

# Big Contextual Data Event Processing on a Software Defined Server: Why, How, and How Fast

---

“The TIBCO Fast Data architecture—supported by an in-memory data grid—is designed to capture all types of data from any source, not only application data. More importantly, it transforms decisions into action by connecting to applications or by starting a business process. Providing an in-memory data grid, the BULL bullion server is the glue, the means by which you can move from looking into the rear view mirror to looking into the future.”

---

—Roman Bartik, managing director, alliances and business development, TIBCO Software

## SUMMARY

Analyzing sets of data helps businesses understand missed opportunities and threats that could potentially be captured going forward. Classical Business Intelligence only supplies information about what happened in the past, with no ability to understand, affect, and act on what is happening now or predict what will happen soon. Big data, while providing more insight, will likely be too slow to meet the velocity challenge.

This is where Fast Data comes in. It provides the ability to detect, affect, and act on situations as they are happening. Event processing is at the heart of a Fast Data architecture.

This paper examines processing and memory requirements that can drive rich event processing scenarios, illustrated through a retail use case. It also explains the benefits and performance of real-time event processing and presents Fast Data use cases in various markets.

## THE IMPORTANCE OF BIG CONTEXT

What is event processing? It's processing a flow of incoming events to identify situations that need attention—and as soon as identified—responding to affect the best outcome. “Best” here really depends on the business value being targeted: risk prevention, costs savings, revenue maximization, or customer satisfaction improvement.

As in human decision-making, the more context available for automated decisions, the better the decision will be. “Context” is built from data—historic, reference, or profile data—which is accessed by decision logic. Because the event processing application needs to respond to incoming information in real time, in-memory computing becomes critical for enabling fast scans through contextual data.

## EXAMPLE OF BIG CONTEXT IN RETAIL

An established retailer sells through a thousand of points of sale. Like many others, this retailer faces fierce competition from global e-commerce services.

How can it avoid being a showroom for e-commerce rivals? How can it efficiently transform shopper visits into transactions? The retailer needs to provide the shopper with a comfortable and trusted zone, and make the purchase decision easier by initiating a personalized interaction.

To enable this personalization and propose the right offers, the retailer collects and stores lots of information about each shopper:

- What the shopper usually buys
- What offers the shopper accepted
- What products the shopper searches for and has browsed recently
- The shopper's age, social status, and other details

“Data is now the power that fuels business. But to convert this abundant raw energy into concrete benefits, you need a chain of expertise. You need to understand the information you’re dealing with, choose and implement the best technology, and learn to identify business challenges—all while ensuring the necessary performance and security. As an integrated operator, we master all the links in the chain, from data infrastructure to data management. That’s Bull’s great strength.”

— Arnaud Bertrand, executive vice president at Bull, an Atos company

TIBCO BusinessEvents is a complex event processing (CEP) solution that abstracts and correlates meaningful business information from the events and data circulating through your information systems. By detecting patterns within the real-time flow of events, it can help you detect and understand unusual activity, as well as recognize trends, opportunities, and problems. Publishing business-critical information in real time to your enterprise systems or dashboards, you can visually see and understand what’s happening in the moment and facilitate proactive decision-making.

Each bit of information makes the personalization richer. Multiplied by millions of individual shoppers, the data rapidly adds up, requiring large data space capacity.

Another key concern for any retailer is managing its stock efficiently and being able to cross check customer information against product inventory across different locations and sales channels.

All this data, as rich as it can be, is however useless if you can’t collect and aggregate streams of shopper events, correlate them with contextual data, and make rapid decisions while the window of opportunity is still open:

- **Event:** Shopper comes into the shop and requests information about product A.
- **Context:** Shopper has lots of loyalty points. Product B, the same type as product A but with enhancements and a higher price, is over-inventory.
- **Decision:** Offer product B at the same price as product A with a debit of X loyalty points.

This example shows typical detection and action logic for a retailer. A real production system will have to detect and act upon dozens to hundreds of similar types of situations.

Having hundreds of rules executing in milliseconds to seconds and scanning terabytes of data leaves no other technical choice than storing data in memory. Therefore, data space directly translates into memory space.

## HARDWARE PLATFORMS AND BIG MEMORY

There are typically two ways to scale in-memory applications: horizontally and vertically.

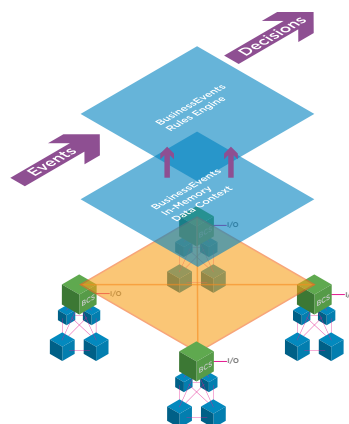
### HORIZONTAL SCALING

Horizontal scaling—using the combined storage of dozens of small machines —works well if you can co-locate data with the logic that uses it. When this co-location isn’t possible, you need cross-node communication to find and access the data, which adds traffic that impacts performance and may limit scalability.

If this is the scenario, vertical scalability using large, modern CPU and memory, like the bullion from Bull, may provide a better option. It simplifies application development and management because data and work do not need to be partitioned.

### VERTICAL SCALING

Vertical scaling—adding processors and storage to individual nodes—allows you to increase capacity on-the-fly, according to need. As an example, the bullion server can be dynamically reconfigured to provide elastic scalability across 2 to 16 processors and 48 GB to 24 TB of memory. Additional memory or storage can be hot deployed. This type of solution can be very affordable because you don’t need to invest in huge capacity up front. In many situations, vertical scalability is an easy choice for handling more complex rules and greater volumes of events.



**RESULTS: TEST 1,  
DATA CREATION**

- 5.4 million objects created/second
- 120,000 events processed/second

**RESULTS: TEST 2,  
LARGE IN-MEMORY  
DATASETS**

- 32,447 events processed/second
- Latency <1 ms

**RESULTS: TEST 3,  
RULES EXECUTION**

- Fewer rules executing: 40,000 operations/second, limited by the network
- More rules executing: 3,500,000 rule evaluations/second per CPU core

**EVENT PROCESSING ON BIG MEMORY****TEST 1**

Real world benchmarks are critical for validating the true scalability of any software and hardware combination. Bull and TIBCO tested TIBCO BusinessEvents® complex event processing platform on bullion machines.

The first set of tests validated the amount of data that TIBCO BusinessEvents can create and manage as one logical segment using large memory.

Two different methods were used:

- 1 Loading data objects from files in chunks of 5,000 objects, the full power of the machine was used. It loaded the entire set of one billion objects at a rate of 5.4 million objects created/second. This benchmark describes the cold start performance when data is loaded from other systems.
- 2 Creating one object when each event is received, representing a retail sales transaction. Each object included about 30 columns. TIBCO BusinessEvents processed 120,000 events/second, sustaining a flat processing rate until one billion objects were created. This benchmark describes the performance the system would deliver when managing up to one billion transaction records or accounts.

**TEST 2**

The next set of tests measured how efficiently the rules-based event processing of TIBCO BusinessEvents can access and manipulate large in-memory datasets. A simple scenario was used, where one rule would trigger for each event received, and execute as follows:

- Get events from the TIBCO message bus (TIBCO Enterprise Message Service™)
- Parse event payload
- Get contextual data from in-memory grid or create it if not found
- Lock it
- Update/enrich the contextual data
- Send the response through the message bus

This action is similar to processing that checks a purchase against a customer record and updates the record. Using three-quarters of machine resources, (one quarter being left to the test drivers), TIBCO BusinessEvents processed at a rate of 32,447 events/second with a processing latency under a millisecond.

**TEST 3**

Last, we validated the performance of pure rules execution with a complex configuration where many rules were evaluated for each incoming event.

Execution of one rule led to execution of another, with the total number varying.

For a low number of rules, the processing rate was limited by event delivery capacity, reaching 40,000 operations/second. Increasing the number of rules creates more CPU intensive processing. Network performance was less of a concern. In this configuration, TIBCO BusinessEvents reached 3.5M rule evaluations/second in each CPU core. This kind of multiple rule evaluation is required in cross-sell or event-driven marketing applications, where many potential offers are tested to find the best offer for a given customer interaction.

## BULLION, THE MOST ADVANCED WORKSPACE FOR FAST DATA

With between 2 and 16 Intel Xeon E7 v2 processors and up to 24 TB of memory, bullion provides an x86 platform capable of unmatched response, flexibility, security—and especially memory. Along with its reliability, availability, and serviceability (RAS) features, it cost-effectively meets all challenges. You can dynamically hot-deploy resources like memory or I/O thanks to a patented blade design and reconfigure bullion software-defined hardware for on-demand capacity expansion and unrivalled elasticity in the face of changing needs. Additionally bullion breaks the Intel memory reference barriers for extreme effectiveness in memory-intensive environments.

## OTHER MARKETS THAT BENEFIT FROM BIG CONTEXTUAL DATA

Previously, we used the example of a retailer to illustrate the need for big context, but this need is certainly not limited to retail. A partial list of examples:

- **Telecommunications:** An operator captures high volumes of events from its network and customer care channels to assess the real-time state of network elements correlated with customer data and interactions. The objective is to provide the best service to customers and maximize efficiency of network maintenance operations.
- **Transportation:** Airline companies hold passenger information, track their reservations and luggage, and correlate with flight and scheduling information. The objective is to react in the best way possible to unpredicted events to continue to provide optimum service and minimized operations cost.
- **Logistics:** Hundreds of thousands of packages, each with different characteristics and requirements, are tracked along their delivery paths, through multiple partners, across multiple geographies.
- **Internet of Things:** A growing number of subscribers and devices, each sending multiple types of information at a high frequency are tracked over a period of time. The objective is to provide meaningful information, enough to detect patterns and trigger alerts or corrective actions.
- **Payments:** Payments delivered through multiple channels and types of payment instruments are aggregated to enrich intelligence on usage, help prevent fraud, and create up-sell and cross-sell opportunities.



**Bull, an Atos company,** is the trusted operator for enterprise data. The Group, which is firmly established in the Cloud and in Big Data, integrates and manages high-performance systems and end-to-end security solutions. Bull's offerings enable its customers to process all the data at their disposal, creating new types of demand. Bull converts data into value for organizations in a completely secure manner. It currently employs around 9,200 people across more than 50 countries, with over 700 staff totally focused on R&D. In 2013, Bull recorded revenues of €1.3 billion. For more information: <http://www.bull.com>/<http://www.facebook.com/BullGroup/>/[http://twitter.com/bull\\_com](http://twitter.com/bull_com)



**Global Headquarters**  
3307 Hillview Avenue  
Palo Alto, CA 94304  
+1 650-846-1000 TEL  
+1 800-420-8450  
+1 650-846-1005 FAX  
[www.tibco.com](http://www.tibco.com)

**TIBCO Software Inc.** is a global leader in infrastructure and business intelligence software. Whether it's optimizing inventory, cross-selling products, or averting crisis before it happens, TIBCO uniquely delivers the *Two-Second Advantage*®—the ability to capture the right information at the right time and act on it preemptively for a competitive advantage. With a broad mix of innovative products and services, customers around the world trust TIBCO as their strategic technology partner. Learn more about TIBCO at [www.tibco.com](http://www.tibco.com).

©2014-2015, TIBCO Software Inc. All rights reserved. TIBCO, the TIBCO logo, TIBCO Software, Two-Second Advantage, TIBCO BusinessEvents, and TIBCO Enterprise Message Service are trademarks or registered trademarks of TIBCO Software Inc. or its subsidiaries in the United States and/or other countries. All other product and company names and marks in this document are the property of their respective owners and mentioned for identification purposes only.  
03/03/15