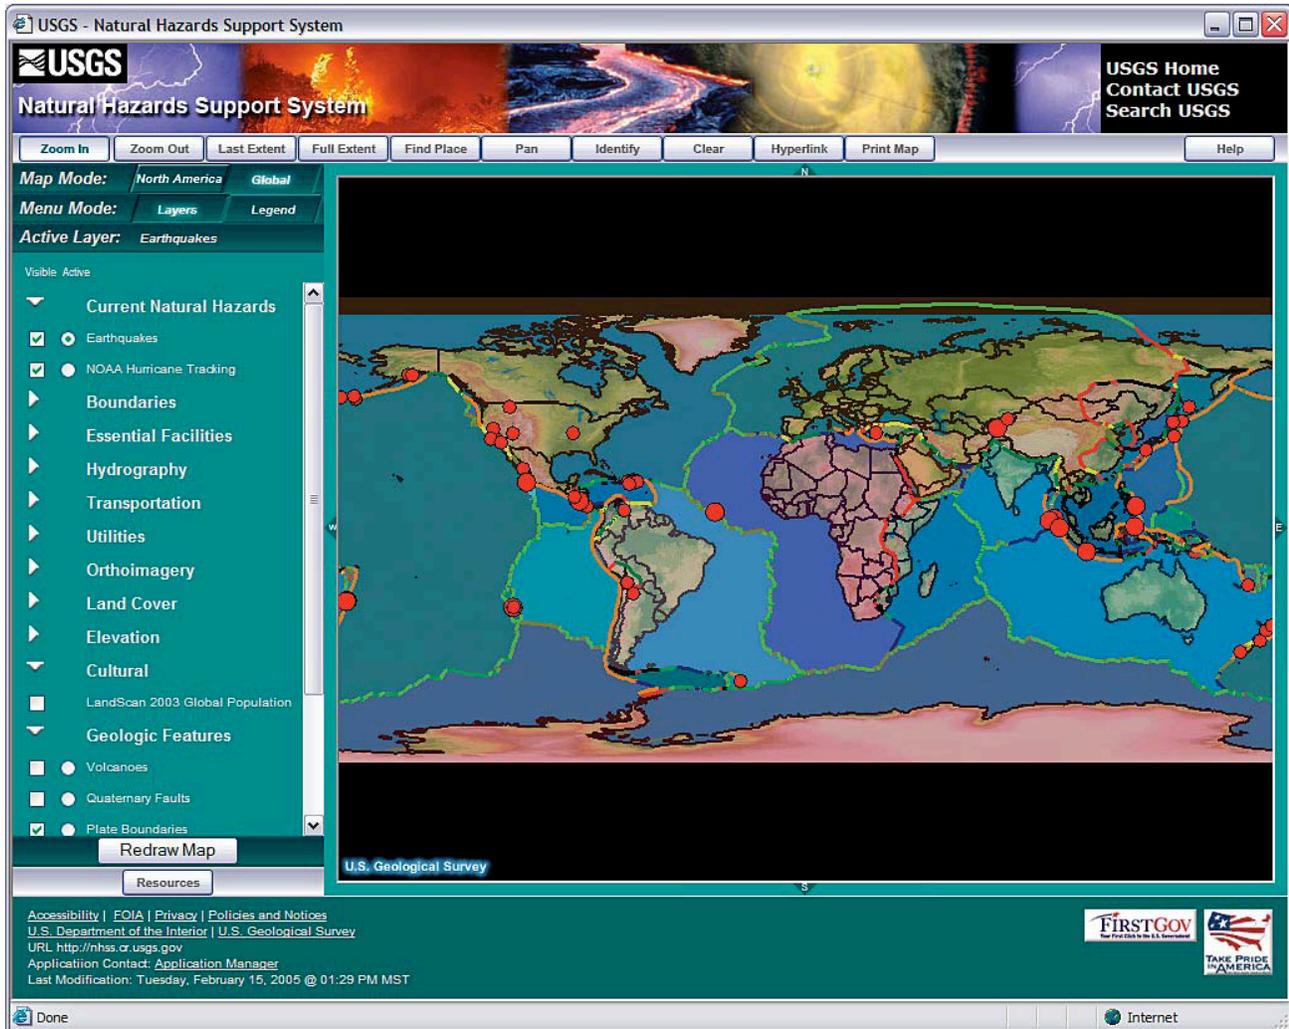


Figure 2: USGS Natural Hazards Support System Interface

The USGS NHSS attempts to make all this sensor data available in a single unified graphical interface – a map of the world. If you click on any location in figure 2 you'll get the information for that location displayed in a popup window. And it will be refreshed as new readings arrive. For example, see figure 3, NHSS refreshes readings from earthquake and weather sensors every 10 minutes, and from volcanoes and hurricane

sensors every 2 hours. It is a near real time system. And it is truly a holistic event processing system.

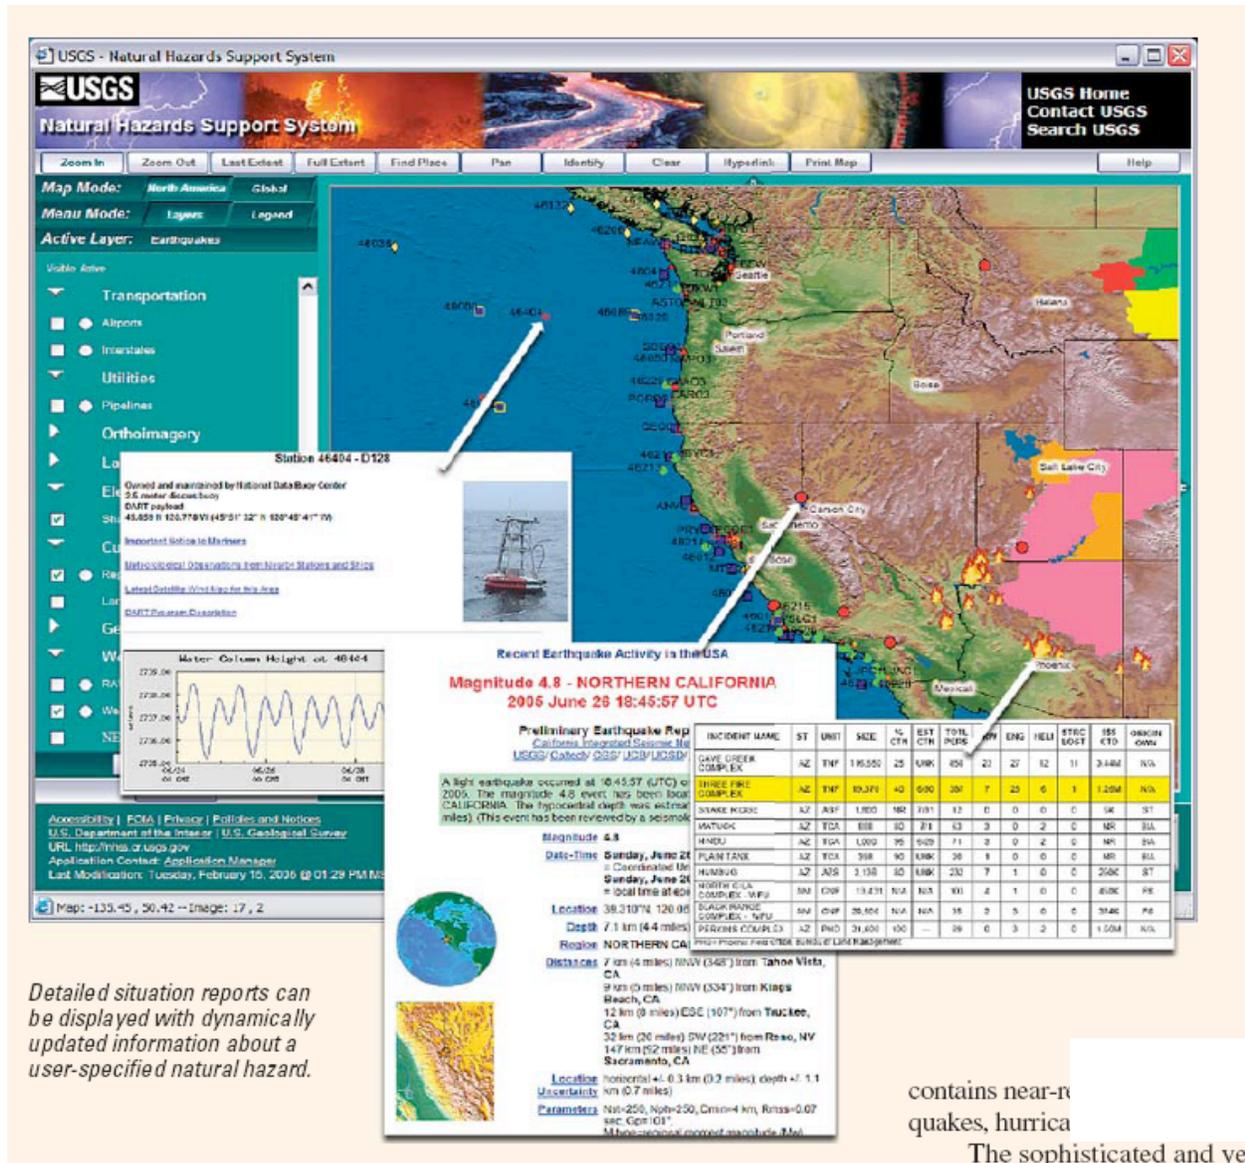


Figure 3: data from sensor systems displayed on the USGS NHSS

interface.

Figure 3 shows a state of NHSS in which there is a 4.8 earthquake in a California location and also a report of recent wildfires in that area. Also, there is a weather prediction of heavy rains coming in the same area. Putting together these inputs from three different earth observation systems, viewers might conclude that there is a probability of mudslides in the region. This illustrates an advantage in unifying the information from different systems. However, viewers might miss one of the data points, and there is no aid to make this conclusion available explicitly, nor to calculate probabilities. This is a missed opportunity! The users of NHSS include local authorities, and first responders such as emergency services. They are not necessarily specialists with a lot of time to spend clicking around the website.

The data provided by NHSS needs to be brought to a higher level. It would require recognition of patterns of events involving timing and spatial properties together with geological and demographical factors. That is something a holistic CEP infrastructure might bring to the table.

There are many examples of holistic systems in the planning or actual development stage out there. But, like the NHSS, all of them are being planned or built without any knowledge of event processing or CEP. There seems to be an implicit belief that networking will satisfy the needs for event communication, correlation, aggregation and abstraction, and the extraction of higher level operational intelligence. This is unfortunate because it means these systems will either reach an impasse or the builders will have to reinvent CEP for themselves. It is simply a matter of lack of awareness.

I will give some of these examples in future articles, each one of them has an important role in the future of our society.