CHAPTER 1

Event Processing and the Survival of the Modern Enterprise

All the world's information is at your fingertips—but can you make use of it?

—Vinton Cerf, 2005

You probably think that every twenty-first-century enterprise uses events and event processing in its business operations. That would seem obvious, given our information-driven world, which is inundated with sources of events from just about everywhere. But you would be wrong! The truth is that a lot of businesses think they use event processing. And, yes, a lot of them do—in their network management and communications. A few businesses go further and use event processing to drive some of their operations, such as supply chain management or consumer relations. And then there are the electronic stock trading and online gaming industries, both of which are totally event driven, but those are niche markets for event processing.

Many times, it turns out on closer inspection that businesses could make much greater use of the events already at their disposal in their business operations and planning. They could do a lot more with today's event processing technology than they currently do to improve the running of the enterprise, their awareness of the business environment, and consequently their business decision making—and it would benefit them greatly in terms of their competitiveness and profitability if they did. Indeed, for some of them, adopting the latest event processing technology

in their intelligence gathering and business planning may become a matter of survival.

This book has four goals:

- Firstly, to explain the concepts of event processing and to answer basic questions such as "what do you mean by an *event*?"
- Secondly, to describe strategies for applying event processing in business and enterprise management, not only to run business operations but also as a tool for business intelligence and a basis for planning
- Thirdly, to describe the progress and probable limitations of commercial event processing technology
- Finally, to explore some of the future trends in event processing and its pervasive supporting role in very large–scale information systems

We include a short survey of how event processing has been a basis for different technology areas from discrete event simulation to business process management over the past sixty years. This gives some background about the multiple roles event processing is playing nowadays in everything from weather forecasting to operating a business or running a government.

The final chapter outlines some of the longer-term developments in event processing technology and the roles it will eventually play in the information infrastructure of our society. Many future applications are quite easily predicted now, and the only surprises are in finding out how long they will take to actually happen!

Four Basic Questions about Events

Here are four questions that every CIO, CTO, and business manager should have an answer to:

1. What are events and which ones are important?
2. Why invest in event processing?
3. What are the main concepts in event processing?
4. What kinds of enterprises have bought into event processing technology so far?

These are pretty basic questions—but they're the kinds of first questions people often ask when they're wondering if event processing has anything to offer them or their business.

What Are Events and Which Ones Are Important?

An event is simply something that happens. That's the common sense meaning, the one in the dictionary. We'll deal with the technical computer processing meaning later. Of course there are many indicators of events. Anything like a message on a news wire service, an email, a sensor signal, a text message, or a newspaper report can indicate an event that has happened.

Events come in all sizes. Some are small events, like getting a text message on cell phone, and others are very big events, like World War II. Historians tend to write about big events, but we live day to day by small events!

Business events are the events that affect our businesses. Those are the kinds of events that concern us here. Now, there is great disagreement as to which events *are* business events, and even more disagreement as to which events are the most important. To illustrate the contentious nature of business events, here is the top ten in *TheStreet.com*'s list of the top 100 U.S. business events in the twentieth century, published in May 1999:[1]

1. Eisenhower creates the interstates: June 29, 1956.
2. Intel invents the single-chip microprocessor: 1971.
3. The Federal Reserve is formed: 1913.
4. The Great Crash of 1929: October 24–29, 1929.
5. Equal pay for equal work: June 10, 1963.
6. Ford introduces the assembly line: 1913.
7. Kaiser's World War II shipyards surpass all expectations of production: 1942.
8. The first Wal-Mart opens: 1962.
9. The current bull market begins: August 1982.
10. Carrier Engineering is founded, beginning the commercialization of air conditioning: 1915.

And, just in case you agree with all of these top ten, which is very doubtful, try to guess number eleven (don't look!).[2] Obviously, an event remains an event, something that happened. But its importance can vary over time, because importance is a subjective measure and relative to what else happens. For example, number 9, "The current bull market begins: August 1982," would have to be reworded nowadays as "the 1980s bull market," because that bull market was replaced by a recession, and

[1] www.thestreet.com/story/747965/the-basics-of-business-history-top-100-events-at-a-glance.html

[2] "11. Reagan is elected: 1980." (TheStreet.com)

then another market about which the experts disagree. Perhaps "the 1980s bull market" would no longer make the top hundred!

Events have *attributes*, such as the time at which they happened, how long they took to happen, and which events caused which other events. Attributes (such as timing) and relationships between events (such as causality) are important clues used in processing events to judge which ones are most important "right now" for the business. We'll deal with attributes and event relationships in Chapter 3.

To illustrate how an event's effects, such as causing other events, can increase its importance, consider for a moment this event from *TheStreet.com*'s top 100 U.S. business events:

> 42. *The Jungle* is published: 1906.

Why on earth is the publication of a novel by Upton Sinclair considered an important event in U.S. business?

Well, Sinclair wrote this novel to highlight the plight of the working class and to remove from obscurity the corruption of the American meatpacking industry in Chicago during the early twentieth century:

> Upton Sinclair originally intended to expose "the inferno of exploitation [of the typical American factory worker at the turn of the 20th Century]," but the reading public instead fixated on food safety as the novel's most pressing issue. In fact, Sinclair bitterly admitted his celebrity rose, "not because the public cared anything about the workers, but simply because the public did not want to eat tubercular beef." Their efforts, coupled with the public outcry, led to the passage of the Meat Inspection Act (1906, amended, 1967) and the Pure Food and Drug Act of 1906, which established the Bureau of Chemistry that would become in 1930, the Food and Drug Administration.[3]

So, event number 42, the publication of a novel, had far-reaching and permanent effects upon American society and the regulation of its businesses by the government.

However, Figure 1.1 illustrates that an event's causal effects can be temporary or may be superseded by later events. Of course, looking back at an event's causal effects is interesting, but for a modern business it is much more important to predict likely causal effects in the future of events—maybe several events—that are happening now! More on this later.

[3]Mark Sullivan, *Our Times* (New York: Scribner, 1996), 222.

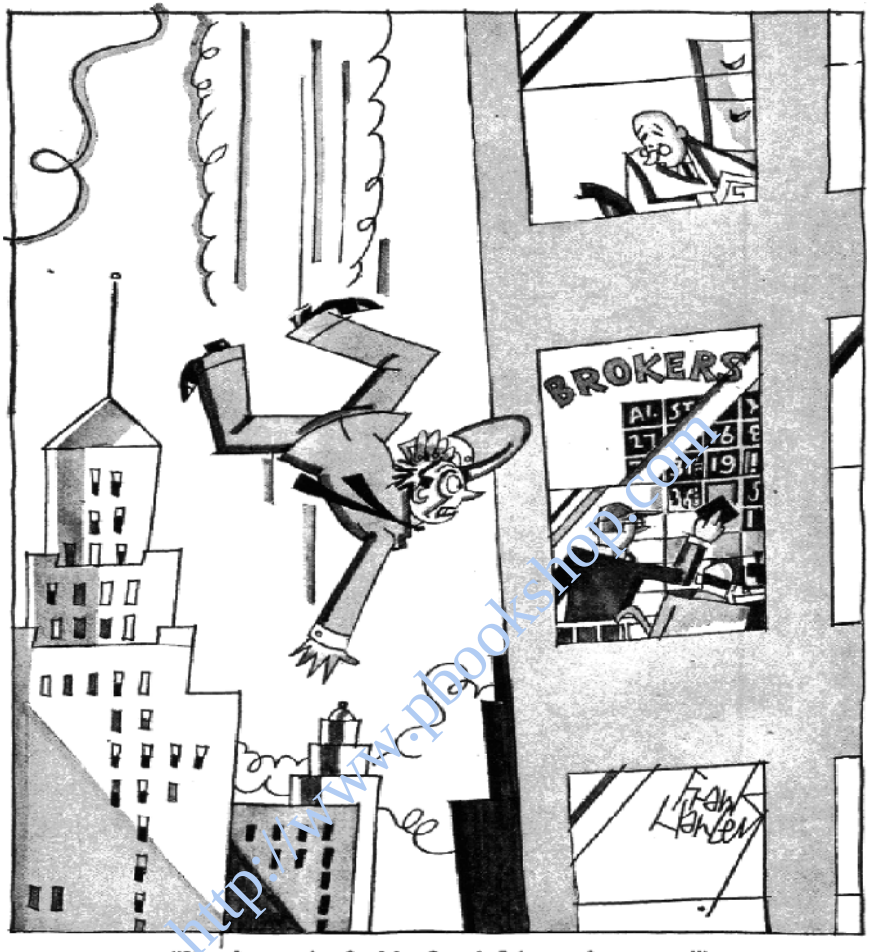

FIGURE 1.1 "Up three points," by Frank Hanley, January 10, 1930
Source: www.archelaus-cards.com/blog/2009/01/12/the-great-depression-in-cartoons-part-1-1929-30/

Why Invest in Event Processing?

Events are carriers of *instant information*. And event processing is about having the ability to take *instant action*. Sometimes the usefulness of the information events carry is short lived. Putting events in databases for action tomorrow is yesterday's mode of doing business.

Events are arriving all the time from a multitude of sources. This is the era of the *real-time event*.[4] Our businesses, our transportation systems, and our government agencies depend upon all manner of events. Some examples are:

- Internet messages
- Stock market feeds
- Online reports from the company's departments
- Customer orders and enquiries on the business's website or by phone
- Cell phone text messages
- GPS tracking updates from a delivery truck fleet
- RFID tag readings from sensors in inventory and shipping
- Satellite-based location feeds from aircraft navigation systems

and many hundreds of other sources of events. All kinds of events are arriving continuously and driving our businesses. Often a business must react immediately to the events it receives just to keep the daily operations moving. So why not explore other ways to make use of them?

Events run our personal lives too! When we go home after a hard day's business there are all the traffic events that we have to put up with. We process those events right then—when they happen—to decide which route to take going home. And at home there's the evening news, weather reports, and even more events. Moreover, these events depend upon other event sources. Weather reports, for example, depend upon satellite communications, radio signals from deep ocean buoys, automated weather stations, and much else. It's an event-active world that we live in.

Whatever the business, many different kinds of events happening "right now" have an impact on it. There are so many events and so many sources of events out there that *the event cloud* is a term that is often used to describe our business environment.

It's an Event Driven World

Just consider a few numbers from the Technology market research organizations:

- There are 5 billion mobile phones worldwide and counting.[5] Many of them are smartphones, and 15 percent are Internet enabled (September 2010).

[4] *Real time* is a much overused term. What we usually mean by it in business processing is *right now*. And that means *this instant*, *immediately*, or *taking as little time as makes no difference*.

[5] The number of mobile phones in use worldwide has topped 5.0 billion, boosted by soaring demand in emerging markets India and China, according to a study by Swedish telecoms giant Ericsson in July 2010.

- In December 2009, texting was being used at the rate of 152.7 billion text messages per month—and that's just in the United States.
- The Internet has 2 billion users, as of June 2010.
- We are also seeing all manner of new kinds of event-enabled twenty-first-century business models: eBay, Amazon, Dell, Google, Yahoo, MSN, AOL, iTunes, Twitter, Facebook, Skype . . . a lot of them are very successful!
- Skype has around 560 million registered users and 8.1 million paying users. People spend an average of 520 million minutes every day talking on Skype.
- Facebook has more than 500 million users.
- Global collaborative communities, some social and some scientific, are springing up: Social Networks, OLPC, Google Earth, the large Hadron Collider, and the like. Some of these will be business-oriented. And their business potentials are being explored, as we speak.[6]

All of this event-enabled stuff represents a new way of doing business. It has been emerging for some time. And now it's here!

The obvious conclusion is that we have to use events. And the real question is how to make the best use of them.

The fact is a lot of us are already using event processing in our businesses. But I'll bet we could do more! Maximizing the use of the information carried by events can be of paramount importance, both from an offensive and a defensive viewpoint. Businesses should try to maximize their use of the events they already receive. And then plan to use more sources of events, ones they don't normally deal with. The question is how.

Let's start by discussing the ways in which events can be useful at the higher management levels of a business. Later in the book, we'll explain how to do it with modern event processing technology.

As we said before, we can think of an event as a message together with additional data. The added data are about who sent the event, where it came from, its duration and the time it was sent, how it got to us, and how it is related to some other events. That extra data, by the way, are very useful when it comes to building event processing systems. It is crucial in helping to figure out things about the events you're receiving, such as whether you should look at some other events at the same time, or whether they're coming from sources you normally trust and rely upon or sources that might be erroneous or corrupted in transit.

[6]Adam Shell, "Wall Street traders mine tweets to gain a trading edge," *USA Today*, May 4, 2011. www.usatoday.com/money/perfi/stocks/2011-05-03-wall-street-traders-mine-tweets_n.htm

You do have to be careful when you use events. But that's a fact of doing any kind of business today, and the beauty of event processing technology is that it gives you the means to do exactly that—be careful!

Patterns of Events

Each event might carry only a little piece of information. To make sense of that information, one may need to look at some other events as well, ones that are related in some ways to the original event. Events tend to arrive in patterns, mixed up with unrelated events.

When some patterns, perhaps containing several events, are looked at together and their total information is analyzed, they can often tell us what is happening or going to happen, where it is happening, and why. It might be a pattern containing only one event that's significant at a given moment. Or it might be a pattern containing hundreds of events, arriving in a millisecond or maybe spread out over days or weeks. We may not know in advance how long the pattern will take to happen. But event processing technology can help us detect the patterns we want to know about.

Also, event sources can be globally distributed and intended for a large number of users; or they may be local, company-specific and confidential to the company. We will deal later with sources of events and which sources may be relevant to a business or enterprise manager.

In the running of an enterprise, *patterns of events* can often play a critical role in providing vital information for taking action *right now*.

Principle 1: Patterns of Events Add Up to Actionable Information

Each event may carry only a little piece of information, but when the information in a pattern of multiple events is aggregated with the contexts in which each of the events happens, some patterns will yield intelligence upon which a business or enterprise must act.

Even so, many enterprises are asleep at the wheel. You will find that event feeds from the sales department, or inventory, or suppliers, or the company's website, drive its operations. The company reacts to these events as they arrive in various ways, sending out products, reordering stock, answering enquiries, scheduling services, and so on.

But this is *reactive* behavior to the events it receives, simply to keep its normal operations in motion. This reactive behavior is the equivalent

of driving a car along the road on automatic while daydreaming, or even worse, dozing off at the wheel. Most enterprises are not making the maximal use of the available events and event processing techniques.

Here are three activities in which detecting and analyzing patterns of events can be an aid to higher level business operations. These are areas where it can really pay off to invest in real-time event processing to maximize the information you extract from the events that are available to you. The same event processing technology can be applied in all three areas:

1. *Knowing* how well your enterprise is performing in relation to your expectations and the competition[7]
2. *Detecting when* what you need to know happens, and then using that knowledge—quickly!
3. *Extracting the information you want* from the cloud of events available to you

But I hear the warning bells already ringing. What's it going to cost? Well, it doesn't need to be a lot of expensive software just to get started or put a toe in the water. We'll get back to that later. And one of the points we're making is that you're probably already using some event processing—it's a matter of going a little further with the events and the software you've got. When you see the payoff, you'll want to analyze whether it is worth investing in more technology.

Know How Well You're Doing

"No man is an island, entire of itself."[8]

A primary reason to invest in event processing is to keep tabs on how well your business is doing. To be sure, you're already using events to drive some of your business processes. Going beyond that normal everyday operation, events can be used to provide an up-to-the-minute understanding of the performance of the business.

Start by using events to track how well your business is doing and to give a *right now* picture of the competitive environment in which you

[7]Corporate performance management is often considered as part of *business intelligence* or BI.

[8]John Donne, "Meditation XVII," 1623.

are operating. Make better use of the events you're already using in your business. For example:

- The events you normally use to drive your business processes
- The events those processes create
- Events you receive everyday from the business environment

Just as importantly, these events can be used to detect *exceptions*—situations where you receive events you don't expect, or don't receive events you do expect—and take appropriate action.

Use All Event Sources

You should aim to gradually extend the business intelligence application of event processing towards the goal of monitoring the information in all your available sources of events. Of course this can't be done in one big installation of new technology. You will need to prioritize which sources you monitor first, second, third, and so on. An event monitoring facility is built up gradually—and is never finished!

You will start by monitoring the events in your own operations within the business. And then extend that monitoring to your interactions with customers through sales, orders, response to advertising, and so on. After that, you might attempt to monitor other event sources, such as the activities of your suppliers, customers, and your competitors. Consider Example 1.1.

Example 1.1

A company plans to introduce a new product. But while production is starting up, some of the company's event sources deliver information that may require changing production plans: (1) unusual back orders causing delays at some parts suppliers, (2) problems hiring skilled personnel at a subsidiary's assembly plant, and (3) transportation workers on some supply routes rumored to be about to strike.

Each of these events happening in isolation may signify a need to modify the company's production or marketing plans. We call them *actionable events*.

But there is another lesson here. When actionable events from different event sources happen together within small time windows, say within a week, the aggregation of those events must be considered, not just each event by itself. Higher levels of business planning and policy may be affected.

In most companies today, this kind of monitoring of multiple event sources is either not done at all or it is done in a very haphazard manner, and the results may not be aggregated. But there are already some industries where it is state-of-practice. News services and financial services are prime examples.

When monitoring of event sources is done properly, it is done using level-wise techniques for efficiency. For example, in the news industry, the first level of incoming raw event traffic is filtered by automated methods for the presence of keywords such as names of companies, newsworthy people, countries, weather events, stocks, and many other topics. One interesting keyword topic that is the subject of experimental use is a lexicon of sentiments that tend to crop up in news items such as "surprising," "hopeful," "unbelievable," "encouraging," "disappointing," and so on. Each sentiment is assigned a positive or negative value that is factored into the total newsworthiness of the posting. This kind of event processing is also being applied in stock trading algorithms.

The traffic that is caught at the first level by the automated filters then goes through a second level analysis for business relevance. This analysis is performed by humans with the aid of various predictive analysis tools. As Example 1.1 shows, this second level may involve aggregating patterns of events from several sources—a typical event processing operation, which can be applied using many of the tools currently on the market. The results go to a third level, which is a second level of human decision makers. It is here that the company uses event inputs to make business decisions and take actions—*right now*.

Automated monitoring of the sources of business events—possibly with humans in the loop—is essential to ensuring that a business is agile and capable of changing plans to save costs or to take advantage of opportunities in a competitive environment. It's not only about running the company efficiently; it's also about instant awareness of changes in the business environment—and what the competition is doing.

Note: Automated monitoring of all sources of business event inputs to an enterprise goes well beyond the kinds of media monitoring offered by some specialty companies. But if you do decide to hire someone to do your business event monitoring for you, then you must be able to judge how good a job they're doing.

Detect When What You Need to Know Happens

A second use of event processing in your business, and a big reason for making some investment in this technology, is to avoid the business being blindsided.

The goals of this activity are different from keeping track of your company's performance and the competitive environment—but you'll use the same technology to do it. The goals are more about keeping up to date with technical and product information than about running the business operations. Example 1.2 illustrates this usage.

Example 1.2

The company in Example 1.1 must deal with another situation during its introduction of the new product. A number of events that it has received, taken together, indicate that a competitor is planning an advertising campaign to publicize an improved technology for a similar product. The competing product will perform faster and cost less than the company's own product.

Obviously the company should have known about the technology improvements as early as the competition did. It should have known about the competitor's plans too. And there's a good chance it could have found the information in the sources of events it had access to but didn't monitor, or didn't monitor for information beyond what was needed to run its business. Learning about it now is probably going to be costly.

The old way of dealing with these kinds of issues was to use industrial spies. But today we use automated event monitoring technology to do some of our spying. And in doing this, sources such as social network sites can be useful. Let's call this process *information detection and gathering*.

A business must keep aware of new technical developments or product-marketing information within its sphere of activity that may affect its planning.

Today, news can happen in seconds. And quite often that information might be right under the company's own nose, contained in the cloud of events coming from the event sources that it either routinely processes to run its business or could use for business intelligence if it chose to do so. Being aware is a problem that event processing can help to solve.

Continuous Watch for Information

This second area of activity, *detecting when what you need to know happens*, is effectively a *continuous search problem*. That is, *being aware* is best implemented by a *continuous search* that never stops until you tell it to!

Being aware of new developments is complicated by the fact that there's lots of information out there that might affect the business, but we never know all of it—only parts. We might only have a few clues about this new information we should know about, such as the topics, possible sources, probable timing, who the likely competition is, or where our business is most vulnerable.

But it is also true that a successful business usually has a pretty good idea of what to watch for, based upon previous experience. One can approach the awareness problem by first taking steps to raise the level of watchfulness within the company. Start by making *watchlists* of the clues the company should look for about what's likely to happen. Use these lists to direct attention toward the events you should know about.

Personnel should be encouraged to use the watchlists.

However, a company, or a government for that matter, cannot rely only on its personnel.[9] A level of automated watching must be considered. If the company decides to invest in event processing technology, the watch lists will be useful input to an application for continuously monitoring sources of incoming events for technology and product information. But beware the limitations displayed in Example 1.3.

Example 1.3: The Lesson of Too-Late Information

The case of the 2009 Christmas day bomber illustrates the lesson of too-late information. The bomber traveled from East Africa to Amsterdam for the explicitly stated purpose of committing an act of terrorism on the United States. His father had given a U.S. consulate in Africa this information about his son. By the time the bomber arrived in Amsterdam, his name had been placed on airport watchlists. However, the process of updating watchlists made revised lists available only once per day. The bomber was able to board a United States–bound airliner because the Amsterdam watchlist had not yet been updated that day.

[9]For example, four days after the Christmas day 2009 bombing attempt, President Obama said publicly that Abdulmutallab's ability to board the aircraft was the result of a systemic failure that included an inadequate sharing of information among United States and foreign government agencies. He called the situation "totally unacceptable" (BBC News, "US President Obama notes 'systemic failure' on jet bomb," December 30, 2009. http://news.bbc.co.uk/2/hi/americas/8434275.stm)

Principle 2: Use It Right Now!

Businesses must design their processes to make the fastest use of actionable knowledge when it arrives. Event processing is about taking immediate action *right now*.

But first of all we must find that information.

Investing in search technology requires careful research by the technology officers of the company. Here are a few observations about *search* and *continuous search* technology.

Continuous search is part of the modern Internet-based way of doing business, an essential business intelligence technology. Everyone uses it with varying degrees of effectiveness, from individuals, to companies, to political organizations, to governments.

Secondly, at the moment we're stuck with far-from-perfect search tools. A useful Internet search engine led to the rise of the modern business world's most successful twenty-first-century enterprise, Google.[10] But let it be said that Google does not have anywhere near a perfect solution today—and won't for a long time to come.

All we have at the moment are rather simplistic kinds of searching tools that will answer only simple questions, deliver a lot of irrelevant answers, and sometimes will not find an answer at all. There's plenty of room for improvement. In fact, a lot of work is going on, usually in secret, to develop more sophisticated search engines based upon using the semantics of data to deliver intelligent answers to searches.

Thirdly, the *continuous search* problem requires a technical leap forward from what Google search can do at the moment. You can build a rather crude continuous search on top of Google. But to do it well, so that you can have confidence in the results, requires new event processing technology. There are two reasons for this:

1. Much of the information that needs to be searched is dynamic, like stock market trading feeds or the Internet. Its usefulness is often short lived. Putting real-time event feeds into a database to be searched later

[10] Google was created by Larry Page and Sergey Brin when they were Ph.D. students at Stanford University as a research project to create a better form of search. The domain Google.com was not actually registered until September 15, 1997, and the company Google, Inc. was incorporated on September 4, 1998.

is too slow—unless it is done very cleverly with in-memory databases, it will lose any advantage of being useful for *right now* action.

2. The information we need is often contained in patterns of many events, not simply single individual events. So we need an *event pattern* technology as the basis for continuous search, which means employing event processing engines in the searching process.

There are some encouraging developments in automated methods of detecting what you need to know as it happens. First, the new search technology we're talking about is beginning to appear, for example, in the form of Internet protocols like the Extensible Messaging and Presence Protocol, XMPP,[11] which can help to build information systems that receive the data they want whenever that data become available.

Secondly, the same event processing technology and principles that are employed in monitoring the company operations and business environment—how well we're doing—should be adaptable to work when the goals of the event monitoring are widened to continuously searching for technical product information. A single investment in event monitoring technology may well solve both problems.

Principle 3: Level-wise Watch for Actionable Information

Event sources should be subject to different levels of monitoring corresponding to the different goals of the enterprise—for example, (1) to know how the company is performing and (2) to detect information on technology and market developments. The goals of monitoring are separated into levels so you don't slow the system by taking the monolithic approach and doing everything at once.

What to Watch

There's another dimension to both the search and continuous search problems that is becoming increasingly important. Much business information is now *real-time*—which for business purposes means *right now* time. That information is increasingly carried in events in various kinds of networks, news media, and messaging systems, which we will call simply

[11] For more information on XMPP, the XMPP Standards Foundation website has an excellent introduction: http://xmpp.org/about-xmpp

event media. New events are arriving all the time. And new information media never stop appearing.

Keeping up with changes in the media that carry event traffic is taking on a dynamic component. Search engines must be able to deal with this challenge. To do this—to direct the searchlight on the latest information sources—event processing techniques will be employed. That's another reason to watch what's going on in event processing technology!

Event Media

All electronic sources of events, such as the Internet, messaging networks, cell phone text messages, newswire and news aggregation services, social websites, and emergency warning services, are examples of today's sources of events.

Event media are constantly changing and evolving, perhaps at a slower pace than the information they carry, but they are changing nonetheless. At the same time, the need to access that information as it becomes available—*right now*—has become an imperative for all enterprises. That is why "detecting it *when* it happens" requires a new kind of continuous search based upon event processing.

Principle 4: The Constant Evolution of Event Information and Event Media

The world's event information, and the media that carry those events, are constantly evolving and changing. The correctness and usefulness of the information is often short-lived. And new media are always appearing.

Event Processing in Use

Here are some short examples to illustrate how event processing—and more accurately, *Complex Event Processing (CEP)*[12]—is being used in solving real-time problems. CEP is a key technology in building event detection and reaction capabilities. We describe situations that develop and require solution within various time windows. And then we describe patterns of events

[12]We place the term CEP in its historical context in Chapter 2 and define it precisely in Chapter 3.

that might be used in event monitoring systems to detect these situations as they happen:

1. A problem in the baggage handling system of an airline (30 seconds—1 minute).
2. The unexpected spread of an infectious disease, indicating a need to trigger national public health control measures (1 week–1 month).
3. Weather conditions cause delays in a company's supply chain, and the company must reorganize the schedules of the sales force in the field (8 hours).
4. Unauthorized message traffic on the company's internal networks indicating the possibility of a spyware installation (variable time windows, anywhere from 5 seconds–4 weeks).
5. A credit card company detects suspicious use of the cards of some of its reliable customers (1–12 hours).
6. A rogue nation or terrorist organization is accumulating the know-how and materials to build and deliver dirty atomic bombs (detection and reaction period may be several years and requires a very wide span of different event sources).

What kinds of events would indicate these kinds of situations? Stop reading for a moment and have a guess! Usually, it will be a pattern of several events, not just one event. Also, detection systems for these kinds of problems will monitor for many different event patterns at the same time.

Much the same event processing technology would be used to detect patterns of events through monitoring the event sources in all of these examples. It would be part of the operating processes being employed in these situations. Here are some possible event pattern detection scenarios:

1. Baggage in an airline baggage handling system is tagged with RFID tags at the ticket counter. Scanners are placed at various positions on the pathways that the baggage takes from the ticket counters to the flights. There are scanners on conveyors and at places where baggage is loaded on and off moving equipment. A scanner emits a tag reading (an event) as it "sees" each tag. This reading contains the scanner's position together with the information on a tag. All the readings are processed by a monitoring system. A pattern of readings such as:

   20 tags read at scanner A in last 10 seconds **and**
   no readings from scanner B in last 10 seconds

   might indicate a problem in the handling system, such as baggage piling up somewhere between A and B.

Different event patterns would be used to detect other kinds of problems, such as some bags being directed to the wrong flights. For example:

Bag destined for Los Angeles loaded on Flight F **and**
Flight F destined non-stop for New York **and**
Los Angeles /= New York
raise ERROR

Similar event pattern monitoring is used to detect problems in many other kinds of processing lines and factory assembly lines.

2. A national public health electronic reporting system receives events from several sources, including electronic reports from hospital admissions and emergency room visits, physician's offices, local health agencies, and sales of prescription items at pharmacies. One target for automated monitoring of this flow of report events would be to detect possible disease outbreaks. Human-in-the-loop epidemic monitoring systems have sometimes proved too slow in the past.[13] Monitoring would use levels of event patterns, so that the system would not be slowed by a monolithic number of patterns to be tested all at once. The first-level event patterns might detect increasing numbers of reports for the same disease on a week-by-week basis for a specified number of weeks. For example:

Number reports of A last week < number reports of A next week
during past four successive weeks

A match of this pattern might then trigger a second level of monitoring. The event data would be run through other patterns that match location data in the reports to detect concentrations of each reported disease in a geographical area or population concentration. For example:

Number reports of A in area X last week > 5 * 25 year weekly average

If the actual numbers are deemed high enough, the pattern matching system would trigger both human interaction and third-level actions, such as feeding the report data to predictive models. Models of propagation for the disease may predict an area outbreak or a

[13]The Canadian GPHIN system is credited with uncovering SARS in 1998 when the Chinese attempted to hide it. www.phac-aspc.gc.ca/media/nr-rp/2004/2004_gphin-rmispbk-eng.php#4

potential for a national outbreak, indicating the need to trigger public health controls.

3. A retailer with a national network of retail outlets and a supply chain system installs event monitors on a variety of different event sources. Interestingly, there may be several different event monitoring systems within a large company, each run by different departments. Separation of event monitoring systems can be a problem, because it is important for the separate systems to cooperate. One system might monitor the supply chain events (e.g., inventory levels and sales figures from the outlets, customer orders, and the company's orders to its suppliers and their responses). For example:

   if supply orders for item X < outstanding customer orders
   then send warning to supply chain manager

   Another system might track the events in its distribution system such the status of its warehouses, distribution schedules and real-time data radioed from GPS systems on its delivery trucks. For example:

   truck A is 30 mins late at scheduled delivery point B

   This second system might also monitor outside event sources, such as national and local weather reports, traffic reports, driver availability, and other sources.

   Situations can happen in which the two event monitoring systems must cooperate by combining their processing results. Suppose the first system receives low inventory level warnings from some outlets, and in the same time window, the second system receives events such as truck breakdowns and weather reports indicating delays in the delivery system. If these events all arrive within a short time window, say eight hours, then cooperation between the systems could trigger various actions, such as advisory warnings to be issued to some outlets, and the rerouting of some supplies.

   if low inventory warning at time T **then**
   if delivery delay warning received within T + 1 day
   then request reroute supplies

4. Message traffic on a company's internal network to an unknown outside Internet address may indicate the presence of spyware or a compromised database and the theft of confidential information. Patterns of events may indicate that spying is in progress or that stolen data are being transferred outside. Such patterns of events can happen

very quickly after data have already been stolen and stored locally. However, the theft activity can go on slowly for weeks "under the radar" if the company does not monitor its IT systems adequately. The company security system could monitor for patterns of attempted data transfers and absence of recent transfer authorizations like this:

> **if** data transfer attempted outside network **and**
> **no** outside transfer authorized **within** last 15 mins
> **then** send warning to security

Notice this example contains two patterns and the second pattern is the absence of an *authorize* event. We'll describe how the absence of events can be detected later.

5. Credit card companies have libraries of event patterns indicating possible use of stolen credit card information. They are usually not very sophisticated patterns, depending upon geographic locations, known centers of card fraud, and very simple patterns of fraudulent use.

   For example, if the card holder lives in the United States and her card suddenly starts being used from locations in Eastern Europe, the card company's monitoring system will create an alert. Usage patterns that involve using a card in different geographic locations within a short period of time will trigger alerts. Any card use outside the holder's residential area may trigger an alert, particularly for foreign travel. Some card company alert patterns trigger high numbers of false warnings and consequently produce lots of irate customers. For example, an immediate card block when a card is used in a jewelry or electronics store or an art gallery without prior notice. Event patterns used in credit card monitoring are sometimes totally simplistic, because the card monitoring systems are built without any general event pattern-matching capability and inadequate memory of recent prior use (state knowledge). They simply test large numbers of use cases.
6. Homeland Security and other national intelligence organizations use a lot of event pattern detection—and they could use a lot more, particularly in coordinating detection across different agencies.

   Security is an area where the critical indicator events come from many different sources and may be spread across separate and competing government departments. The detection processes for security event patterns can be long lived, taking years to match completely. It is a problem that is compounded by false reports and erroneous events. Event pattern detection demands collaborative efforts among organizations that do not normally collaborate very well. Automation of event pattern detection is an obvious direction toward a solution

to this long-term global-scale monitoring problem. An example of a possible event pattern in this area would be:

if Classification of person A **is** terrorist threat **and**
report R received of A traveled to location B **and**
report R passes credibility test **and**
Flight F from B to USA **is** in progress **then**
alert homeland security **about** A onboard F

Each of the four events triggering an alert event in this hypothetical example would result from the aggregation of other, simpler events.

The time taken for the information content of several events to "add up" to an actionable situation varies greatly with the area of activity. But once the information is deemed credible, the enterprise must react—no matter whether it is a small business, a global enterprise, or a government. The ability to immediately detect and recognize patterns of events that have implications for decision making within the enterprise, *as and when they happen*, can sometimes be a matter of survival. And certainly it is always a competitive advantage. Every enterprise today should possess these kinds of event-processing capabilities.

Principle 5: Constant Vigilance

Every enterprise should be capable of monitoring 24/7 those event media, both internal and external, that carry information likely to indicate the need for changes in its business processes, strategies, or organization.

The Human Element and Other Sources of Errors

No discussion of event processing systems can be complete without a word of warning about the human element. Human errors, including very often the error of omission or the error of lack of coordination between actors, are some of the strongest arguments in favor of adopting automated event monitoring technology. A case in point is the previously mentioned mishandling of information in the Christmas 2009 airline bomber case.

Automated systems do go wrong. But when one looks at the reasons why, they often turn out to be due to error by their human users. For example, incorrect event patterns—patterns that do not contribute to detecting the intended situation—can be input to an automated monitoring system.

Another common source of errors is due to noise in the event sources themselves (e.g., RFID readers can give blurred and incorrect results). The information carried by events can be corrupted, thus misleading the event monitoring. Networks that transport our information deal with this by building confirmation processes into their communication protocols.

Guards against errors will be built into event monitoring systems. Error checking should be part of event processing technology—the ability to go backward and find the history of high-level events. Whenever a result of event monitoring and processing comes to a management decision level and involves financial or other resource expenditure, there must be an automated process of double checking.

Processes and methods of reducing errors in automated event processors and in their use must be adopted. Error detection has to be done in *right now* time.

There are many approaches to this issue. Any event-processing system will have the usual event logs and will support their use to retrace the history of processing that led to a questionable decision and also to do retrospective analysis after the fact. There will be various kinds of checking that can be turned on or off in the event processors. In normal use, these checking systems will probably be turned off for speed and efficiency. But an important event processing result will be rechecked (e.g., for errors in event pattern matching).

Extract What You Want to Know

The third reason to use event processing, *extracting what you want to know from the cloud of events available to you*, is all about how to cope with the cloud of too much information. It is related to the other problems we have discussed and can be considered as a necessary part of solving them.

Event feeds come from everywhere, and many of them are unstructured and unfocused. They simply carry events. Some of the events are unreliable, a lot of the events are unrelated to what the enterprise needs to know, and other events are relevant or even vital. It is an interesting area for development of new event processing techniques, somewhat akin

to news service analysis or sentiment analysis in financial news reports. It presents a challenge both to commercial enterprises and to governments. Some examples follow:

1. An international package distribution company operates fleets of airplanes and trucks and a network of distribution centers. Its business processes are intensely event driven and clocked. Timing is all-important to remaining on schedule and honoring service level agreements (e.g., overnight delivery). Its operations receive events from many sources, including its own fleets and distribution centers. Sources outside of the company are also tracked. These include traffic and travel advisories, labor disputes, and weather reports, much of which may turn out to be irrelevant. But all of those event sources must be monitored to extract any information that could indicate an impact on the company's operations. The potential impact on its schedules must be predicted and the affected operations replanned. The company has also to evaluate every situation that it considers for false alarms, because they are costly.
2. With more than 2 billion people traveling by air every year, an outbreak or epidemic in one part of the world is only a few hours away from becoming an imminent threat elsewhere. Worldwide epidemic warning systems must be improved so that the time lag between detecting the incubation of a possible epidemic anywhere in the world and alerting public health organizations is much shorter—on the order of days. Present reaction times in government-sponsored medical networks are sometimes on the order of weeks or months. The systems include human experts in making analysis and decisions. Also, these systems are hierarchical and involve time-delaying decisions at various levels, from local to regional to national. We must build event-monitoring warning systems that can detect epidemic threats while making minimal use of humans in the loop.

   For example, new systems are being developed based upon SMS cell phone reports from local field agents in rural areas of Southeast Asia. These systems also leverage popular social networking traffic. Rumors and hearsay have proved to be the earliest indicators of emerging epidemics, far earlier than medical reports in fact! websites such as Twitter are being monitored for indications of emerging diseases. Techniques include detecting increased use of disease and illness keywords in tweets. Lexicons of keywords are being developed and refined.

   The reason to use social interaction websites is to detect possible emerging epidemics as early as possible. Cell phone use has exploded in areas where some of the most infectious diseases incubate (SARS,

swine flu, West Nile valley fever, malaria, etc.). Many rural areas in Asia and Africa have little recourse to medical networks, but they have a rapidly expanding use of cheap cell phones.

But there are problems to be solved. The cell phone traffic is in multiple languages and will obviously include data of varying degrees of accuracy and focus. The data must be filtered, translated, and analyzed for relevance to a disease outbreak, geographic location, and frequency. Voice-to-text transformation technology would add greatly to usefulness of the earliest sources, especially in third-world countries. Automated methods must be used in all of the early processing phases to allow high-traffic volumes. The resulting events must be correlated and summarized for the decision makers.[14]

And above all else, false alarms of an epidemic outbreak must be avoided at all costs, because they can result in disastrous global economic consequences.

3. Homeland security is a paradigm example in which information needs to be extracted from a large variety of sources and aggregated into actionable intelligence. Its operations receive event feeds from direct sources, such as INS passport and visa entry point data, as well as from U.S. government departments, foreign governments, airlines, and so on. There are also many indirect sources. At the other end of the scale is intelligence gleaned by monitoring the daily activities of sets of individuals already resident in the country. Phone traffic may be taken into account. The sources vary in their reliability.

   The scale of the event input is one aspect of this detection problem. Another is time. Any single investigation may involve monitoring event inputs over long time periods. And intelligence has to be recognized from fragments of data that may appear at first to be unrelated or may be widely separated by both source and time. False alerts are expensive, but failures may be catastrophic.

As we previously mentioned, reaction times may be very different in each situation. For government agencies, both local and national, the situation may require alerting emergency services to meet a natural disaster within 24 hours, or averting a national medical crisis within a few weeks of the first indicators, or defending against a cyber attack on the power grid within seconds.

Commercial businesses must also uncover and evaluate the information in the event cloud in a timely manner. This will often have a direct

[14]The legality of this kind of monitoring is beyond the scope of this book but might involve exceptions for both governmental agencies and NGAs.

influence on the company's bottom line, maybe not immediately, but a few months or a year later. Again, reaction times from receipt of information to action may vary widely. Financial trading requires decisions in milliseconds. On the other hand, patterns of business events can take weeks or months to add up to information that forces adjustments in various managerial issues, such as marketing and sales strategies, or changes in the supply chain organization, inventory order schedules, stock investments, and so on.

Similar situations to the ones illustrated in these examples are happening all the time. Event pattern detection will play a key role in detecting actionable situations in large volumes of event inputs.

There is another factor to consider. The commercial offerings in this technology are improving rapidly. There's a continuous war of event processing technology going on. Enterprises that already have event processing in their business systems should keep their event processing capabilities up to date. And for those who don't—now's the time to start! The technology officers in the company should know what is going on in real-time event processing. This is not just about databases any more.

Principle 6: Keep Your Event Processing Technology Current

Event processing technology is evolving and improving. Enterprises must keep their event processing technology up to date and capable of coping with the changing problems they face.

This principle is about operating in the event driven world. There is a gap between the latest theories of event processing, such as Complex Event Processing (CEP), and the existing commercial event processing products in the marketplace. Commercial CEP technology is not static; it is evolving and improving all the time. Improvements are driven by competitive pressures and by financial incentives to solve new event processing problems.

Getting Started

We've described three of the more advanced reasons to add event processing technology to your business operations—*knowing*, *detecting*, and *extracting* various aspects of business intelligence. But there are many day-to-day

reasons to add event processing too—just to make your operations run smoother and faster, for example.

You will find a lot more on event processing technology in the later chapters. Also, Chapter 5 is a broad survey of the different business applications of event processing and the markets for it—essentially, this chapter answers our fourth question: *What kinds of enterprises have bought into event processing technology so far?*

Now let's suppose you decide to experiment and introduce event monitoring and processing into your company. What technology do you need? And what about the cost? Event monitoring is going to require investing in a lot of expensive software, isn't it?

Not necessarily! A lot of event monitoring can be done with quite simple software that'll run on the usual computers in your business. In fact, you might well have the software you need already deployed for other applications within your company—you're just not using its event monitoring capabilities. Or the new software you need might be easy to add on to what you've already got. A database with a graphical interface would do to get started. But you have to set it up in a way that doesn't lose too much of the *right now* aspect of your event processing.

When the company gets serious about applying event monitoring to its operations, some technical knowledge on the part of some personnel in the company will be needed.

What do they have to know? The basic concepts in event processing and complex event processing. A couple of technically minded personnel will get you started. In these times of personal digital assistants (PDAs) and smartphones, event processing is almost a natural language.

The concepts in complex event processing play a major role in building applications of event processing. Sooner or later, one gets into the "build it or buy it" decision. Some CEP technology out there today is freeware. But you have to fit that freeware into your IT systems to use it effectively. That's where the two or three technical people for a few weeks are needed.

Eventually, you may come to think it's worth investing in commercial CEP products and the vendor's consulting on their use. You can do what we're talking about in this book with an investment that will certainly yield a positive return—if you understand just a little bit about CEP. We'll come back to this topic when we've covered some of the basic CEP concepts. So, read on!