

Federated Event Systems: The Event Web

K. Mani Chandy and Michael Olson
California Institute of Technology

Event-driven applications that are constructed as compositions of Web applications offer significant benefits. Just as mashups compose Web services to create added value, so too can compositions of event-driven applications create added value.

This article first reviews concepts about event driven applications, sense & respond systems, and presents metrics for evaluating and comparing such systems; then a few examples of federated Web applications are presented; and finally two examples of Business Application Monitoring (BAM) dashboards dealing with politics and with startup companies are presented as illustrations.

Event Driven Applications

An event is a significant change in the state of an enterprise or its environment. An event may be a toxic spill, the detection of radiation material in a container, a candidate winning an election or a steep drop in company stock price. Whether a change of state is significant or not depends on the user.

A sense and response (S & R) system detects or “senses” events and then responds to them appropriately and expeditiously. An S&R system sensed the arrival of a hurricane in Bangladesh and responded by evacuating people in the hurricane’s path, saving thousands of lives. Supply chain applications sense delays in arrivals of key components and respond by identifying alternate sources.

Event driven applications in general, and S&R systems in particular have *persistent* goals. For example, an S&R system for baggage handling has the persistent goal of ensuring that all bags get to their appropriate destinations. By contrast, a person surfing the Web may change goals frequently – searching for information about jobs at one instant and movies the next. One of the challenges in designing event driven applications is developing methods for specifying persistent goals: how do business users describe the events that should trigger a response and the appropriate response?

Almost all S&R systems are decision support systems: they help people who make final decisions. An enterprise’s organization and the quality of its people are often more important than technology in determining the efficacy of a sense and response application. Effective integration of S&R technology into people’s daily operations is critical.

The costs and benefit metrics by which S&R systems are compared are somewhat different from those of other IT systems. These measures include the following 5 metrics called the *REACT* metrics after their first letters:

1. **Relevance:** The relevance of information sent to the user by the system depends on the person's state. A tsunami warning is highly relevant to a tourist on a beach where the tsunami is about to strike but less relevant to the same person a day later when she is hundreds of miles away. Alerts about strikes in France may be accurate and timely but irrelevant to purchasing managers who are unaffected by French strikes.
2. **Effort:** The amount of effort required by business users – the operational people who execute responses – is a vitally important metric. They – stock or commodity traders, first responders in emergencies, supply chain managers – don't have the inclination or the time to define and fine-tune specifications of events or train machine-learning algorithms.
3. **Accuracy:** Inaccurate information, such as a false positive, has expensive consequences. A false tsunami warning gets beaches to be evacuated unnecessarily. High frequencies of inaccurate information result in the information being ignored. Alarm fatigue is a serious concern.
4. **Completeness:** Incomplete information has very high costs; for example, the absence of a hurricane warning can result in thousands of deaths. False negatives have much higher costs than false positives in most situations.
5. **Timeliness:** The benefit of detecting an event and actuating a response depends on the time available to execute an appropriate response. There is little added benefit in having millisecond responses when responses in minutes will do. There are, however, significant costs in responses initiated too late. Identifying the ranges of delay that are acceptable is an important aspect of designing S&R systems.

Discussions about IT systems are couched in terms of how well they measure up to desired goals, and we will use the REACT metrics as our measures. The metrics are evaluated in the context of the goals of the application; too often systems are measured in terms of a few numbers such as milliseconds of response time and numbers of messages processed per second.

An S&R system includes sensors and other data-acquisition components, processing agents, and responders. In BAM applications the initial system response displays information on a dashboard. An S&R system also has mechanisms by which business users specify the conditions for which they want responses, and the types of responses. The success of S&R systems depends critically on the ease with which each business user can tailor a system to that user's specific needs. One way to make specification easier for the business user is to build system versions for specific types of users as we will discuss shortly.

Federations of Events

Many people, all over the world, use Web search engines to find information, music, and videos. Searches use request-reply mechanisms: the user types in a phrase – the request – and the engine replies with links to information relevant to that phrase. Search engines have evolved to *Web monitoring*: the Web equivalent of BAM. Users specify persistent

interests and systems automatically update dashboards and send alerts. In some cases, systems respond by updating calendars or by executing other activities. Web monitoring services are becoming increasingly common and are likely to become ubiquitous. Receiving a message on a cell phone that an event has occurred is more convenient than searching the Web repeatedly.

The Web was primarily thought of as a passive system in which users initiated activity – such as searches – to which the Web, and its associated engines, responded. Now the Web is increasingly thought of as active; the Web and associated software agents find, organize and display information on behalf of users. The speed of transition of the Web from passive to active depends critically on the ease with which end users – business users, school children, and retirees – can tailor Web monitors to their unique needs.

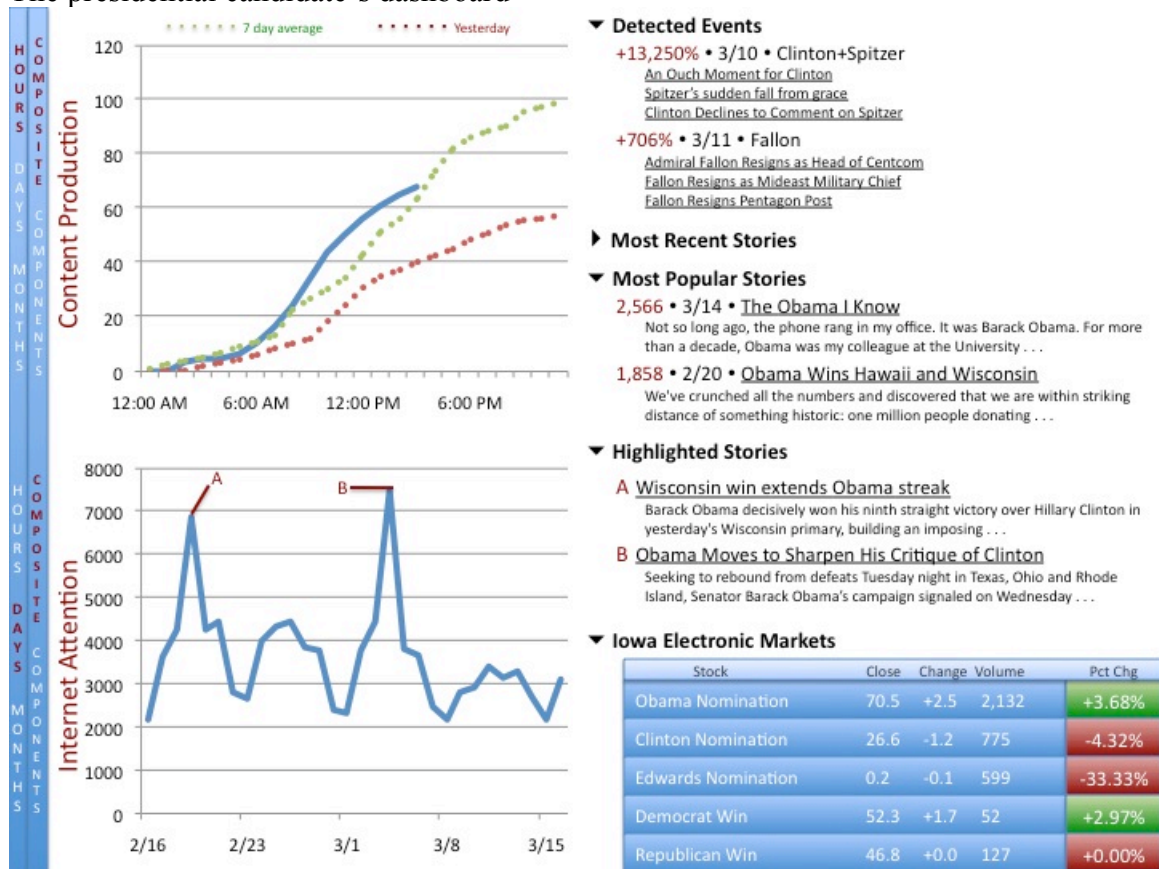
Web S&R applications are S&R applications obtain most of their data from the Web. We don't have space here for a detailed analysis of each of the REACT metrics for different designs of Web S&R, but we encourage readers to evaluate the Web S&R applications that they use in terms of these metrics.

Web data sources can be active or passive. A client can poll a server, using a request-reply protocol to extract information. Alternatively, information can be pushed to a client by means of RSS or ATOM streams or by using other data protocols. A passive data source can be made to have an active interface by an agent that polls it periodically and actively sends information about changes in successive polls.

There are a variety of active Web data sources. Google Alerts and Yahoo! Alerts are two examples of similar services. Google Alerts notifies you by e-mail when the results to a web search change; in other words, by working with your e-mail account Google creates a persistent query out of your web search. Yahoo! Alerts provides alerts on a variety of events from news, television shows, stocks, flight fares, and fantasy sports to missing children reports. In addition, Yahoo! Alerts allows you to receive notifications by e-mail, instant message, or cell phone. Yahoo! Pipes can also be integrated into the alerts system, which means that you can compose your own alert out of almost any content that is publicly available on the Web.

Some data hosts have also made the move to providing active notifications of updates. For instance, by connecting to the persistent SixApart update stream, consumers of the stream receive real-time updates to the SixApart universe – including LiveJournal, TypePad, and others – as they happen. Active data sources are still fairly limited, but there are many more passive Web data sources for which active interfaces can be built. Digg, del.icio.us, Technorati, and Twitter are all examples of passive data sources that offer APIs to access their information; in doing so, they enable easy integration with other data sources.

The presidential candidate's dashboard



The presidential candidate's dashboard pulls together many different data sources to attempt to give a one-screen overview of the current political state. The graphs in this dashboard highlight two things: the volume of content produced that discusses relevant issues and the amount of attention that content has received in the Internet community. When there are substantial movements in either the amount of content or the amount of attention, events are generated that point to the likely cause of the spike.

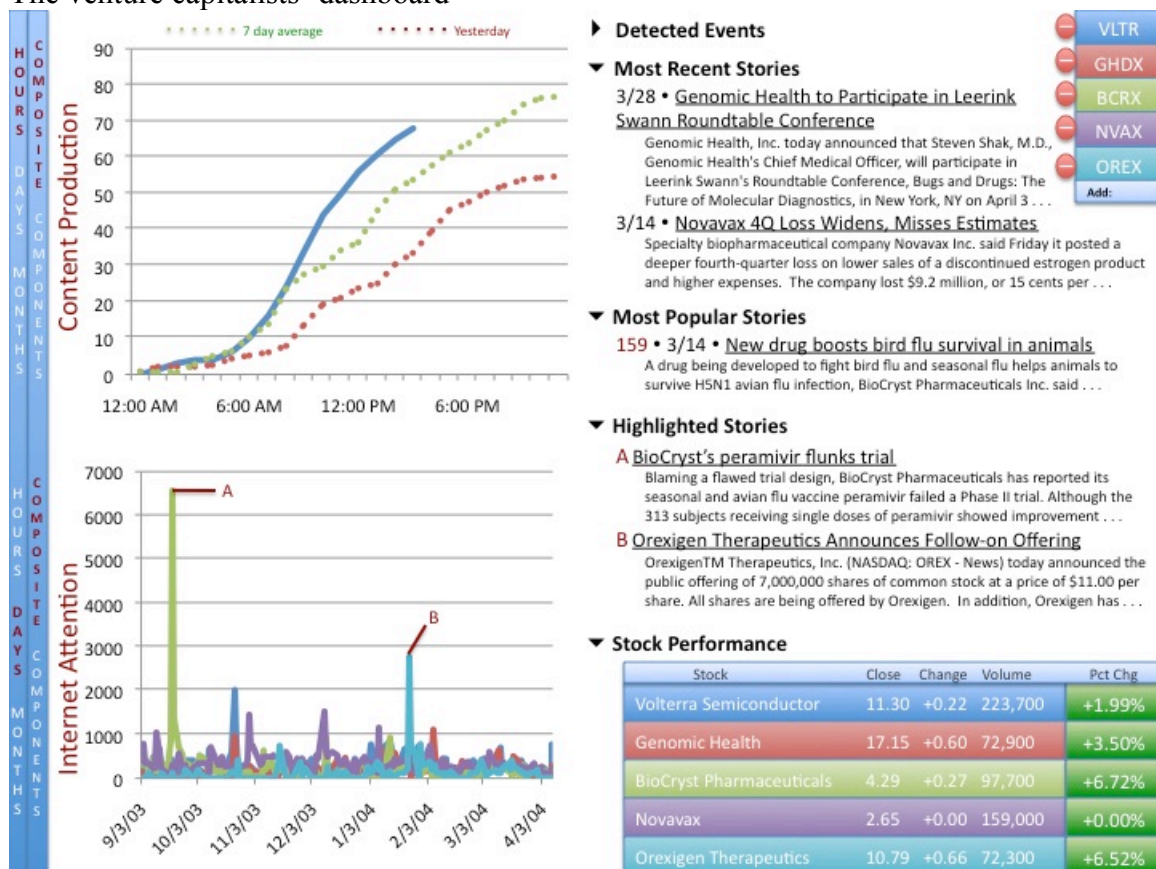
Both the content graph and the attention graph are composite statistics; they are composed from many different data sources. The content index deals with news articles, YouTube videos, and some blog posts, while the attention index deals with blog posts that reference the content, del.icio.us links, Digg hits, Wikipedia edits, Twitter references, and the number of comments on all of the above.

Each graph can be designed to monitor whatever keyword or topic is most relevant. The above graphs focus on the keyword "Obama"; hence, the Highlighted Stories section attempts to find stories that explain the peaks in attention that Obama received on specific dates. The dashboard can also identify what content is the most popular at the moment based on the same components as the attention index. When a particular story receives a substantial amount of interest, it will appear in the "Popular Stories" section.

Paired keywords have a high potential for identifying events. This is because the frequency of conversation on two normally disparate topics barely rises above the level of noise, but when an event occurs that ties those topics together, the shift in volume is of a magnitude that cannot be missed. Above, we used the keyword pairing Clinton and Spitzer to demonstrate this shift in volume. The percentage used identifies the percent change in the volume of content with both keywords.

Additionally, the dashboard incorporates other useful data sources, like the Iowa Electronic Markets for presidential nominations, and could easily incorporate even more. Given the number of data sources available, condensing the data into a readable format is more important than having all of the data available. This is why indices were used that combine many different data sources. In a live version of the dashboard, it would also be possible to drill down into the specific components of activity and content, explore how a story became popular, and follow the story's conversation as it made its way through the Web.

The venture capitalists' dashboard



This dashboard incorporates the same components as the presidential candidate's dashboard, but, rather than focusing on politics or a keyword like "Obama", it selects companies relevant to a particular venture capitalist. The output then is relevant only to those companies. We can see how the content produced about them, and the attention

they receive, shift over time. We can also watch as their stock prices fluctuate, and attempt to correlate stock performance, and change in stock prices, with events that occur such as news stories.

You can see from the highlighted story about BioCryst that attempts to explain the spikes in attention and content can be made for these companies as well. In the live version, adding or removing companies from the comparison, as well as juxtaposing graphs like stock price and attention, would provide for interesting analysis and monitoring of company performance and popularity. A dashboard made for a particular sector, watching not only companies in that sector but also the general attention that sector receives, would have great descriptive power for emerging sectors.

This brief article makes the point that event processing and S&R in general, and Web dashboards in particular, will become ubiquitous in a few years. A broad analysis of event processing will be presented at the upcoming Gartner Conference on Event Processing to be held September 15 – 16, 2008 in Stamford CT (see www.gartner.com/it/page.jsp?id=616710).

Acknowledgment: Research on event processing for societal applications is supported, in part, by the Caltech Social and Information Sciences Laboratory (SISL).