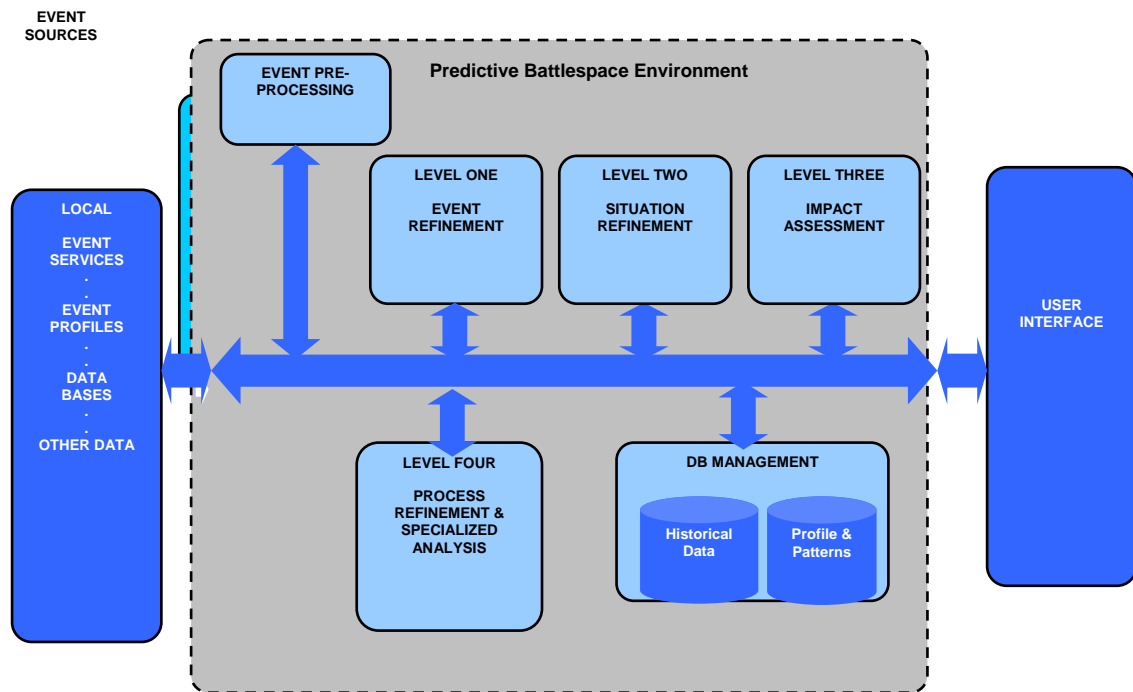


# The Predictive Battlespace

## A Strategic Thought Paper

Don Adams, Chief Technology Officer, Worldwide Public Sector





**Abstract** - The art and science of multi-sensor data fusion has emerged as the underlying foundation for Predictive Business<sup>®</sup>, including applications in telecommunications, finance, transportation and defense. All of these have common threads requiring complex inference processing solutions and require the management of real-time events from distributed sensors, agents and other processing components, including historical data-at-rest repositories. Distributed coordination-based architectures such as TIBCO<sup>®</sup> Enterprise Messaging Service (EMS) or TIBCO<sup>®</sup> Rendezvous (RV) provide the underlying communications infrastructure that enables complex event and high performance rule-based processing services. In this paper we discuss Predictive Business in the context of a pressing defense need to remove an enemy's decision in a Predictive Battlespace with a focus on the distributed processing architecture. The focus is on the JDL data fusion model and TIBCO<sup>®</sup> BusinessEvents in the context of complex event processing applied to the future state Predictive Battlespace.

**KEYWORDS: COMPLEX EVENT PROCESSING, EVENT-DRIVEN ARCHITECTURE, EVENT STREAM PROCESSING, JDL DATA FUSION MODEL, MULTISENSOR DATA FUSION, PREDICTIVE BUSINESS, RULES-ENGINE, RULES-BASED SYSTEM, PREDICTIVE BATTLESPACE**

## Introduction

In his latest book [1], *The Power to Predict*, Vivek Ranadivé, founder and CEO of TIBCO Software, discusses how *Predictive Business*<sup>®1</sup> is enabled by the fusion of historical knowledge with real-time information. The combination of data-at-rest with real-time data-in-motion is at the heart of understanding the momentum behind TIBCO's event-driven architecture (EDA). We elaborate on the concepts of *Predictive Business* in the context of an overarching processing architecture that helps TIBCO customers and employees realize solutions in context to different business domains. This paper focuses on the use of Predictive Business as part of the evolution from Network Centric Warfare to a Predictive Battlespace while grounding the discussion to the application in counter terrorism.<sup>2</sup>

## Network Centric Warfare and Network Enabled Capabilities:

Over the past decade Network Centric Warfare and Network Enabled Capabilities have largely proven themselves as they improved both the efficiency and effectiveness of combat operations. Both have allowed more effective projection of force while minimizing support and logistics expenses. This is commonly referred to as the “Teeth-to-Tail” ratio of operations. So what is next in this largely evolutionary process?



For a while we are likely to see more of the same as commanders like proven approaches that make their forces more effective with lesser costs and saving of lives. The C4ISR community has learned a lot from the study and advocacy of NEC and NCW precepts and theories. These have been studied to death. One of the things to evolve from NEC operations is an awe-inspiring array of sensors at every echelon. Every forward deployed

<sup>1</sup> TIBCO® and Predictive Business® are registered trademarks of TIBCO Software Inc.

<sup>2</sup> Please note that the same high-level processing architecture applies to solutions in many other Predictive Business® related areas, including opportunistic trading, network and telecommunications management, intrusion detection, real-time diagnosis, and more.



asset seems to have some secondary sensory mission. Unfortunately, these sensors are largely localized to a unit, or at best within one echelon. This limitation means that the assets in the Battlespace remain largely reactive to immediate local sensor triggers. The availability of sensors thus has only limited impact on their opponents OODA Loop.

#### **JOINT DIRECTORS OF LABORATORIES DATA FUSION MODEL**

The Joint Directors of Laboratories (JDL) data fusion model has been demonstrated to be directly applicable to detection theory where patterns and signatures discovered by abductive and inductive reasoning processing (for example data mining) is “fused” with real-time events. The JDL processing model [2], has survived the “test of time” as the dominant functional data fusion model for decades. The vast majority of the most complex real-time event processing architectures are based on the JDL model. Vivek Ranadivé indirectly refers to this model<sup>3</sup> when he discusses how real-time operational visibility, in context with knowledge from historical data, is the foundation for Predictive Business.

It is of interest to note how the heart of the JDL data fusion model is a communications infrastructure that looks remarkably like TIBCO’s marquee “information bus.” It may also be of interest to note how the JDL model directly corresponds to the concept of a service-oriented architecture

There are three main “meta” architectural components of the JDL model. Perhaps the most important of these components are the events. Events can be local and external, and originate from myriad sources in many formats. There is no requirement that events must use the same syntax or format, as we will understand in the context of *Level 0 Event Preprocessing*. There is also the core complex event, and/or the so called event stream, processing architecture, which the JDL model depicts as the “data fusion domain.” We will use the terms, *data fusion domain*, *data fusion*, *multi-sensor data fusion*, *complex event processing* and *event stream processing* somewhat interchangeably, as the concepts are independent of vendor implementation and have similar technical objective and outcomes. Finally, there is the user interface, providing operational visibility into every technical and business process depicted in the model.

The business objectives of organizations may differ in context, but the overarching technical goal to enable the business objectives are the same:

*Correlate information gained from abductive and inductive reasoning processes with real-time events to infer current situational knowledge and predict both opportunities for, and threats to, the enterprise to maximize assets and minimize liabilities.*

---

<sup>3</sup> Refer to pages 46 and 47 of [The Power to Predict](#) [1].




## Predictive Battlespace

This technical objective is at the heart of TIBCO vision for *Predictive Business*. It is also at the heart of concepts and conceptualizations of the Predictive Battlespace. To expand the basic precepts of Predictive Business into a Predictive Battlespace Environment, we need to add other common elements from both C4ISR and commercial environments. These include sensors of every description, data warehouses, Business Process Management and workflow, specialized analytical processors ( imagery, image, facial recognition, license plate readers, thermal imaging, signal processing, and classic multi-intelligence sources), inference engines supporting rule-sets , state-machine and temporal elements, integration with every aspect of legacy C4ISR, and DII administrative systems, and distributed secure kiosks.

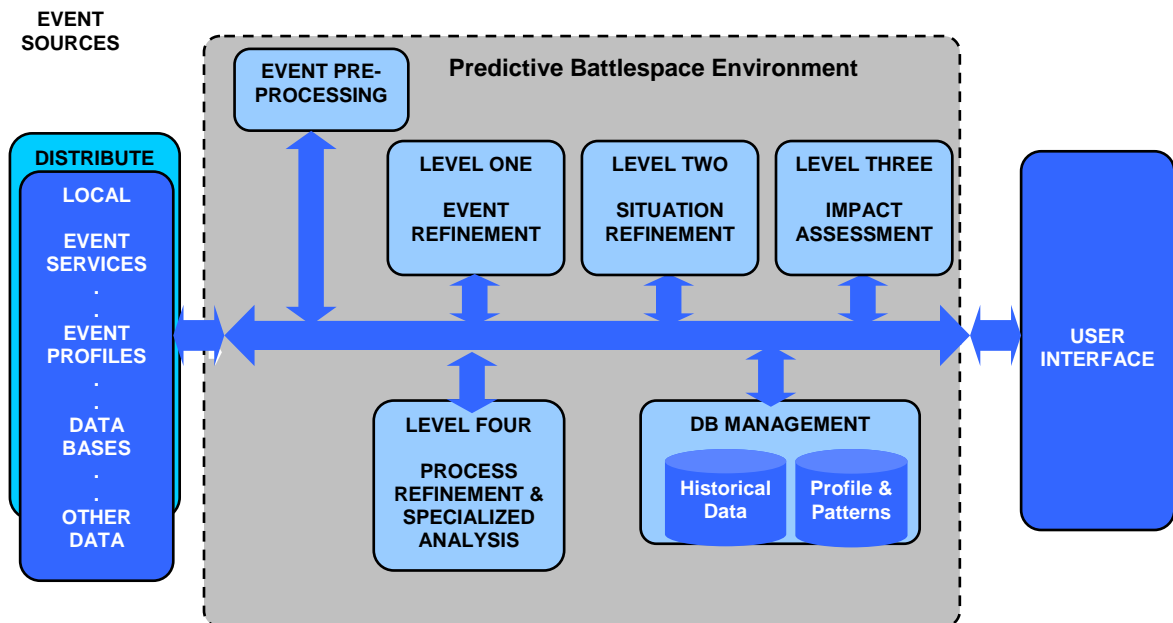
This is a totally fictional scenario for a notional Predictive Battlespace Environment (PBE) that current technology (page 11 has an architectural drawing of a notional PBE) could provide as an augmentation to the MoD DII systems. This augmentation into the C4ISR space would build on recently proven success of Network Enabled Capabilities (NEC) in NATO and Network Centric Warfare from the US to turn a responsive combat capability into a predictive combat capability. The PBE builds on commercial success in large-scale employments of Complex Event Processing, initially described by Stanford University professor, Dr. David Luckham in his book, **The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems** and the corollary Predictive Business solutions from TIBCO Software

### Scenario begins:

'M' slips out of what serves as a safe house in this 'little slice of paradise' no-man's land between two conflicted countries and jumps into an ageing Mercedes in the middle of a line of six vehicles. He is smiling and praying because today's meeting will be a major step in expanding his already significant influence among his growing ultra-violent radical following. Little does he imagine that he may have a bad day. Unseen, unheard and otherwise undetectable, a modified AeroVironment Global Observer, "eye-in-the sky" caught several images of him in the gap between safe house and car, as a routine part of its one-week overhead imaging mission. . The high altitude imagery sensors, drawn to the heat signatures of the collection of vehicles matching a typical pattern for a local terrorist convoy, caught everything in its digital sensors and relayed imagery to the Predictive Battlespace Environment.



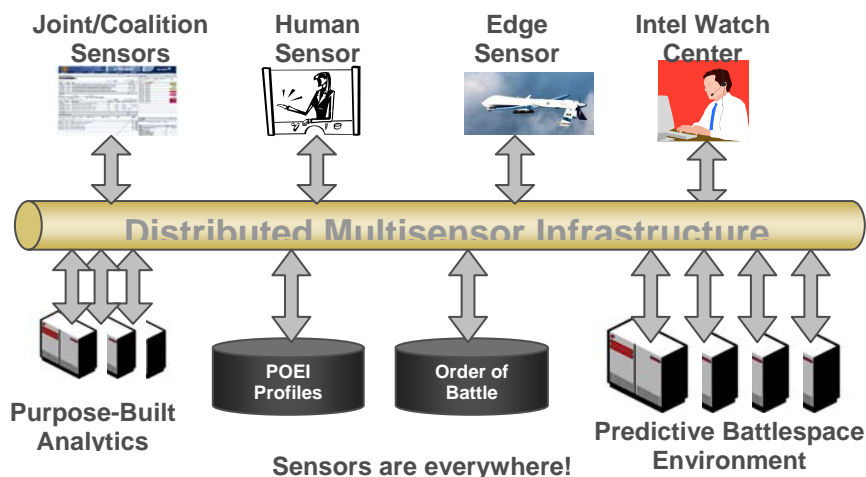
Over a hazardous and uncontrolled, unfriendly, frontier an advanced Global Observer intelligence asset passes unobserved at an altitude of nearly 60, 000 feet, silently recording video, still and thermal imagery,. The sensors would detect matching parameters of motion, image definition and thermal patterns indicating potential terrorism or other threat. The Global Observer, tasked as part of the sensor input for the UK MoD DII Predictive Battlespace Environment, blindly records anything in motion or with heat signatures matching its downloaded rules of action (ROA). Its recordings are processed in real-time by imagery, infrared and Video IQ systems at a mid-tier in the PBE hierarchy.



The mid-tier analysis server escalates an imagery tag to a higher level analysis and simultaneously to an Intelligence Watch officer tracking Persons of Extreme Interest (POEI). The higher level analysis software adds clarity to the visible analysis of the Watch Officer that this is a convoy of vehicles in a configuration commonly used by terrorists in this region potentially moving ‘M’ a POEI. Triggered by the Watch Officer, the PBE re-tasks the intelligence asset to monitor, track and record the convoy when in its targeted area. In addition to normal automated image analysis, the feed is also delivered to the Watch officers and others they may designate.

The convoy stops in front of a restaurant in a mixed commercial and residential area and a group of men step out, chats briefly, and proceed into the restaurant. While they are milling, imagery is immediately processed by the Video Facial IQ engines running under control of the PBE at level 4 (since there was a previous alarm active on the convoy). The watch officer simultaneously zooms in on a face and is fairly certain that the individual is ‘M’, a Person of Extreme Interest. He sends an alarm message out to the duty officer, Flt. Lt. Heather Jenkins who is away from the watch.

In response to the insistent and unique alarm from her secure handheld message terminal, Flt. Lt. Heather Jenkins bolts out of the Daily Summary briefing she is attending. She runs to the nearest DII kiosk and authenticates with the three factors required for the role she is entering. She is immediately presented with all of the critical current task information including a cached version of the video delivered to all kiosks she was likely to use based on her appointment calendar. This entirely AJAX rendered Watch Console worktop is secured in case she has to exit immediately leaving no content or other sensitive information in the DII Kiosk.



While Flt. Lt. Jenkins and the watch officer were slowly (by computer time) performing these actions, the PBE researches the Order of Battle (OB) and all current FRAGOs, assembles and analyzes all current information on the POEI, known associates and predicts a pattern of their behavior in this scenario and publishes a Most Reasonable Course (MRC) recommendation.

Flt. Lt. Jenkins looks at the PBE rule-sets that triggered the alarm and the recommendation of the watch officer, the MRC provided by the PBE and issues an immediate concurrence and the pre-FRAGO provided by the PBE for approval by the Joint Task Force Commander (JTFC). All video and other dense content is pre-queued to all positions where the JTFC may require access to them.

Called out of a staff meeting the JTFC accesses his command console in much the same way as Flt. Lt. Jenkins accessed the Kiosk and is instantly presented the key aspects necessary for his role in the decision to re-task an armed UAV on station within range of the POEI. He reviews the rule-set, imagery and the MRC with the senior Agile Mission Team members, and authorizes the issue of the necessary FRAGO to attack and destroy the POEI and adjacent entities. In an earlier time, this process would have taken 30 minutes to an hour or more. When augmented by the PBE, the kiosks and the DII infrastructure, all actions were completed in less than seven minutes.




An armed Aurora Excalibur controller (already alerted by the PBE as part of the pre-frago processing) rapidly dials-in the coordinates of the restaurant so that her baby can fly autonomously as she continues to study all aspects of the location, imagery of 'M'

arriving and the profile information forwarded by PBE under security release authorization of the JTFC for her target.

Flying just sub-sonic, the Hellfire armed Excalibur will reach the target coordinates in nine minutes. History and informants report 'M' is a person of habit, who normally exits a building only after three cohorts have done so without drawing fire. He always scans the surroundings and the sky as he stands in the doorway. Once confident, he walks out quickly with a group around him approaching the car in front of his in line before stopping for ritual pleasantries. Mingling and talking, he will then slip largely unnoticed into his armored Mercedes vehicle; the cohort will disperse into theirs and depart immediately. Helpful for a sniper or special operations team attack, it will not help him today. The Excalibur controller has weapons release authority for 'M' and surrounds.

Two forward deployed clandestine special operations teams with ultra-light, silent, Chang Industries SkySeer, hand-held, digital video craft, have also been tasked and as the Excalibur approaches, they are launched, on an up to one hour mission to document the attack. The electric powered SkySeers relay real-time imagery to their controller who relays it into the PBE via satellite. The telephoto lens on the side-facing camera system of the two SkySeers allows the JTFC a clear image of the target in case there has been a misidentification in the process



In the Watch Center, the JTFC and intelligence staff watches the mission as it unfolds before them in digital imagery from the SkySeer and the Excalibur cameras, overlaid with all multi-intelligence sources and processed information. The JTFC can at any point where he is not certain of the target, push the button in front of him and recall the entire mission.

The Excalibur controller using all intelligence assets as basis decides to wait for the gaggle to reach the cars and begin milling before releasing the Hellfire. She approaches



from behind the restaurant keeping out of ear-shot and line of sight while 'M' exits the building. Continuing to glance at her PBE Dashboard, in case of a recall or 'wave-off, she prepares every aspect of her attack, knowing there is only one chance for success. Working from the real-time SkySeer, gods-eye-view displays, she waits for the optimal second to dive her baby towards the congregation and release the Hellfire Missiles directly at 'M'. 'M' did not have a good day, his last in this lifetime. The PBE creates a perfect record of every mission parameter, imagery frame, multi-sensor fusion outcomes and the human factors in the decision tree for posterity.

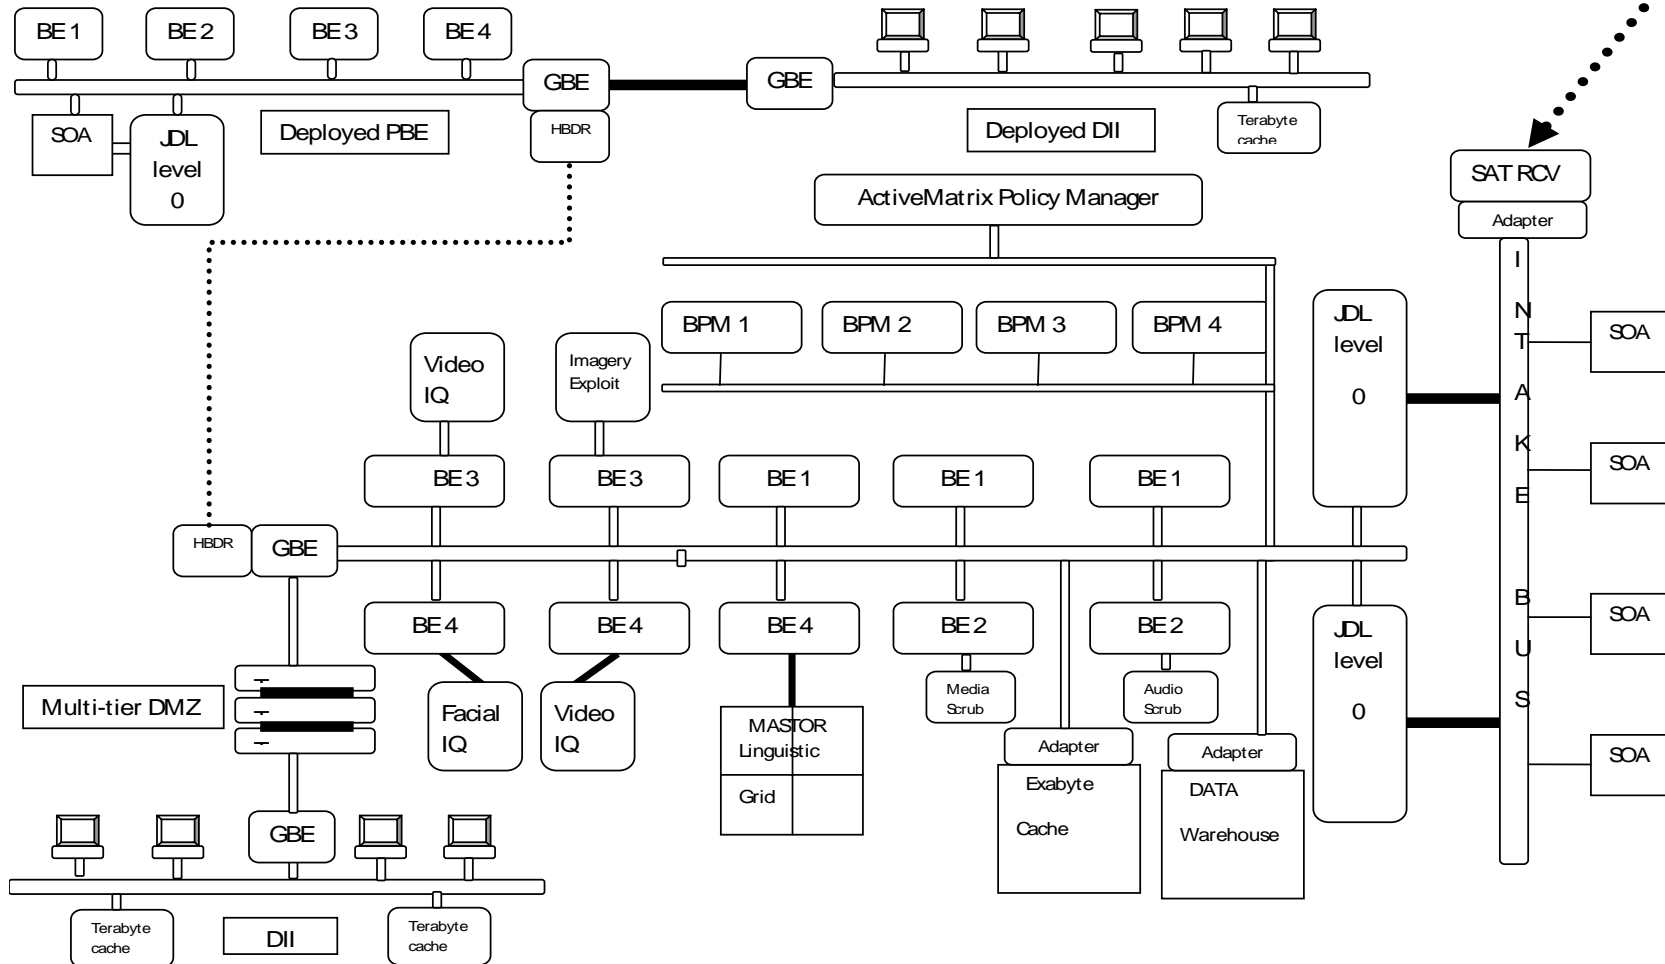
With the fusion of information from the multi-sensor event stream and the fusion of multi-source intelligence on 'M' and his cohorts, a deployed PBE employed by a group of highly skilled professionals and decision makers a threat to all mankind has been removed efficiently with little risk to friendly forces.

In a mission debrief, the PBE delivers pre-packaged and printed documentation of all information, rule-set triggers, processing logic and human decisions. These packages are studied and discussed by experts. From their analysis, the rule-sets, state-machine parameters, and all temporal aspects can be adjusted. Infrastructure can be changed to include location and orientation of DII Kiosks, training aspects related to accuracy of online calendars of key personnel used for pre-positioning of cached imagery and fusion data from all-source data warehouses. These newly changed aspects can all be tried, tested and exercised without impacting ongoing PBE support to the mission.

## End of Scenario:

As shown above the addition of the Predictive Battlespace Environment deployed as an extension and enhancement to the MoD DII helps the transition from responsive to predictive.

The following sections reflect on the TIBCO Software Inc. BusinessEvents product and its Predictive Business architecture and its direct applicability as a core element in the Predictive Battlespace.





## **Predictive Business Architecture**

The business outcomes of *Predictive Business* are realized when an enterprise can leverage knowledge and information gleaned from historical data-at-rest as patterns that are applied to real-time events and situations. Figure 2 depicts a high-level communications architectural view of how business events are published to subscriber event-services, such as rules or inference engines. In turn, the output of these rules or inference engines can also be events, complex events, situations, threats and/or opportunities which are published to subscribers.

The communications model, which some industry analysts refer to as an *enterprise services bus (ESB)*, provides an organization the capability to take advantage of business intelligence, real-time business events and other information sources, such as advanced sensors and rule-based processing. Of interest to note is that the communications infrastructure is but one technical requirement for *Predictive Business*. Another required technical capability for *Predictive Business* is a processing model. This is the rationale for introducing the JDL processing model into a discussion of both the Predictive Battlespace and complex event processing.

The high level requirements for *Predictive Business* are fairly straight forward when viewed in context with an established inference processing architecture such as the JDL data fusion model. This model is summarized in the paragraphs below:

### **Level 0 - Event Preprocessing**

Event preprocessing is often referred to as data normalization and feature extraction. Heterogeneous events from sources across the extended value chain, both external and internal to an enterprise, exist in many data formats, accessible from local methods and application interfaces. Level 0 preprocessing is a generic term for normalizing data for upstream event processing. The terms *data normalization*, *data cleansing*, *event normalization*, and *object preprocessing* are terms that are often used interchangeably. In context, they refer to similar processes of preparing data for further “upstream” processing.

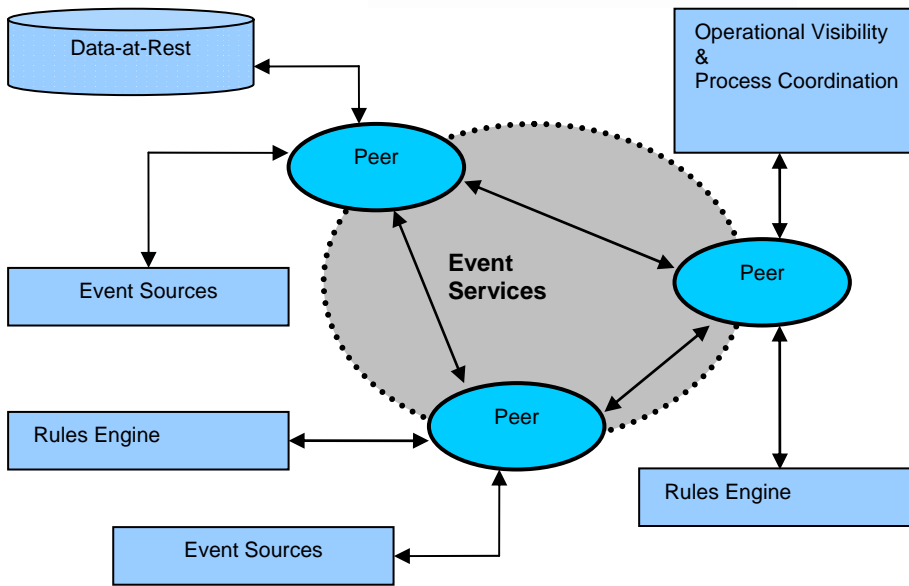


Fig. 2. Example Architecture for Predictive Business

TIBCO customers often use TIBCO Adapters and BusinessWorks to normalize data before presenting the data to the next processing stage within an enterprise service bus. In addition, a high performance rules-based engine such as TIBCO Business Events, discussed in more detail later, can also be used for data normalization and event preprocessing.

For example, a high performance network service that passively captures all inbound and outbound network traffic from a multitude of sensors and sources may normalize traffic based on temporal sessions prior to further processing. This is a real world example of Level 0 event preprocessing for the core elements of a notional Predictive Battlespace business scenario. As one might imagine, this can be a challenging problem when there are thousands of transactions per second in real-time. TIBCO has a proven track record in high performance distributed processing.

### Level 1 - Event Refinement

In reality, the event processing challenge does not end with event preprocessing. Data normalization and basic feature extraction is a housekeeping task – dirty work – but critical if the next stages are to be successful. After event normalization, other challenge looms on the processing horizon. Level 1 processing involves selecting events for inclusion in an event track or trace.<sup>4</sup>

<sup>4</sup> This selection process often involves the use of a real-time scorecard, discussed in more detail later.



In the context of relevant event detection, this is often the task of identifying possible *elements of opportunity or risk* by methods of association and correlation - classification by pattern matching in the normalized raw event stream. Generally, this process has three steps:

1. Hypothesis Generation;
2. Hypothesis Evaluation; and,
3. Hypothesis Selection.

*Event Hypothesis Generation*, in the context of our example of relevant terrorist event detection, tags groups of events which *may* represent an aspect of a possible terrorist activity from the remote sensor arrays on the network. The input is the normalized event stream and the output is a matrix or scorecard of possible events of interest.

*Event Hypothesis Evaluation*, the next step is to evaluate the event hypothesis generated to rank the events based on likelihood. For example, the output determines which events have a higher likelihood of representing terrorist related activity.

*Event Hypothesis Selection*, continuing in the context of relevant event detection, hypothesis selection attempts to associate a name or classification to the suspected sensor or source event.

This level, or stage, of event processing normally requires a high performance rules-based pattern matching algorithm. TIBCO BusinessEvents uses the RETE Algorithm, discussed in a bit more detail in the next section, regarded as one of the most efficient algorithms for optimizing mainstream rule-based systems [5].

## **Level 2 – Situation Refinement**

Situation refinement represents a higher level of inference than object or event refinement where estimation and prediction of event states happen based on statistical associations between events.<sup>5</sup> Level 2 processing is commonly referred to as *relation-based state estimation*. The state of the aggregated events is represented by a network of relations among individual events.

Another way to view this stage of processing is that we have previously taken normalized data from the raw data or event stream and extracted basic features, processing the event stream by matching with patterns obtained from historical knowledge. This historical knowledge was extracted from *a priori* data-at-rest and “fused” with the real-time event stream.<sup>6</sup> After an event is identified it can be correlated with other real-time events,

---

<sup>5</sup> In the prior inference stage (L1) the process performed high speed pattern matching with a highly efficient rules-based engine to detect objects of interest in real-time. The outcome of L1 provided both an estimation and prediction of low level entity states based on inferences from the actual raw data stream.

<sup>6</sup> There might be a bit of overlap and disagreement in emerging terms and concepts at this point in the discussion. Hence, for purposes of this paper, and given the liberty as the writer, I will refer to what “just happened” in Level 1 Event Refinement as “*event stream processing*” (ESP) and what “happens now” in Level 2 Situation Refinement as “*complex event processing*” (CEP) – but this is just a subjective opinion. In fact, a good argument could be made that ESP and CEP are the same thing. Objectively, there is not

models, patterns and data, illustrated in Fig. 3, to infer and predict more complex event and situations.




Fig.3. Discovering Situations from Detected Events

The correlation of events into more complex events is what is referred to as situation refinement in the JDL data fusion model. Here, we use the term “situation refinement” and “complex event processing” interchangeably because the output of this level of processing has the same goal, inferred situations, or complex events, from aggregated individual events.

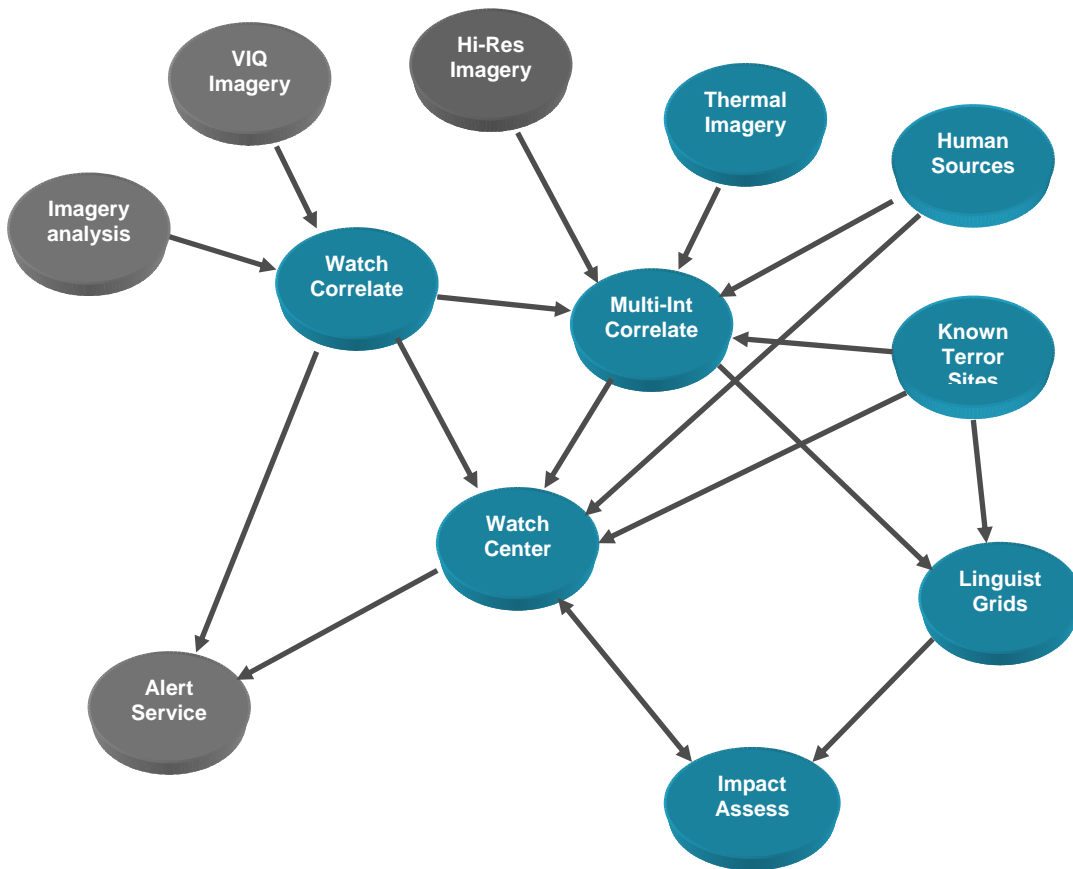
Fig. 4 represents how an organization might detect the event of interest from one or more sensors or sources. In this figure external event sources are represented as gray nodes and internal event sources are represented as blue nodes providing an illustrated example of an aggregate network of relations among individual events.

Each node in the graph is connected to other nodes by arrows which represent a cause-and-effect (causal) relationship. Nodes may be associated with conditional probabilities. As a graph, Fig. 4 represents the organization’s *a priori* knowledge of how they might detect a terrorism related event in the context of sensor rationalization and fusion, which is a situation inferred by the correlation of aggregated events.<sup>7</sup>

---

enough definition and vocabulary of either CEP or ESP to provide the required “absolute” context; and folks from different vendor camps are competing for position and influence in the marketplace. This is one reason I have chosen to ground the discussion in the JDL data fusion model, which is much better defined and less subject to marketing hyperbole.

<sup>7</sup> Detecting situations with the method represented in Fig. 3 was invented by Thomas Bayes in 1731. Bayesian probability is the name given to a family of related interpretations of probability, which have in common the application of probability to any kind of statement, not just those involving random variables. “Bayesian” has been used in this context since about 1950 [6]. There are numerous other mathematical techniques for detecting situations. Bayes’ theorem is the dominate technique in fraud detection and similar classes of detection problems including most modern email SPAM filters.



**Fig.4. Detecting Threats and Opportunities**

### **Level 3 – Impact Assessment**

After detection and situation refinement, organizations are generally interested in ascertaining or predicting the intent of those responsible for the current situational threat. Processing situations based on models to infer intent is referred to as Level 3 Impact Assessment. Level 3 is defined as the estimation and prediction of events on situations of planned, estimated, and/or predicted threat (or opportunistic) actions by the participants.

Therefore, Level 3 processing is normally implemented as a prediction based on higher level inferences based on Level 2 associations. At this stage of the model, we may estimate the impact of an assessed situation, which includes likelihood and/or cost/utility measures associated with potential outcomes of a player's planned actions. From this inference, loss projections and liabilities may be estimated.

For example, organizations are generally not only interested in detecting terrorism trigger events from the sensor arrays in their networks – they would also like to know, to the extent possible, the intent of the terrorist or his pattern of movement and the potential



damage to the organization, in our example a defense organization, if the terrorist is successful.

#### **Level 4 – Process Refinement**

Event process refinement is an element of resource and task management. Functionally, Level 4 is where adaptive data management and real-time resource and process control takes place.<sup>8</sup> The focus of Level 4, in contrast to Levels 0-3 in the JDL model, is overall planning and control, not detection, estimation and prediction. Level 4 process refinement is the process of assigning resources to tasks, performed either by people or automated processes. Level 4 is also where variables and models are updated, removed, added, tuned or otherwise refined.

#### **Database Management and Operational Visibility**

Supporting the overall model, are databases of features and patterns extracted from historical data and other supporting databases. These databases may also contain models created by domain experts. For example, there may be a database of known IP addresses of Internet sites of known and suspected terrorist organizations or a database of known terrorists, their habits, associates and false identities used to hide the true identity of the individual terrorist agents. It goes without saying that operational visibility at all levels of the inference process is important.<sup>9</sup>

#### **Rule-Based Systems & Business Events**

Rule-based systems (RBS), often referred to as *expert-systems*, are widely used to model the behaviour of domain experts. For this reason, RBS are used extensively in a wide variety of business applications such as customer relationship management, fault and medical diagnosis, mortgage evaluations, credit card authorization, fraud detection, and C4ISR environments. These systems use declarative programming to tackle problems involving control, diagnosis, intelligent behaviour, and problem solving by describing *what* the computer should do rather than the exact procedural instruction on *how to do it*. Rule-based systems contain rules from a specific domain and use these rules to derive solutions to a wide range of business problems.

---

<sup>8</sup> Readers may be aware that TIBCO® Hawk® provides a mature communications infrastructure for monitoring, command and control of distributed complex event processing components.

<sup>9</sup> In some interpretations of the JDL model, operational visibility is referred to as a fifth level.

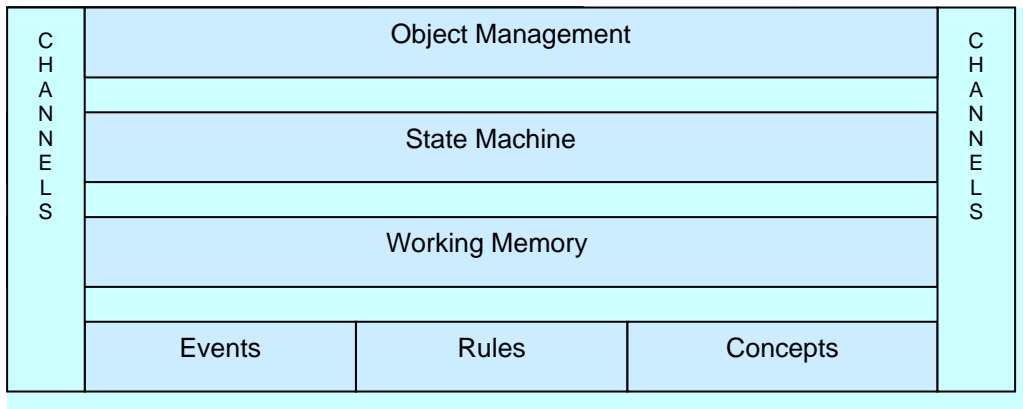


Fig. 5. Business Events High Level Architecture

*Rule-engines*<sup>10</sup> are processing engines which are designed to follow rules that are written in their language. These engines are used to execute complex decision-making logic on data and event streams to make decisions based on tens, hundreds or thousands of *facts*. Rule-engines accomplish this by decomposing large sets of rules into an efficient network of nodes which can process and react to *facts* far more efficiently than can be programmed using procedural programming. Therefore, rule-engines, if properly designed, scale well for numerous classes of problems, including complex event processing [7].

The RETE Algorithm<sup>11</sup> is a well-known efficient RBS implementation that also creates a network of nodes. Each node in the network represents one or more test conditions found on the left-hand side (LHS) of a rule set. At the bottom of the RETE network are nodes representing individual rules. When a set of events filter all the way down to the bottom of the network, the set has passed all of the tests on the LHS of a particular rule and this rule set becomes an *activation* and potentially generate one or more events as input to other processing or alerting solutions. The associated rule may have its right-hand side executed (fired) if the activation is not invalidated first by the removal of one or more events from its activation set [9].

TIBCO BusinessEvents (BE) was designed to provide organizations the capability to execute the same complex decision-making logic, without the programming complexity, and is based on the RETE Algorithm. Using a declarative programming model, business architects can define rules that will execute on both individual and combinations of

<sup>10</sup> This section is intended as only a brief review of rules engines, the RETE Algorithm, and a few key concepts related to TIBCO® BusinessEvents. Please refer to the references for a more detailed discussion or explanation.

<sup>11</sup> As mentioned earlier, the BusinessEvents rules engine is based on the RETE Algorithm. The RETE Algorithm is an efficient pattern matching algorithm for implementing rule-based systems designed by Dr. Charles L. Forgy in 1979. Subsequently, RETE became the basis for many expert systems, including JRules, OPS5, CLIPS, JESS, Drools and LISA.

events and facts in working memory. The Business Events high-level run-time is architecture represented in Fig.5.

The core rules-engine architecture of BusinessEvents is illustrated in Fig. 6 [8]. A TIBCO BusinessEvents rule has three components. The *declaration* component is used to specify which *Concepts*, *Score Cards* and *Events*<sup>12</sup> the rule requires (as well as naming attributes). The *condition* component represents facts that evaluate to a Boolean value.

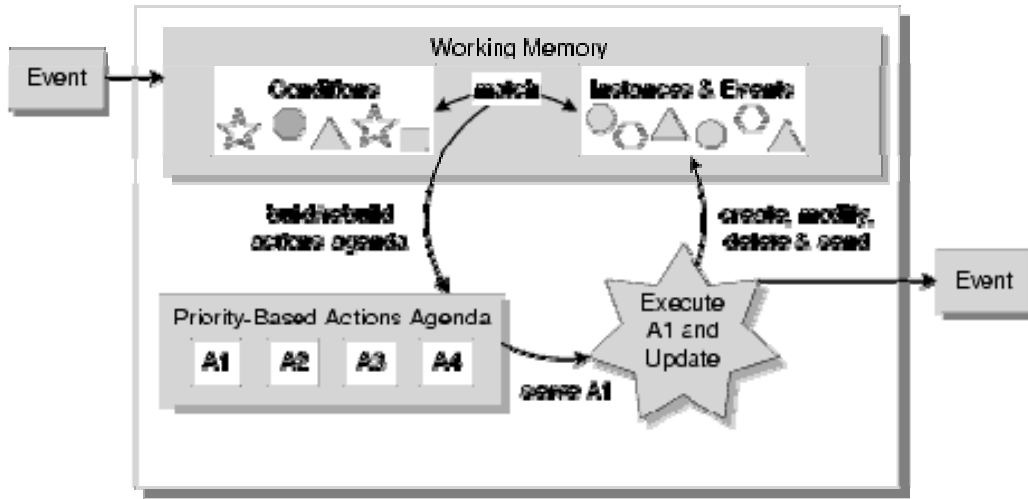


Fig. 6. Business Events Rules Engine

All conditional statements must be true for the rule’s action to be executed. The *action* component of a Business Events rule is a collection of statements to be executed when all the conditions are true.

A Business Events *Event* represents an instance of an *event definition*, an immutable activity that occurred at a single point in time. An *event definition* includes properties evaluated by the BE rules-engine, including TTL and comprehensive information related to the event. A *Concept* in BE is an object definition of a set of properties that represent the data fields of an application entity and can describe the relationship among entities. Likewise, an instance of an event is persistent and changeable; whereas an *event* expires and cannot be changed.

A *Score Card* is a BE resource that serves as a container for global variables and can be used to track information across the application. For example, in the JDL data fusion model, matrices are used for assignments at all levels of the inference processing model.

<sup>12</sup> TIBCO® defines an *event* as an immutable object representing a business activity that happened at a single point in time. An event includes information for evaluation by rules, meta-data that provides context, and a separate payload — a set of data relevant to the activity.

As illustrated in Fig. 7 multiple *Score Cards* can be used, conceptually, for this dynamic assignment at Level 0 through 4 of the JDL model [2].<sup>13</sup>

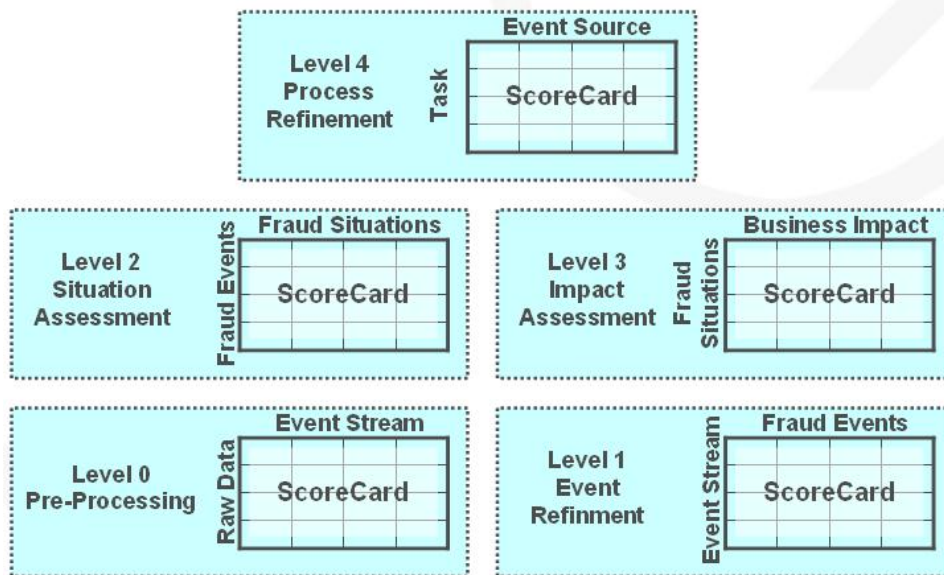


Fig. 7. Scorecards and JDL Processing

The multi-level relevant terrorist related event detection scorecard example of Fig. 7 illustrates how an event stream might be processed into “interesting events” by extracting events from the event stream. These events are ranked, or scored, according to a likelihood estimation of relevance to a pattern of known terrorist activities. These events, ranked according to likelihood, may be further processed to infer more complex events and situations. These complex events and situations are then evaluated for mission impact. The entire event stream should be reconfigurable in real-time to provide the capability to fine tune all detection and estimation processes.

In the interest of the reader’s time, this section only scratched the surface with regard to the run-time capabilities of Business Events, rules-based systems, rules-engines and the RETE algorithm.<sup>14</sup>

<sup>13</sup> Note that although some applications of the JDL model show scorecards used for Level 0 – 4 association and state estimation variables. *Score Cards*, from a BusinessEvents implementation perspective, may or may not be the most efficient method to extract features from an event stream. Scorecards are introduced here as both a generic concept (scorecard) form the JDL model used for explicit associations in performing state estimation and BE Business Events *Score Card* resource. In practice, BusinessEvents implementations of association and state estimation will more-than-likely use the BE *Concepts* resource.

<sup>14</sup> Inference and multi-sensor data fusion can be a difficult subject to grasp. You are not alone if you feel this way! That is one of the reasons I have grounded parts of the discussion in fraud detection - to hopefully provide an element of business context to the dialog. My impression is that terms like *complex event processing* and *event stream processing* are useful because they “sound better” (cool!) than the more traditional terms like *data fusion*.



## 4 Brief Discussion

TIBCO® Business Events defines complex events as an abstraction that results from patterns detected among aggregated events. The reader will readily note that the CEP [7, 8] definition of a complex event corresponds to the JDL Level 2 inference abstraction referred to earlier as a *situation*<sup>15</sup>. Fig. 7 also illustrates how building inference in a relevant terrorist event detection scenario may be represented by using the built-in Business Events *Score Card* resource.

Hopefully, this paper has illustrated that real-time sensor data, some might refer to as the “event cloud”, and across the extended value chain can be processed into event streams using rules-based processing. From event streams it is possible to extract features in real-time based on matching patterns to the event stream. Features of the raw data or event stream become “events (objects) of interest” which can be scored according to likelihood estimates constructed from high speed pattern matching algorithms with almost zero latency.

Events may be aggregated and correlated in run-time to infer complex events also referred to as “situations.” Complex events and situations discovered in run-time can be processed with patterns developed from historical data to predict future business events and situations. This brings us back full circle to *Predictive Business*, the theme of TIBCO’s CEO, Vivek Ranadivé, latest book, [The Power to Predict](#).

Real-time events, combined with patterns and features extracted from historical data and expert domain knowledge are the foundation for businesses to anticipate exceptional situations, estimate the impact on both the business and the customer, and take corrective actions before exceptional situations becomes problems. In other words:

*Predictive Business* leverages business assets to maximize opportunities and minimize future organizational liabilities.

Relevant terrorist event detection is only one of many examples of how the concepts of *Predictive Business* can help C4ISR organizations minimize threats and liabilities to their combatants, the enterprises infrastructures that serve them, and in turn, the citizens and politicians who place trust in their organizations.

## Acknowledgments

Dan Ziman, Puneet Arora, Saul Caganoff, Nick Leong, Heinz Schaffner , Suresh Subramani, of TIBCO Software Inc.  
Tim Bass, Silk Road

---

<sup>15</sup> Kindly refer to an earlier footnote on an editorial comment on the terms “complex event processing” and “event stream processing.”



## References

- [1] Ranadivé, V., "The Power to Predict," McGraw-Hill, NY, NY, 2006.
- [2] Hall, D. and Llinas, J. *editors*, Handbook of Multisensor Data Fusion, CRC Press, Boca Raton, Florida, 2001.
- [3] Bass, T., *Service-Oriented Horizontal Fusion in Distributed Coordination-Based Systems*, IEEE MILCOM 2004, Monterey, CA, 2 November 2004.
- [4] Bass, T., "*Intrusion Detection Systems & Multisensor Data Fusion*," Communications of the ACM, Vol. 43, No. 4, April 2000, pp. 99-105
- [5] *Working Group on Rule-based Systems*, International Game Developers Association, The 2004 AI Interface Standards Committee (AIISC) Report, 2004.
- [6] [http://en.wikipedia.org/wiki/Thomas\\_Bayes](http://en.wikipedia.org/wiki/Thomas_Bayes)
- [7] Luckham, D., The Power of Events, Addison Wesley, Pearson Education Inc., 2002.
- [8] *TIBCO BusinessEvents User's Guide*, TIBCO® BusinessEvents, Software Release 1.2, September 2005.
- [9] *Jess®*, *The Rule Engine for the Java™ Platform*, The Rete Algorithm, Version 7.0b5, DRAFT, 5 January 2006.

## Additional Reading

- Waltz, E. and Llinas, J., Multisensor Data Fusion, Artech House, Boston, MA, 1990.
- Hall, D., and Llinas, J., "*An Introduction to Multisensor Data Fusion*," Proceedings of the IEEE, Vol. 85, No. 1, IEEE Press, 1997.
- Bass, T., "*The Federation of Critical Infrastructure Information via Publish and Subscribe Enabled Multisensor Data Fusion*," Proceedings of the Fifth International Conference on Information Fusion: Fusion 2002, Annapolis, MD, 8-11 July 2002, pp. 1076-1083.
- Carzaniga, A., Rosenblum, D., and Wolf, A., "*Design and Evaluation of a Wide Area Event Notification Service*," ACM Transactions on Computer Systems, Vol. 19, No. 3, August 2001, pp. 332-383.
- Carzaniga, A., "*Architectures for an Event Notification Service Scaleable to Wide-area Networks*," PhD Thesis, Politecnico di Milano, December 1998.



## **BIOGRAPHIES**

Don Adams (dadams@tibco.com) is currently a Vice President, Chief Security Officer and Chief Technology Officer – Worldwide Public Sector, at TIBCO Software Inc. In this position he provides expertise in security, government strategy and emerging technologies related to the TIBCO family of software solutions and service offerings. Prior to TIBCO, Mr. Adams was the Chief Technology Officer of TriStrata Inc. where he set the overall security philosophy, design and systems architecture for the revolutionary TriStrata Secure Information Management System. Prior to TriStrata, Mr. Adams spent six years at Sun Microsystems, where his last position was Principal Architect - Security and Networks. While at Sun Don participated as architect or chief architect on over four and a half billion dollars worth of government contracts they won. Prior to Sun Microsystems, Mr. Adams spent a highly decorated 23-year career in the United States Air Force. Mr. Adams started his career teaching at the Air Force Cryptographic Systems School in San Antonio Texas and spent the majority of his career in design, architecture, operations and maintenance of command, control, communications, computer and Intelligence (C4I) systems.. Mr. Adams was recently published as one of the contributing authors of the McGraw Hill Homeland Security Handbook. His chapter on critical concepts for IT in homeland security covers both Sense and Respond and Predictive Response for enterprise, counterterrorism and healthcare emergency response.