



Algorithmic Operations: Integrating Real-Time Analytics into your Business

Richard Tibbetts

CTO, TIBCO Event Processing

QCon San Francisco 2014

Changing
Consumer
Expectations

Internet
Connected
Devices

Machine Learning
and Predictive
Analytics

Empowered
Super Users

Technology
Initiatives
Outside IT

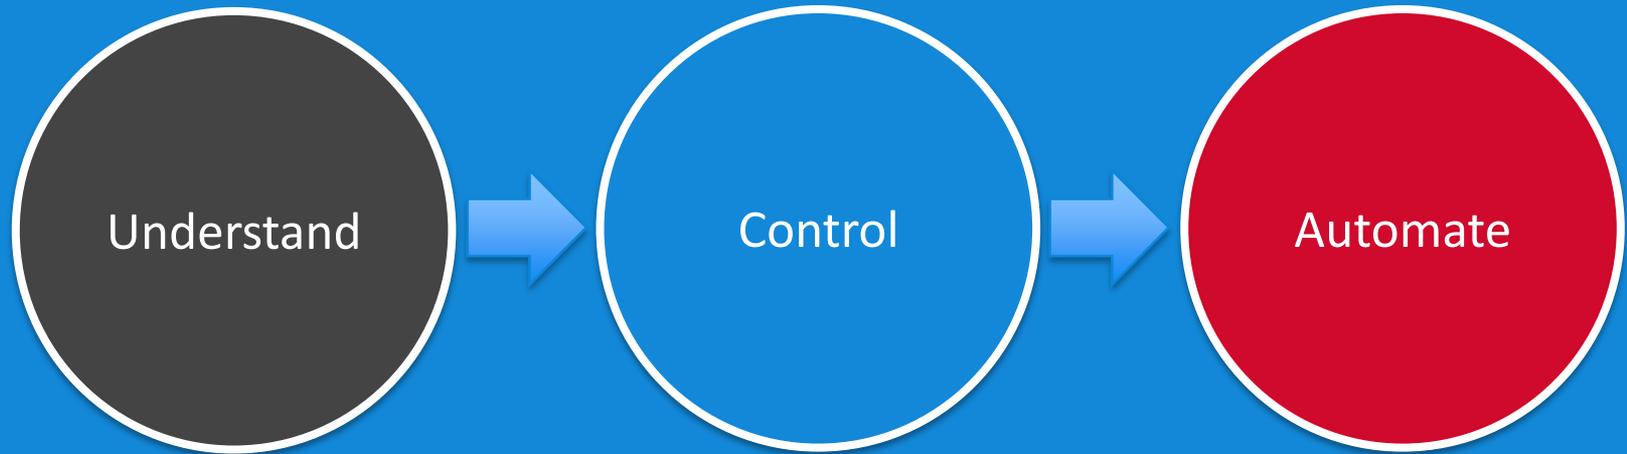
Disruptive
Full Stack
Startups

Mobile Devices
and Apps

API Economy

How will you transform into a digital business?

To automate your business with analytics,
you must give the business
visibility and control





Digital Business: From Big Data to Fast Data

Operational Intelligence and Algorithmic Operations

Lessons Learned from Key Fast Data Use Cases

TIBCO Event Processing



Big Data: There can be an overwhelming volume





| Most solutions focus on accumulating and analyzing



TIBCO™

| Especially analyzing things which already happened





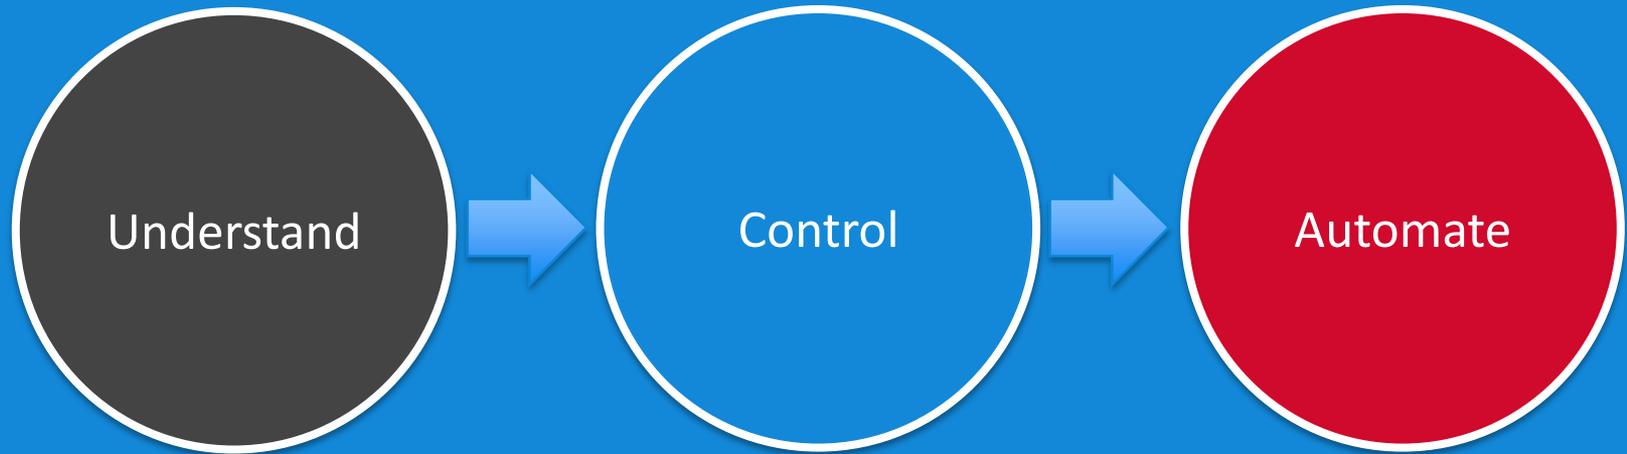
But how will you access the data when you need it?



Fast Data

FAST DATA IS PROCESSING BIG DATA
IN REAL-TIME TO GAIN INSTANT
AWARENESS AND INSTANT ACTION

To automate your business with analytics,
you must give the business
visibility and control



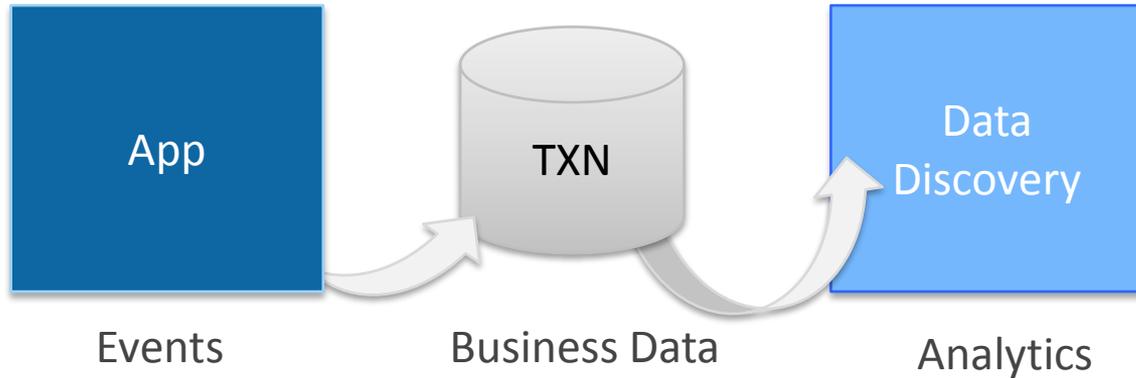


Digital Business: From Big Data to Fast Data

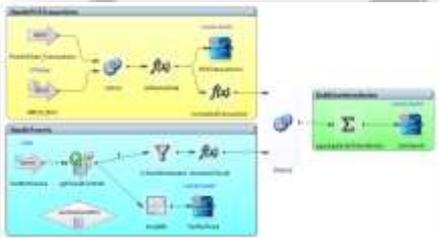
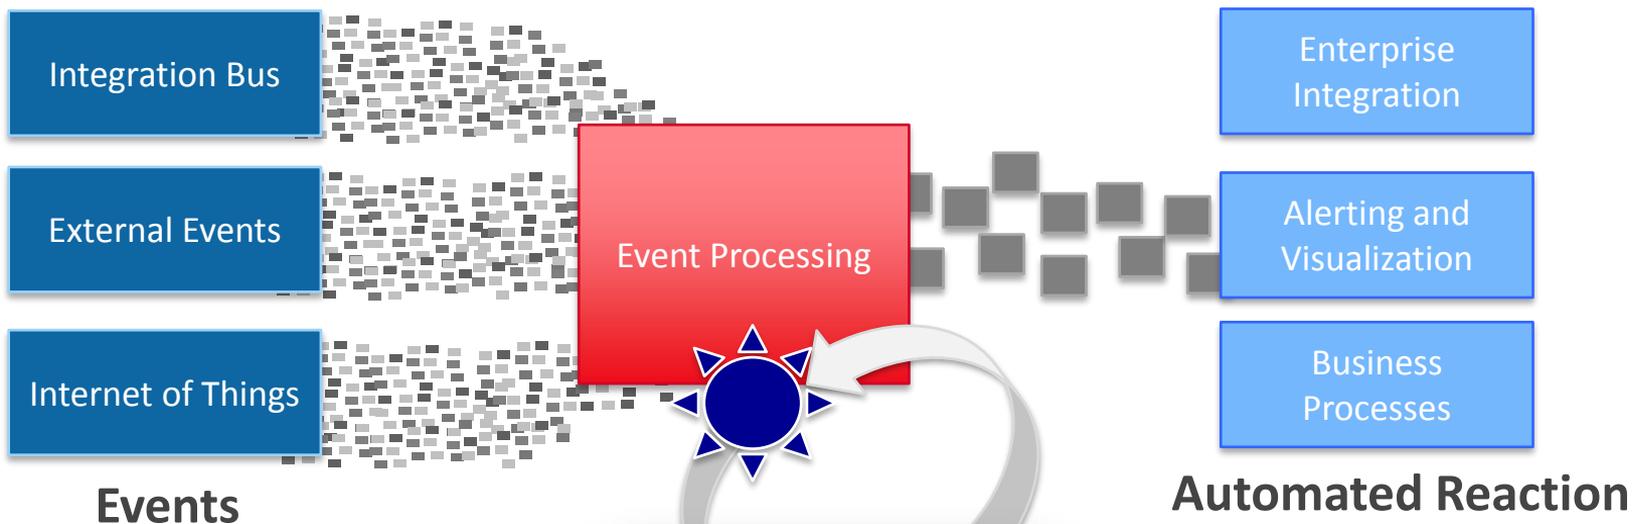
Operational Intelligence and Algorithmic Operations

Lessons Learned from Key Fast Data Use Cases

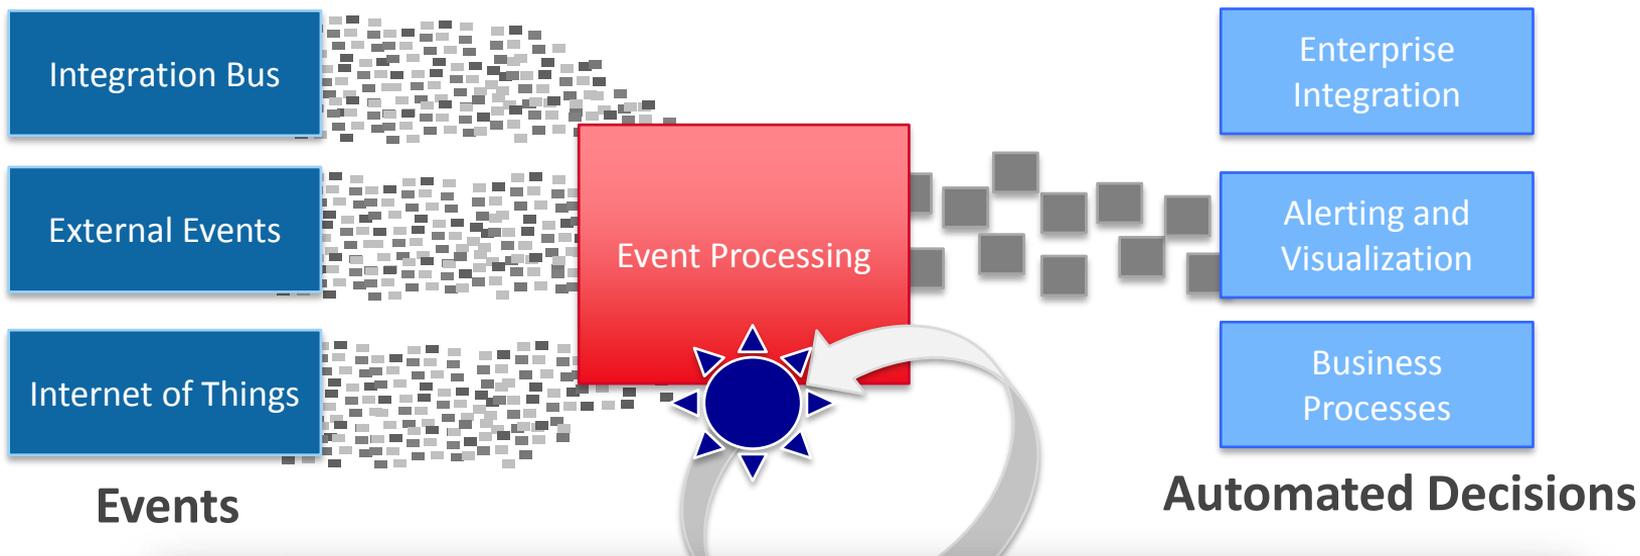
TIBCO Event Processing



“Who cares if we find out we lost a customer after he left?”

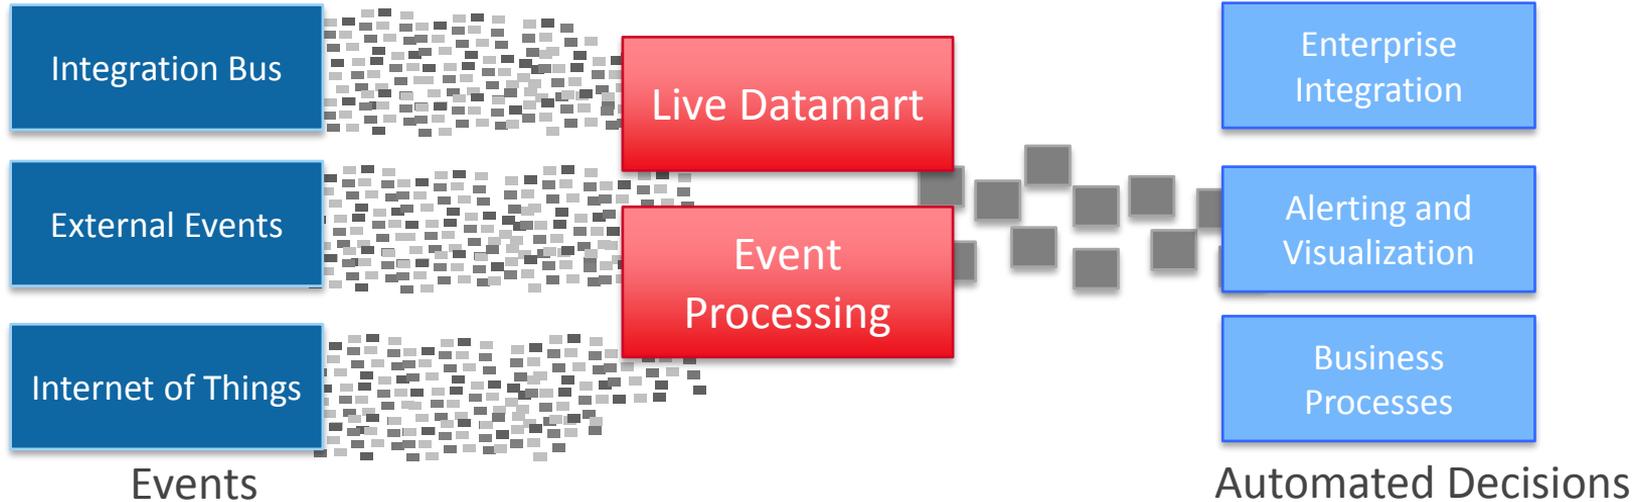


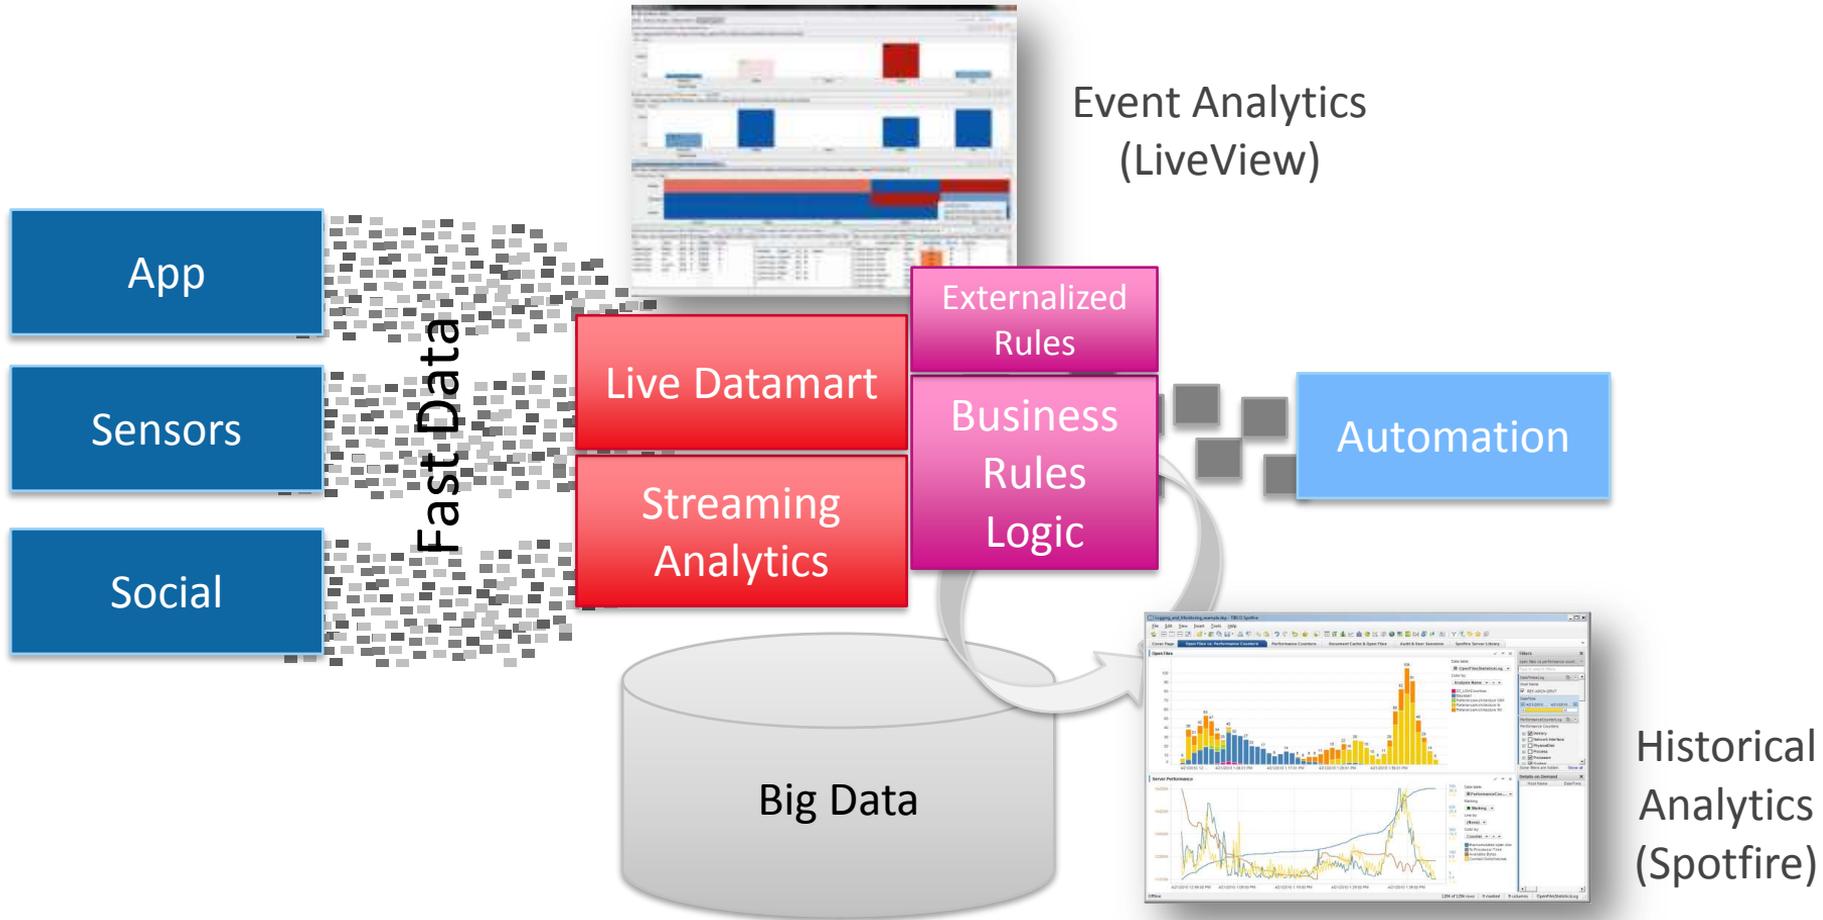
Streaming Analytics & Event Driven Rules



ID	- Condition Area					- Action Area
	location.sensitive	location.level	person.interesting	personLocation.duration	personLocation.madeCall	location.level
1	true	"NORMAL"	false	>= 300	true	
2	true	"NORMAL"	true	>= 300		"ELEVATED"
3	true	"ELEVATED"	false	>= 200	true	
4	true	"ELEVATED"	true	>= 100		"CRITICAL"
5	true	"CRITICAL"	true			

Event Rules





Operational Intelligence

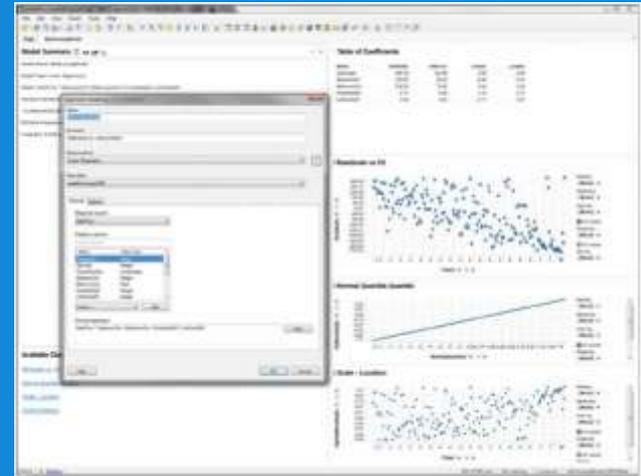
Human decisions in real time
informed by up to date information



The Challenge: The right
information, ready to be used by
the right people at the right time.

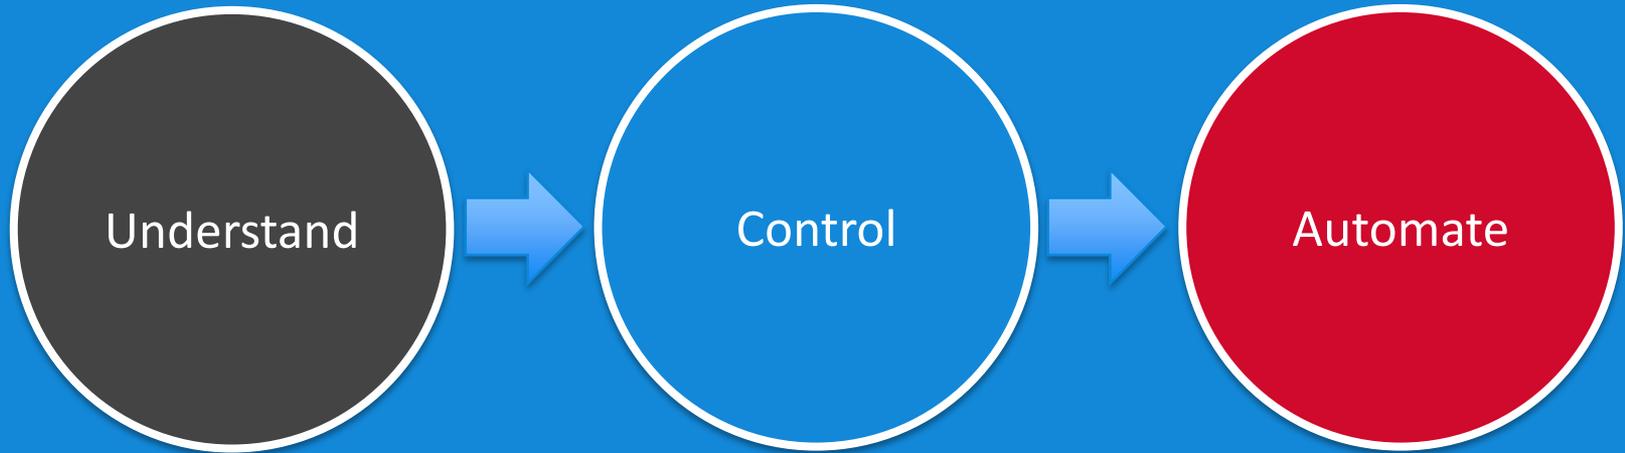
Algorithmic Operations

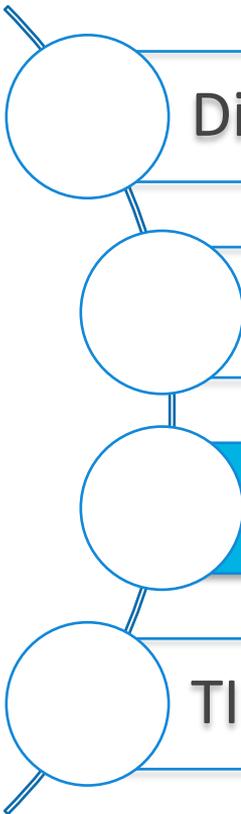
Automated action based on analytic
models of history combined with
live context and business rules



The Challenge: Empowering business
stakeholders to understand and
control the selection of models

To automate your business with analytics,
you must give the business
visibility and control





Digital Business: From Big Data to Fast Data

Operational Intelligence and Algorithmic Operations

Lessons Learned from Key Fast Data Use Cases

TIBCO Event Processing

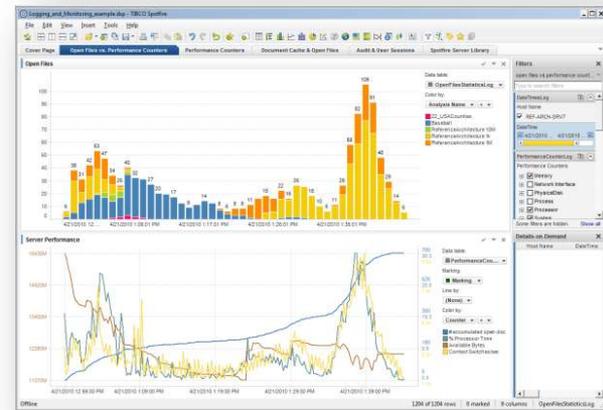
- **Trading: Execution Optimization**
- **Telco: Customer Churn Prevention**
- **Banking: Real-Time Compliance**
- **Marketing: Real-Time Campaign Attribution**
- **Dealing: Counterparty Behavioral Analysis**
- **Manufacturing: Yield Improvement and Cost Reduction**

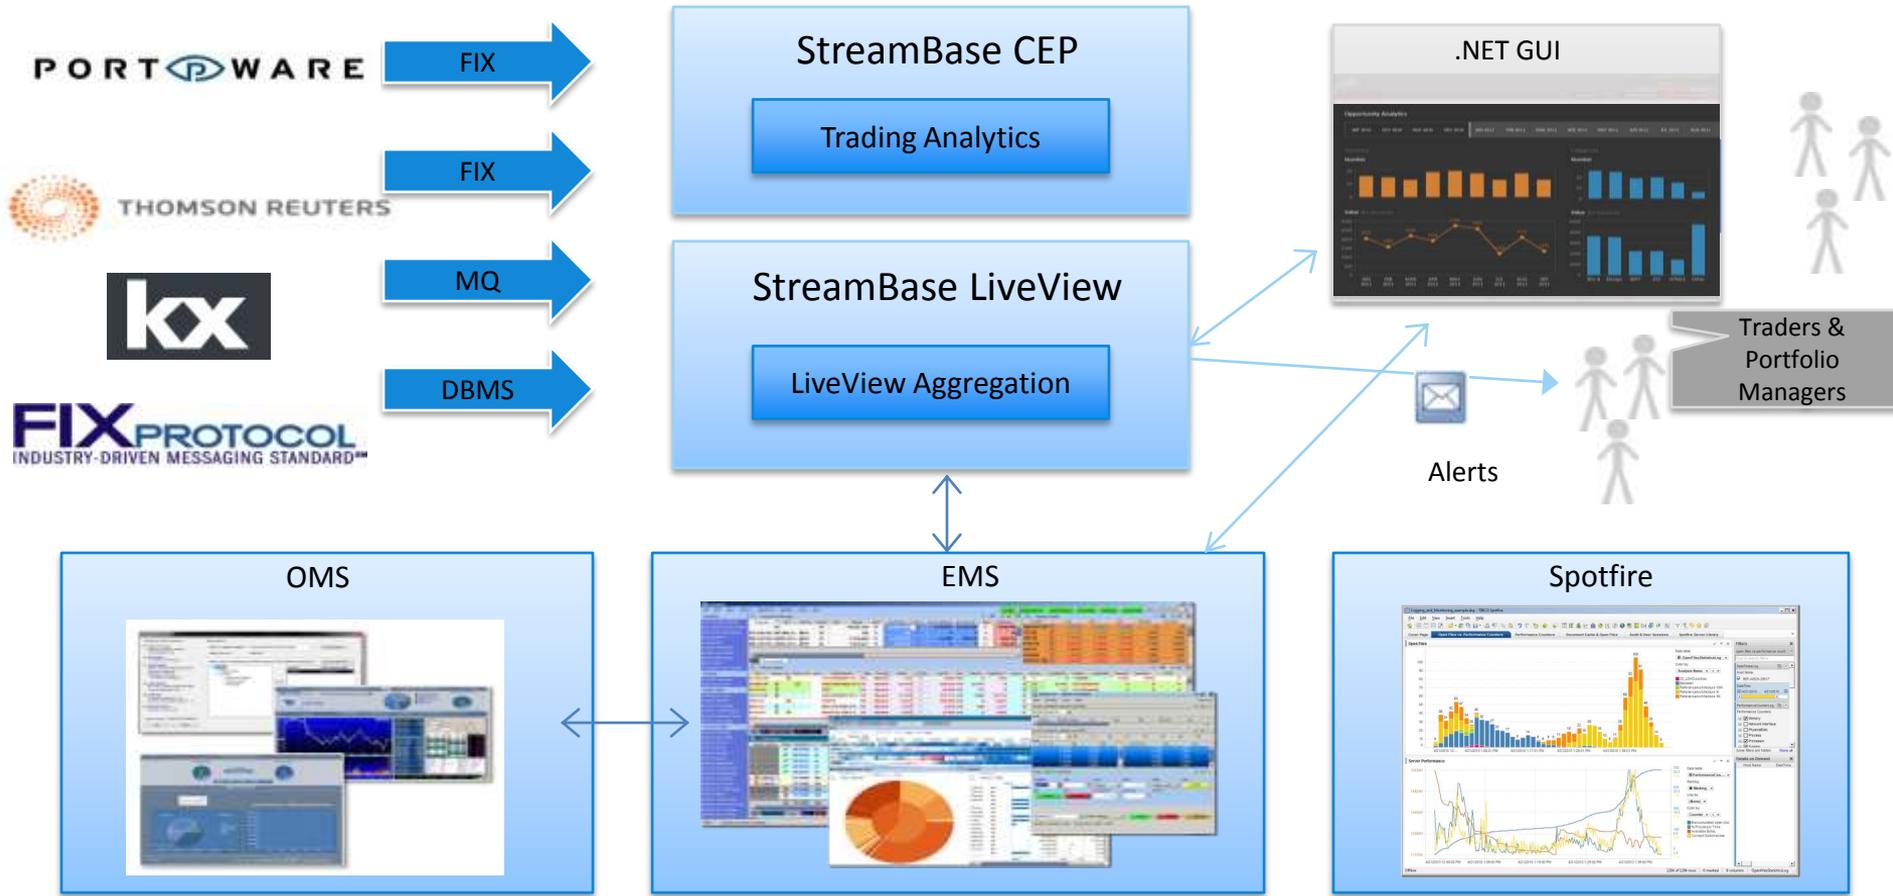
Execution Optimization

“With StreamBase & Spotfire, we can analyze everything that impacts trade performance, and adjust to it on-the-fly.”

- Head of IT, Asset Management

- **Situation: Major asset manager needs more insight**
 - Which trades are slipping? How close to target are we for participation rates? How can we adjust our trading in the market now?
- **Problem: How to get better price discovery?**
 - Need real-time analytics using market data, interface to OMS, EMS (Portware), and internal systems
- **Solution: StreamBase, LiveView, and Spotfire**
 - Developed a live visualization system with StreamBase CEP, LiveView, a custom .NET visualization front end for traders & PMs, and Spotfire for historical analytics
- **Impact: Improved trading performance**
 - Solution allows traders and PM's to understand where their trade performance is going, and adjust trading strategies during the day





Lesson #1

Identify opportunities being missed by operational staff because they are too small or too hard to optimize, and automate those first for quick wins.



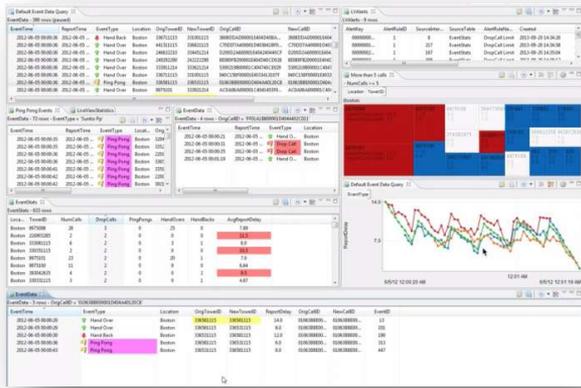
Telecommunications: Customer Churn Prevention

“If I can see network issues in the moment, I can fix them before my customer notices, and **automatically improve call center efficiency** by letting customers proactively know we’re working on the problem.

CIO, major telecom



TIBCO EP for Telco Customer Churn Prevention



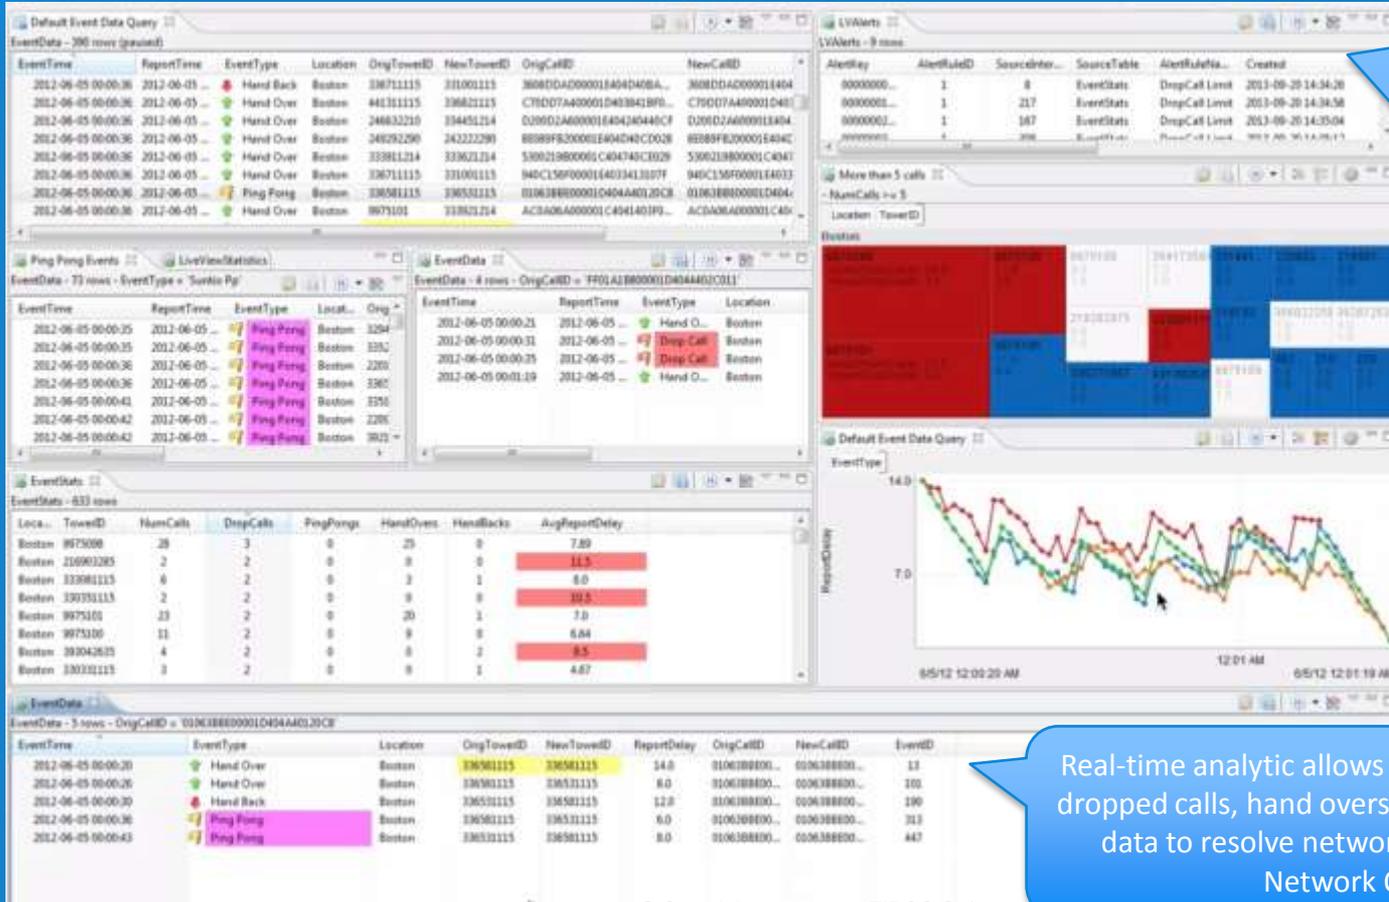
- **The Situation Before: Customer Churn Projections to Cost \$28M+ / Year**
 - Telecom flying blind in real-time with respect to network outages that can effect valuable customers and social influencers
- **Negative Consequences: Call Centers Can't React to Problems in the Network as They Happen**
 - Systems don't talk to each other in any way other than batch, so the call center handles calls for network issues that are already identified
- **Positive Business Outcome: Predictive Real-Time Customer Service**
 - Call center systems proactively inform effected customers effected by network issues
 - Network systems proactively re-route high-value customers around effected network areas in real-time
- **How We Do It: TIBCO Event Processing for Real-Time Analytics in Telco Networks and Call Center Action**
 - High-speed network and user event analytics, visualization, alerting, and action to optimize customer call center behavior

"If we detect and correct network problems before the customer notices, we could reduce revenue loss based on churn by as much as \$28M a year."
- Head of IT, major telco provider

- **One Telecommunications Firm Customer Churn Study...**
 - Studied 8M customers, 7 billion service calls for a 3 week sample
 - Found 1M clusters of callers, 120K “Dropped Call Watch List” clusters
 - Identified 40K network influencer “signatures”
 - Found 4000 watch list customers who already cancelled service and can influence others; impacted an additional 18K customers when they cancelled
- **Impaction: Projected \$28M in Lost Annual Revenue**
 - Customer loss directly linked to network outages whose impact might have been mitigated if real-time visibility was available
 - By providing influencers better service in real-time, and optimizing call centers to support high-value customers, revenue loss can be dramatically reduced
 - By predicting which customers will call due to service outage and intercepting calls, customers will know immediately that the Telecom is on top of the issue. ***Builds trust.***



Call Center Analytics, Driven By Network Insight



High priority alerts for Network Operators. e.g., “These 100 high value customers are being effected by these network events. Give them high priority access and fix the network.”

Live aggregation and slice and dice of call hand backs, dropped calls, and Ping-Pong events – all signs of problems

Real-time analytic allows operators to view in real-time the dropped calls, hand overs between towers – these is critical data to resolve network issues in the moment, in the Network Operations Center.

Lesson #2

Empower operational staff with contextually relevant information, decision making authority *and* recommended action

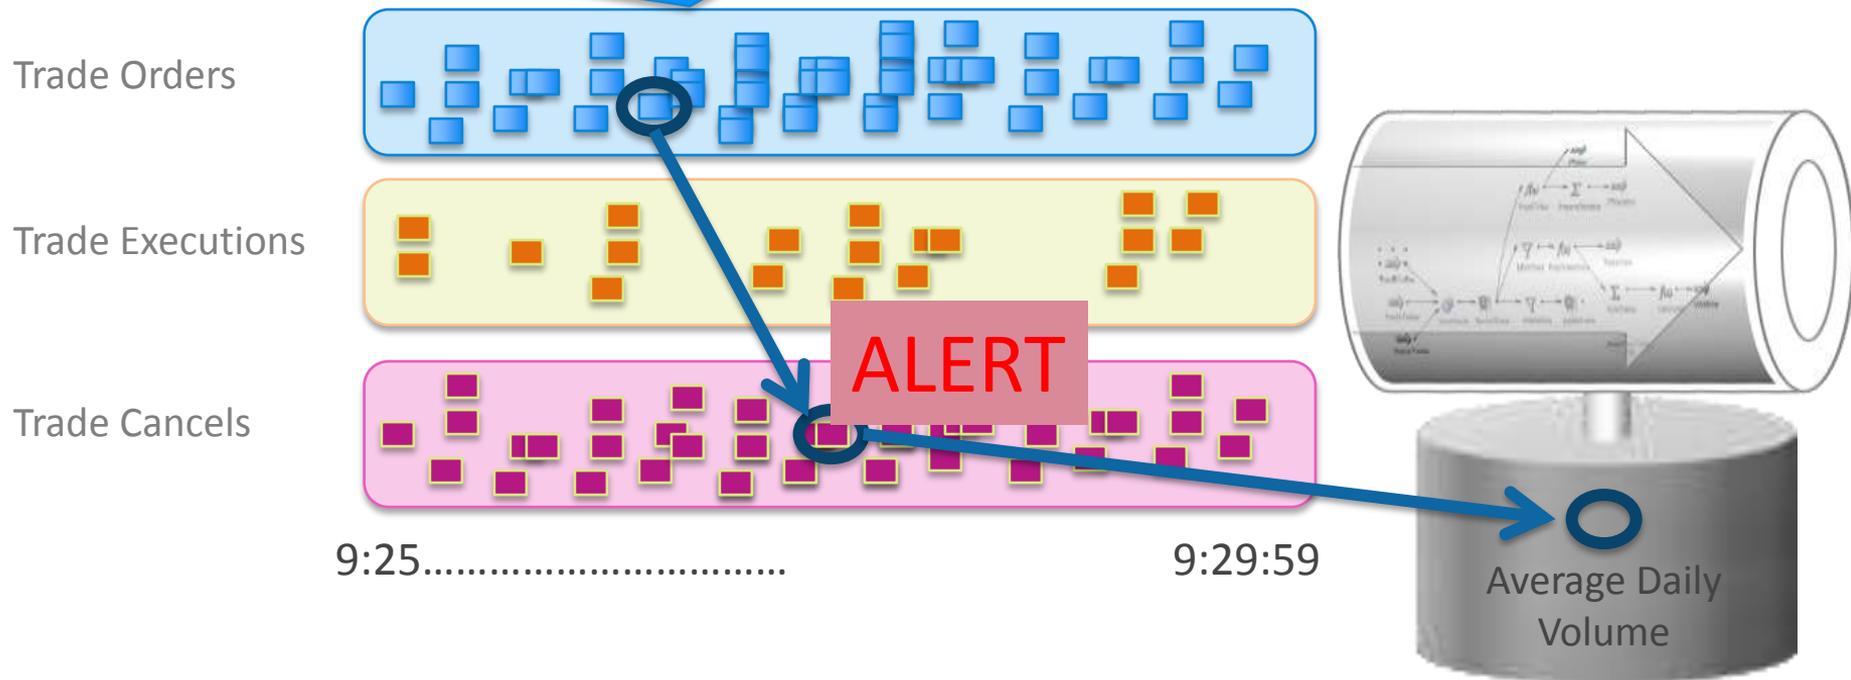


Continuous Compliance

“In December 2012, **Knight Capital lost \$460M in under 40 minutes**. That changed everything. Now, it’s no longer acceptable to run our business based on end-of-day reports.”

- Head of Risk Management, top 3 bank

“WHEN any security has **SIMULTANEOUS TXNS** with **THE SAME PRICE**, within **30 SECONDS**, flag the TXN as a potential **WASH TRADE**”



CRSS

Risk Alerts Case Management Add Tool

Firm Count: 53

Buy-Side Clients Sell-Side Clients

3/27/2013 (Wed)

Displaying All Firms/Alerts

Firm Id	Firm	Client Type	Large Order	Selling	Ramping On-Close	Ramping On-Open	Marking The-Close	Wash Sales	Layering	Spool 1	Spool 2	Spool 3	Firm Total
	596 Firm506	BuySide			130 15								192
	4626 Firm4626	SellSide			120 15								178
	8661 Firm8661	SellSide			180 15								129
	2513 Firm2513	BuySide			72 3								74
	2897 Firm2897	SellSide	1 10		62 4								49
	6232 Firm6232	SellSide			23 0								23
	5529 Firm5529	SellSide			99 38			52 36					153
	3221 Firm3221	SellSide			33 0								33
	5204 Firm5204	SellSide	3 10		9 0								12
	3181 Firm3181	BuySide			6 0					2 0	3 0 0		113
	838 Firm838	BuySide			7 0								7
	2947 Firm2947	SellSide			4 0								4
	3173 Firm3173	SellSide		5 10	4 0								5
	2100 Firm2100	SellSide		1 10	7 0								4
	4282 Firm4282	SellSide			1 0								3
	4761 Firm4761	SellSide			3 0								3
	5530 Firm5530	SellSide			1 0					1 0	26 0		38
	5933 Firm5933	BuySide			2 0								2
	4168 Firm4168	SellSide			3 0								3
	2094 Firm2094	SellSide			2 0								2
	1295 Firm1295	SellSide			2 0								2
	5367 Firm5367	SellSide			1 0	40 0							46
	4268 Firm4268	BuySide			1 0	1 0							2
	4854 Firm4854	SellSide			1 0	1 0							2

Edit Filter

All Firms Alert Count 539

9:30:00 AM 9:30:00 AM Get Clear Filtered Rows: (440)

Firm3181 Alerts Assigned to Cases 94

Ramping On Close Displaying Unmanaged Alerts

Code	RC Factor	Alert Time	Sending To	Symbol	Symbol Eff	30-Day Adj	Side	Previous Close Price	Exec Price	Exec Shares	Order Qty	Order Price	Client Name	Client Type	Op
IB RC	0.212704	3/27/2013 9:49 AM	9:49:01 AM	GRPN		14,333,617	Sell	5.9825	4.71	100	75,000	0	Firm4626	SellSide	
IB RC	0.212704	3/27/2013 9:49 AM	9:49:01 AM	GRPN		14,333,617	Sell	5.9825	4.71	100	34,300	4.5	Firm2513	BuySide	
IB RC	0.212704	3/27/2013 9:49 AM	9:49:04 AM	GRPN		14,333,617	Sell	5.9825	6.71	100	75,000	0	Firm4626	SellSide	
IB RC	0.210032	3/27/2013 9:49 AM	9:49:25 AM	GRPN		14,333,617	Sell	5.9825	4.72	2,500	2,500	4.72	Firm5951	SellSide	
IB RC	0.211832	3/27/2013 9:49 AM	9:49:25 AM	GRPN		14,333,617	Sell	5.9825	4.72	100	100	4.72	Firm1129	SellSide	
IB RC	0.209361	3/27/2013 9:49 AM	9:49:41 AM	GRPN		14,333,617	Sell	5.9825	4.73	800	2,300	4.73	Firm5951	SellSide	
IB RC	0.209361	3/27/2013 9:49 AM	9:49:41 AM	GRPN		14,333,617	Sell	5.9825	4.73	180	2,300	4.73	Firm8661	SellSide	
IB RC	0.209361	3/27/2013 9:49 AM	9:49:41 AM	GRPN		14,333,617	Sell	5.9825	6.72	1,500	1,800	4.72	Firm8661	SellSide	
IB RC	0.212704	3/27/2013 9:50 AM	9:50:00 AM	GRPN		14,333,617	Sell	5.9825	4.71	100	75,000	0	Firm4626	SellSide	
IB RC	0.212704	3/27/2013 9:50 AM	9:50:00 AM	GRPN		14,333,617	Sell	5.9825	4.71	100	75,000	0	Firm4626	SellSide	
IB RC	0.212704	3/27/2013 9:50 AM	9:50:00 AM	GRPN		14,333,617	Sell	5.9825	4.71	100	34,300	4.5	Firm2513	BuySide	
IB RC	0.216047	3/27/2013 9:50 AM	9:50:21 AM	GRPN		14,333,617	Sell	5.9825	4.69	1,433	2,300	4.69	Firm5951	SellSide	
IB RC	0.216047	3/27/2013 9:50 AM	9:50:21 AM	GRPN		14,333,617	Sell	5.9825	4.69	100	2,500	4.69	Firm5951	SellSide	
IB RC	0.216047	3/27/2013 9:50 AM	9:50:21 AM	GRPN		14,333,617	Sell	5.9825	4.69	867	2,500	4.69	Firm5951	SellSide	

Edit Filter

* (V) [Double] Like "MCS" And [Symbol] = GRPN *

Ready

Username: hawkie

Connection Manager Status

CRSS Status

Connected to

Lesson #3

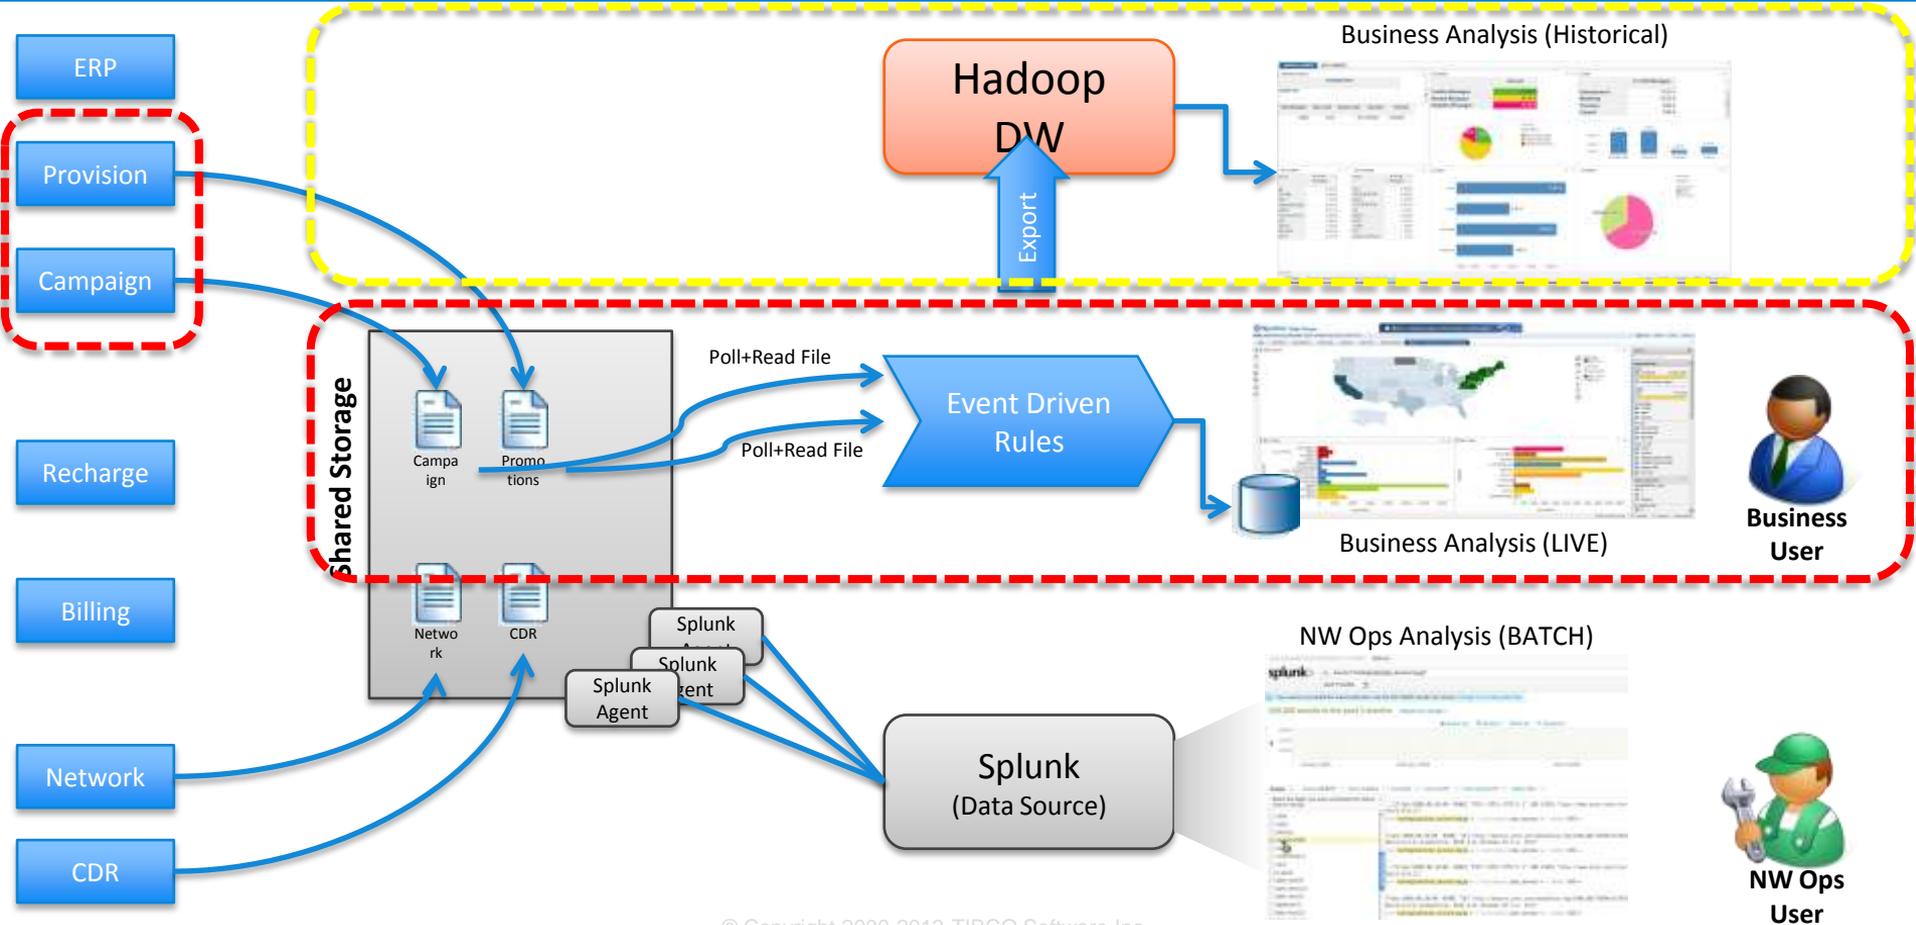
Alert and react to operational problems while they are still fresh to prevent risk snowballs and to keep outcomes positive



Marketing: Real-Time Campaign Attribution

- A leading cellular operator in Asia, which serves more than 130 million customers. The operator has experienced double digit growth over the last 3 years in prepaid, postpaid, and value added services
- **The Problem:** Bridging the gap between marketing campaigns launched and potential revenue that is generated
 - Complexity of data correlation between current campaign system and surrounding environment meant manually scripting and calculating the potential revenue
- **The Solution: Automated Campaign Tracking**
 - Integrate data from four systems
 - Campaign provision report
 - Campaign spikes: Alert on spikes in offer provisioning

- Replacing manual work with automated process
- Campaign revenue visualization by mapping campaign data to multiple sources
- Streambase provides a expandable platform for correlating different streams
- Live view provides source for tracking data movement in real time
- Spotfire can be used for war room reports
- Why TIBCO?
 - TIBCO solution will help view campaign flowing data (LiveView) and also correlate multiple streams in StreamBase (4 sources of data) and view revenue reports in Spotfire (Data to be made available in DB)
 - The existing campaign solution cannot provide reports. TIBCO can handle both the live view and historical – which will help Telkomsel understand data and spikes in a much better fashion
 - Live view will also help in raising alerts whenever campaign systems breach the threshold of bonus provision (This can be extended to multiple sources)



Lesson #4

Use event processing and rules to correlate events in real time to build context from data in decoupled or siloed systems

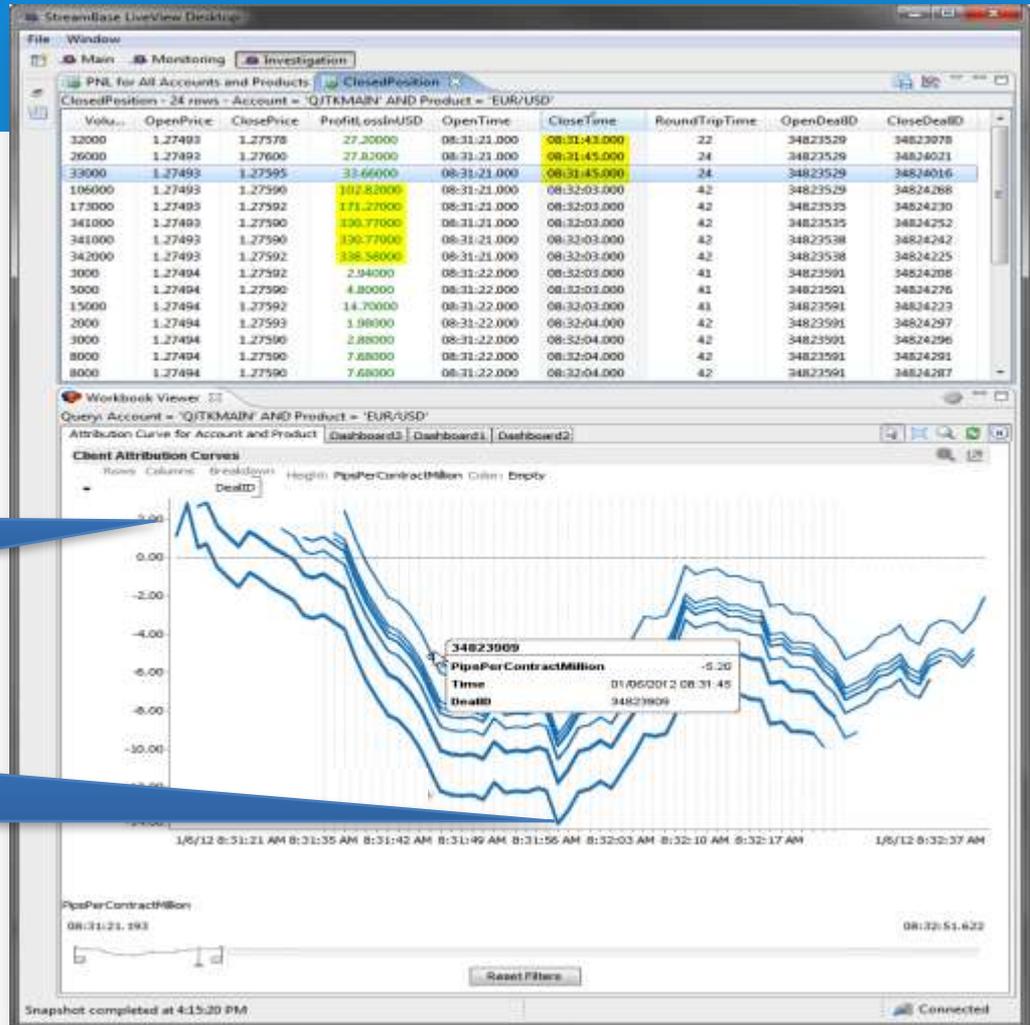


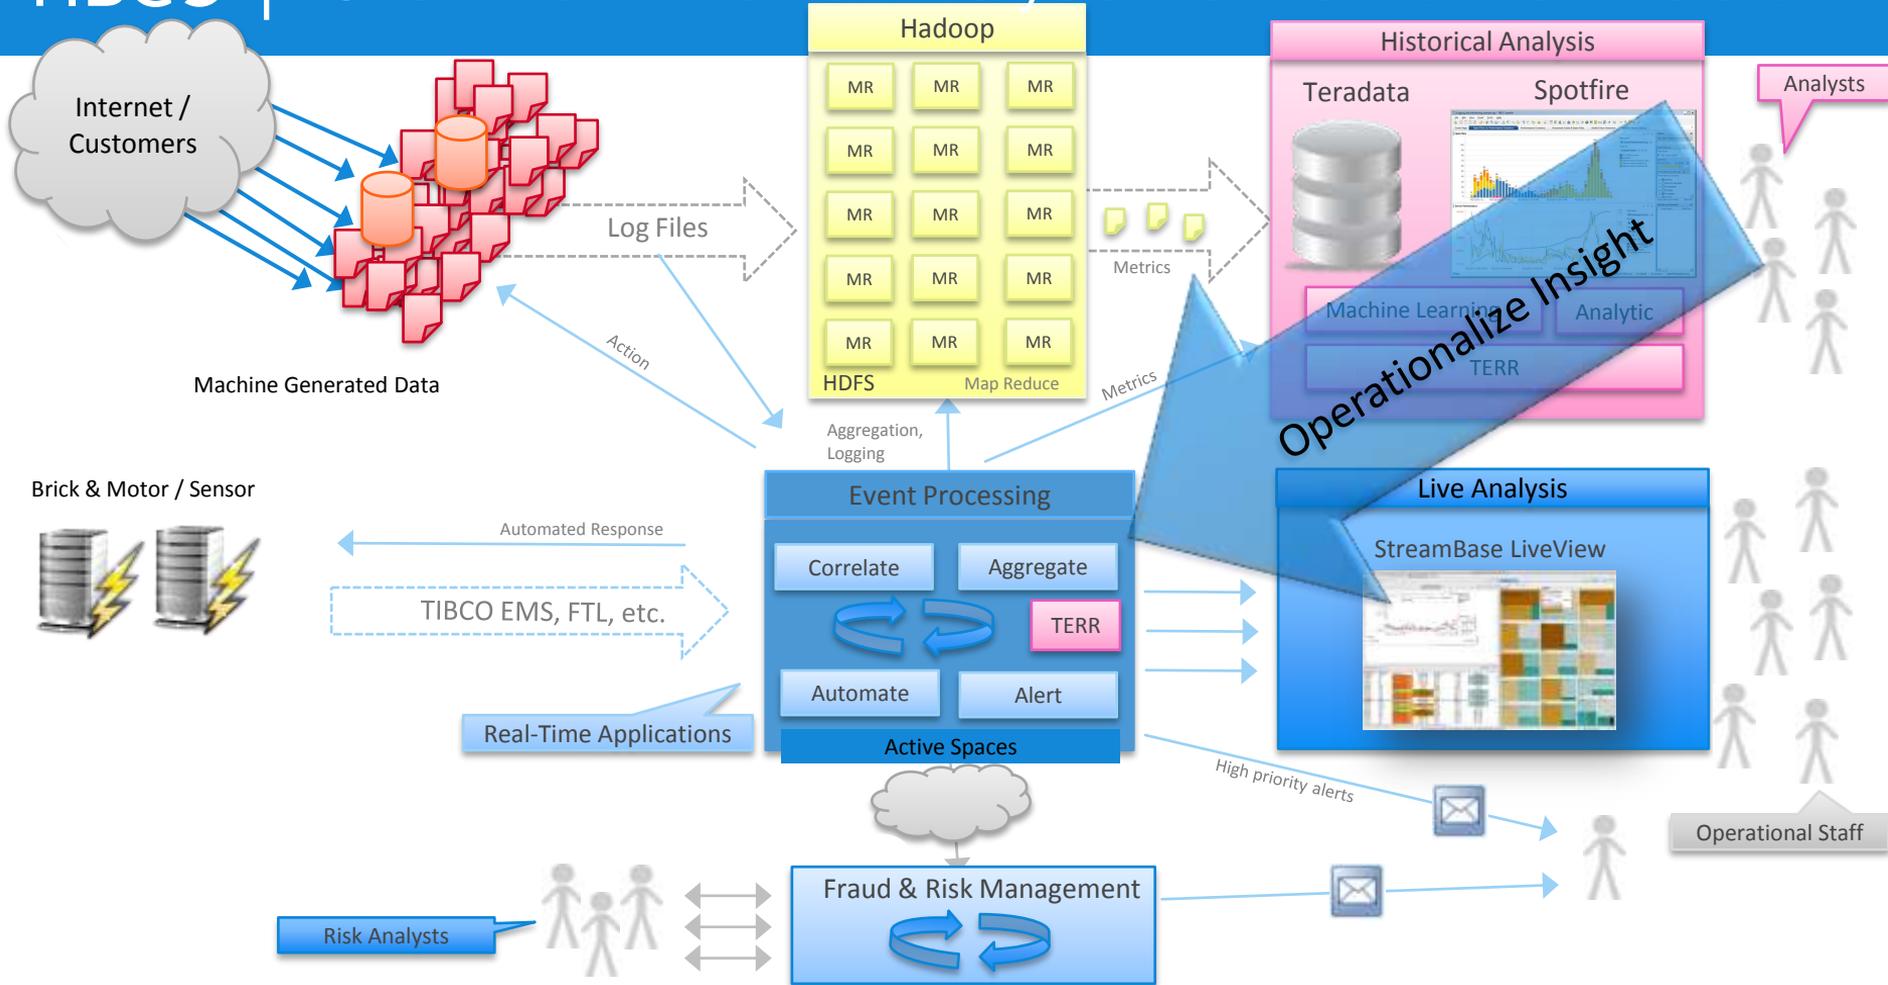
Counterparty Behavioral Analysis

- Identify customers acting on superior information
- Route trade flow elsewhere

At 8:31:21 AM, client sells 7 times in 10 seconds. At this moment, the firm made money

35 seconds later, the firm has lost over \$10,000 on these trades because the client had better market data





Lesson #5

Understand your data analytics pipeline and make the same data available to operational intelligence and algorithmic operations

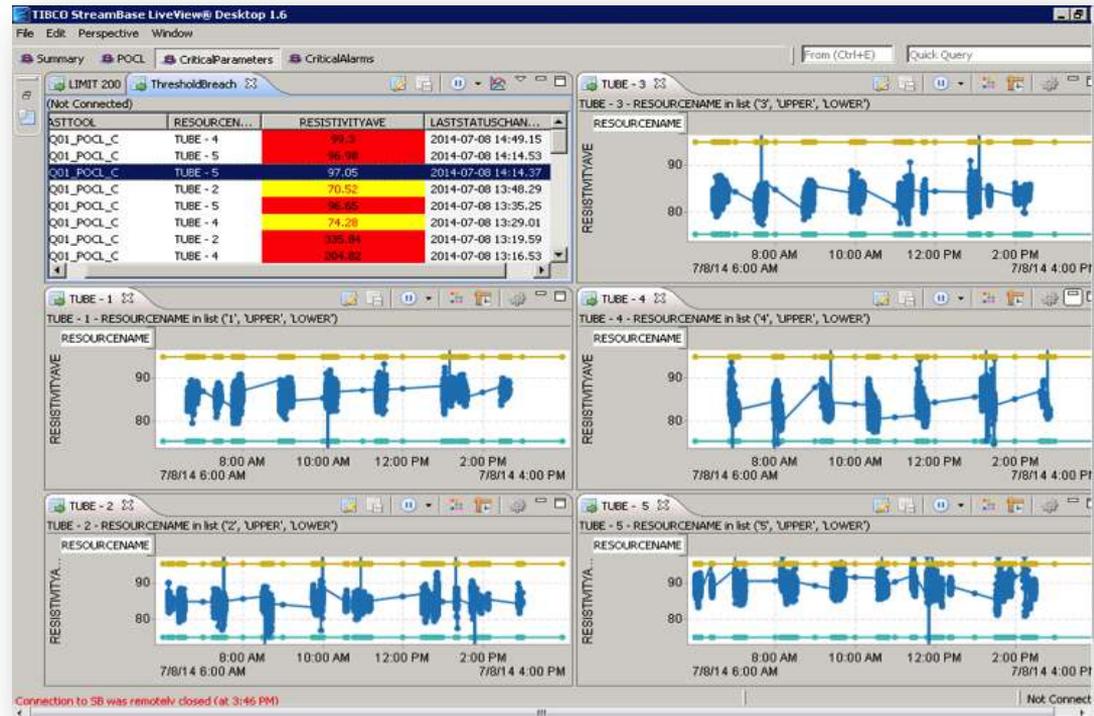
Lesson #5

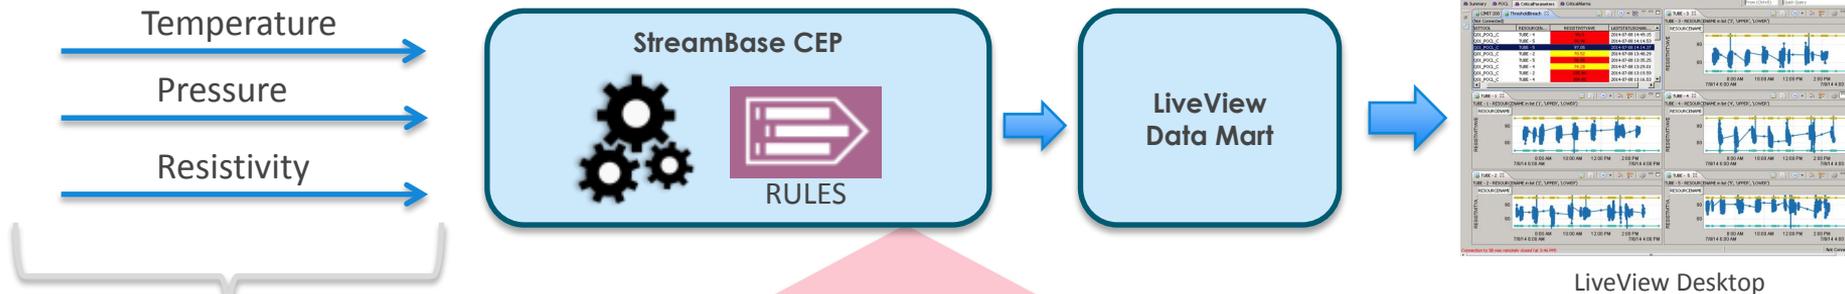
Understand your data analytics pipeline and make the same data available to operational intelligence and algorithmic operations



Manufacturing: Yield Improvement and Cost Reduction

- **About The Customer**
 - Solar Panel Manufacturer
 - One of many manufacturing facilities
- **Manufacturing Yield Management and Optimization**
 - Spotfire & LiveView for manufacturing process yield management and optimization
 - Monitor temperature, moisture reading sensors in real-time with LiveView
 - Real-time predictive detection of manufacturing problems





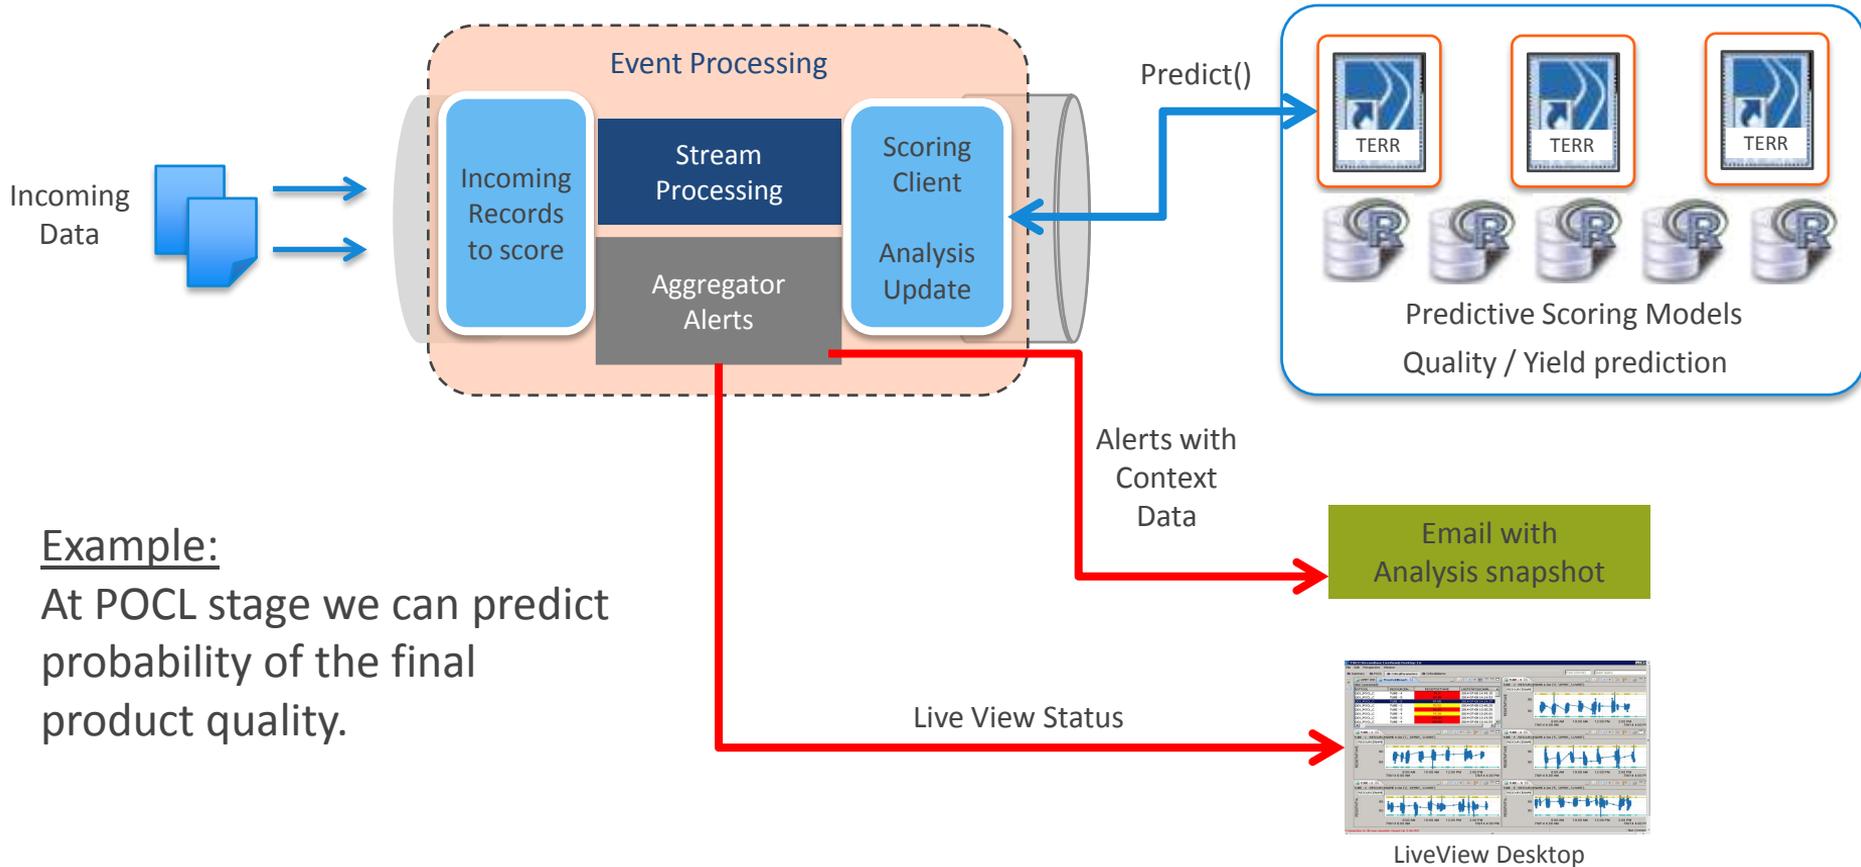
Streaming Data

LiveView Desktop

IF avg (temperature) in the LAST 5 MINUTES >100C

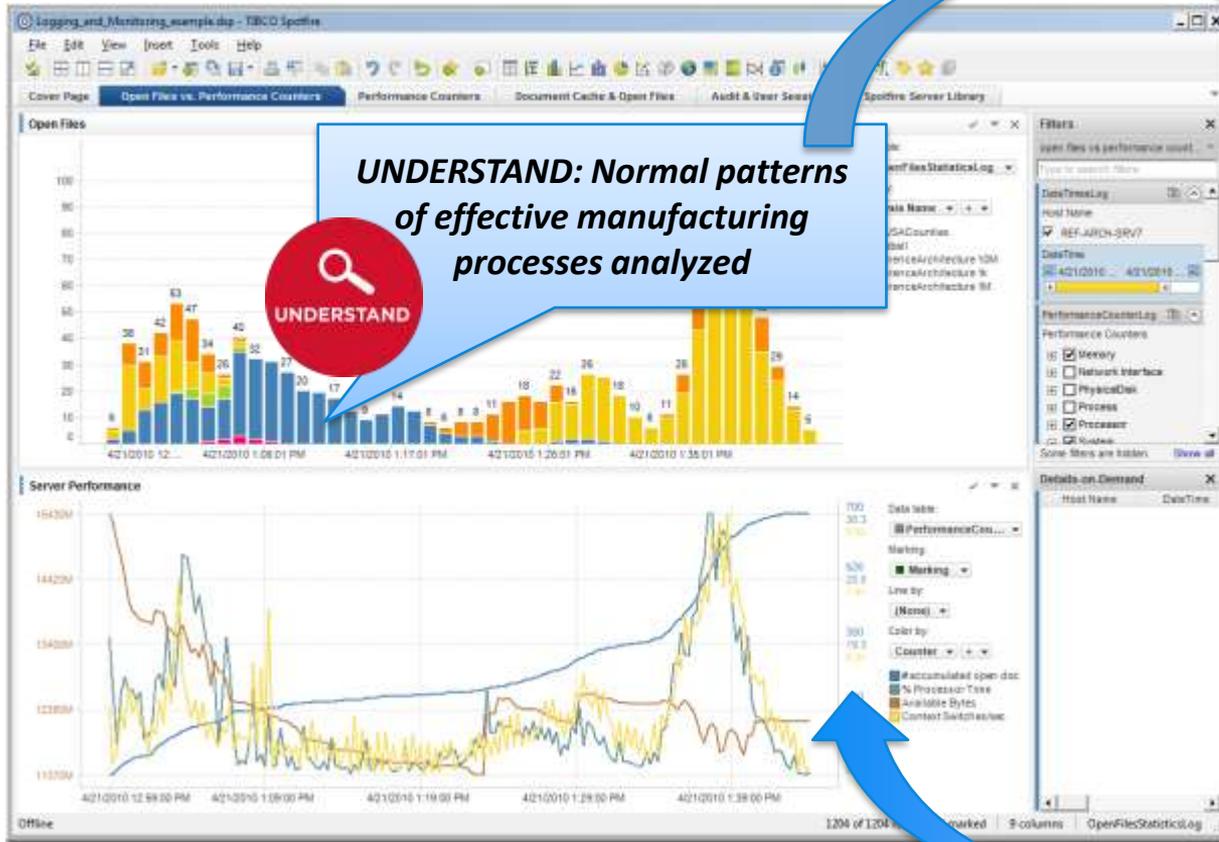
AND avg (pressure) in the LAST 10 MINUTES > 1000 PSI

THEN RAISE ALERT (“DANGER: Maximum Operating Parameters Exceeded”)



Example:

At POCL stage we can predict probability of the final product quality.



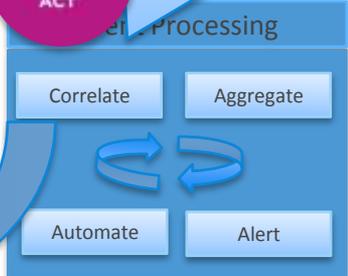
UNDERSTAND: Normal patterns of effective manufacturing processes analyzed



CORRELATE: Real-time factory behavior compared to the model



ACT: Shut down equipment, schedule additional measurements,



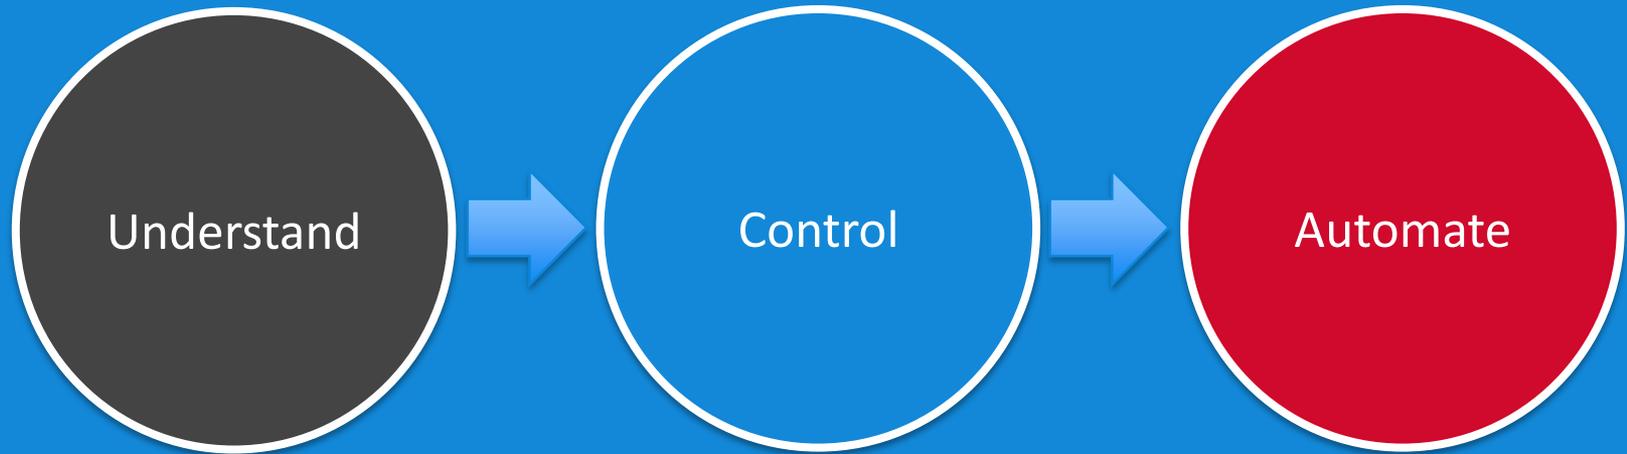
Lesson #6

Follow up on automated problem detection with not only automated action, but also root cause analysis using historical data tools



1. Optimize and automate small decisions first, for quick wins
2. Resolve problems while they are fresh to prevent snowballs
3. Understand your analytics pipeline and have the same data in real time
4. Use event processing to correlate events from multiple systems
5. Empower staff with decision making authority and information, and also provide recommended action
6. Follow up not only with automated action, but root cause analysis

To automate your business with analytics,
you must give the business
visibility and control



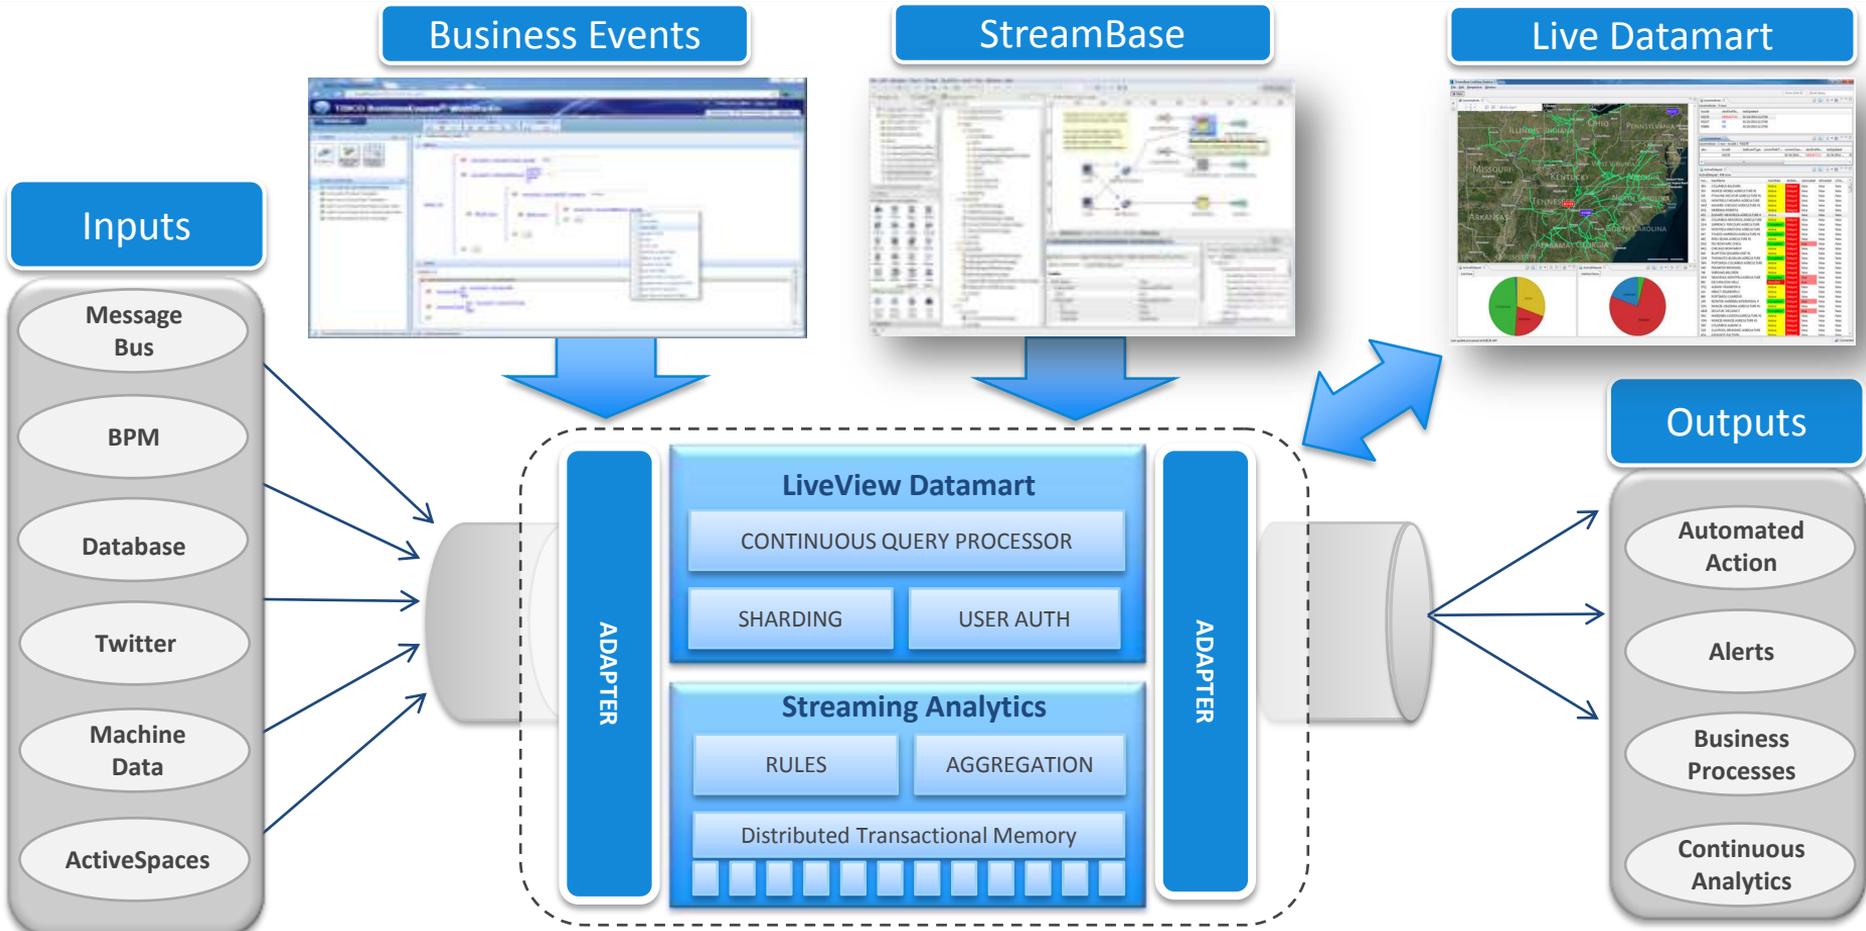


Digital Business: From Big Data to Fast Data

Operational Intelligence and Algorithmic Operations

Lessons Learned from Key Fast Data Use Cases

TIBCO Event Processing



- **Visual Programming**
 - Visual programming
- **Stream & Database Connectivity**
 - 150 adapters, Event-driven database integration
- **Event-Driven Rules**
 - Business Events
- **Streaming Analytics**
 - StreamBase
- **Integrated Analytics**
 - R, TERR, Matlab
- **Big Data Processing**
 - Data grid (ActiveSpaces), distributed transactional memory
- **Big Data Integration**
 - Hadoop (Flume), Amazon Kinesis, Twitter Storm
- **Social Media Integration**
 - Twitter, GNIP
- **Fault Tolerance**
 - Managed clusters, continuous availability
- **Cloud Deployment**
 - TIBCO Silver Fabric
- **Event-Driven BPM**
 - With AMX-BPM
- **Simulation and Backtesting**
 - Replay / testing
- **Visual Debugging**
 - Trace and break-point debugger
- **Enterprise Testing**
 - Test recording, Junit, Fitness, Continuous Integration
- **Developer Community**
 - StreamBase Component Exchange, certification
- **Language Integration**
 - Java, .NET, HTML5
- **Management Tools**
 - GUI, command line, API, JMX

“Business value doesn’t come from building faster applications; it comes from *building applications faster.*”

Event Processing Connectivity 2014: 150+ Streaming Source & Sinks

- **Enterprise Messaging**
 - TIBCO Rendezvous
 - TIBCO EMS
 - TIBCO FTL
 - IBM MQSeries
 - Informatica 29West
 - Solace
 - Tervela
 - JMS (publish / subscribe)
 - TCP/IP
- **Big Data**
 - Hadoop (Flume)
 - Amazon Kinesis
 - Twitter Storm
 - **Spark**
- **Data Grid / Cache**
 - ActiveSpaces
 - **TIBCO Distributed Transactional Memory**
 - **Mongo**
- **Integration**
 - **BusinessWorks**
- **Internet of Things**
 - **OSI Pi**
 - **MQTT (MQ Telemetry Transport)**
 - **Insteon**
 - **X10**
 - **XBee,**
 - **Google Glass**
 - **Live cameras for face detection**
 - **Beacons**
 - **SDR (software defined radio) for plane tracking**
 - **AR Drone quadcopter**
 - **OBD-II car on-board computer diagnostic adapter**
 - **Xbox**
- **Social**
 - Twitter
 - GNIP
 - TIBBR
 - Jabber / IR
 - RSS
- **Database**
 - Vertica
 - Oracle
 - SQL Server
 - MySQL
 - IBM DB2
 - Sybase ASE / IQ / RAP
 - Active Oracle log reader
 - Active SQL Server log reader
- **Technology**
 - Binary file (read and write)
 - CSV files and sockets (read and write)
 - Email (read and writer)
 - IP packet capture (HTTP / IRC / POP3 / SMTP)
 - Log files
 - Regular expression file and socket

150+ adapters to Streaming,
Social, Hadoop , IoT,
Integration, Technology data
sources

“The LiveView Datamart is like a data warehouse for data that’s constantly in motion, and changing all the time.”

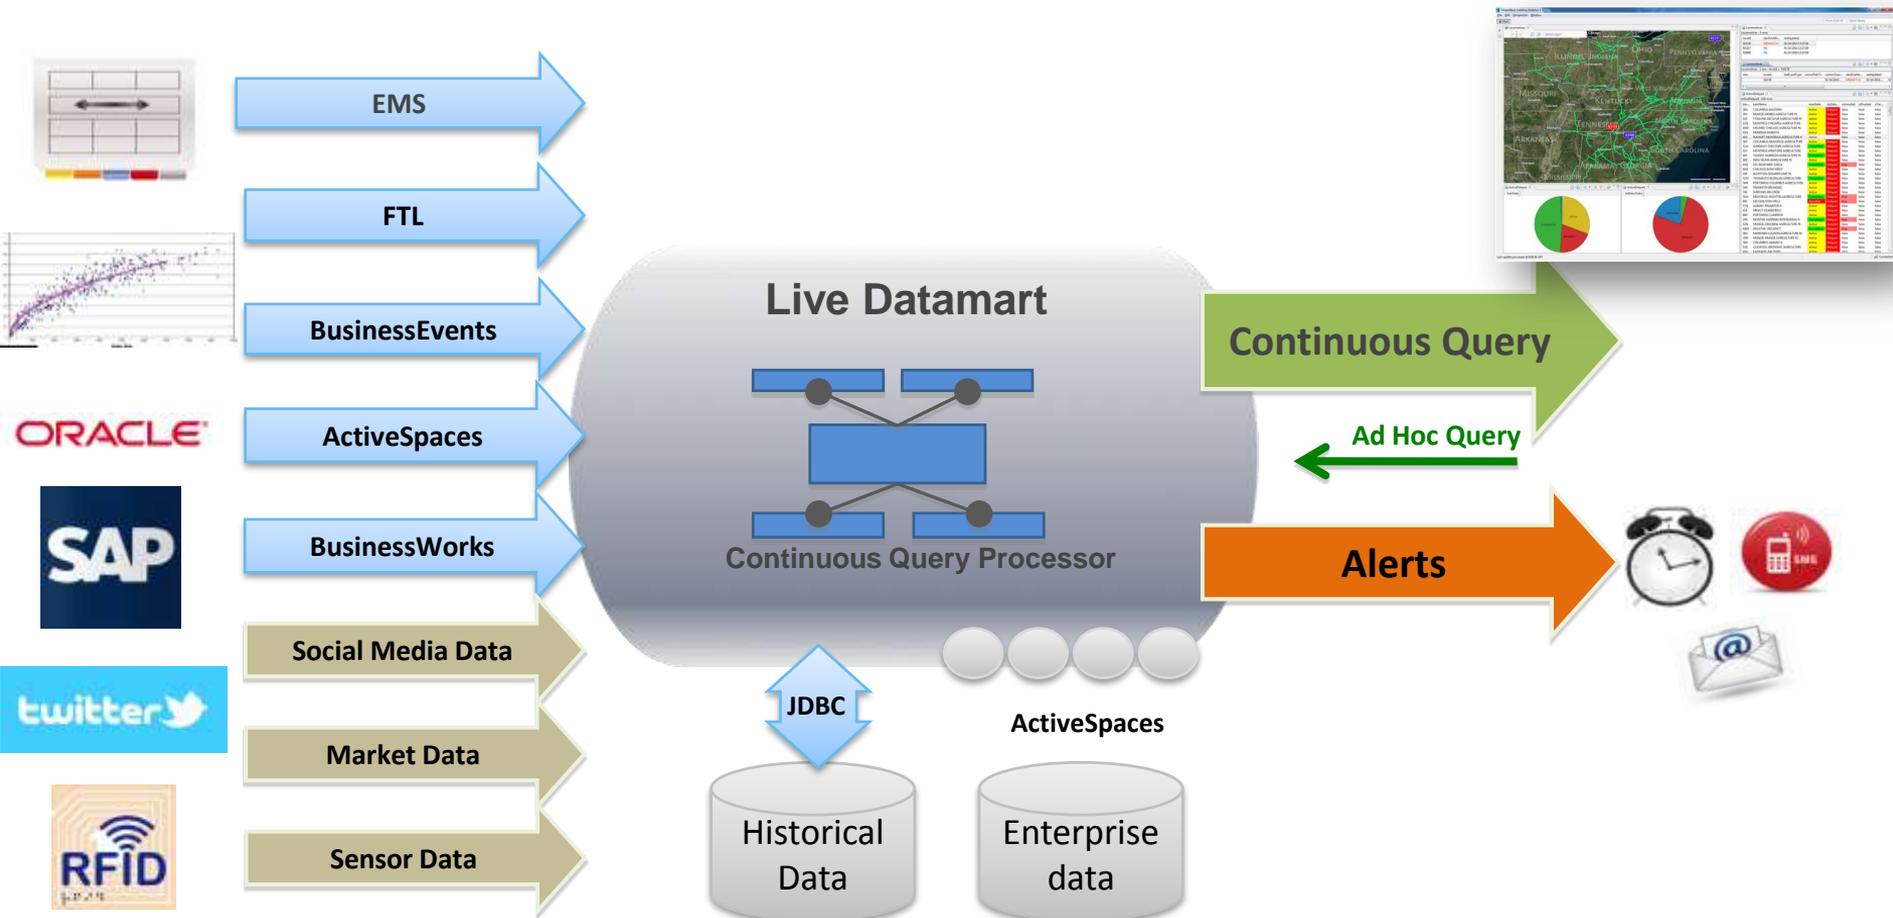
THE END OF THE END OF DAY REPORT

The screenshot displays the IngestBase LiveView Desktop 3.3 interface, which is a live data monitoring dashboard. It features several key components:

- Drilling Operations Center:** A map of a city area with various colored markers (red, blue, green) indicating operational points.
- Fraud Surveillance:** A table listing addresses and associated data points, such as '2000 24th Street #200' and '1806 Market Street'.
- Trade Surveillance:** A table listing addresses and associated data points, such as '1806 Market Street' and '1-01 Tandy Place'.
- Network Operations Center:** A map showing network connections and data flow between various nodes.
- CMO (Continuous Monitoring):** A dashboard showing real-time data and trends.
- Logistics Operations Center:** A dashboard showing logistics data and trends.
- Continuous Compliance Monitoring:** A dashboard showing compliance data and trends.

The interface also includes a table of data points on the right side, with columns for 'City', 'State', and 'ZIP'. The table lists several entries for San Francisco, CA, with ZIP codes ranging from 94120 to 94124.

City	State	ZIP
San Francisco	CA	94120
San Francisco	CA	94123
San Francisco	CA	94109
San Francisco	CA	94107
San Francisco	CA	94113
San Francisco	CA	94122
San Francisco	CA	94103
San Francisco	CA	94102
San Francisco	CA	94104
San Francisco	CA	94105
San Francisco	CA	94108
San Francisco	CA	94111
San Francisco	CA	94105
San Francisco	CA	94103
San Francisco	CA	94103
San Francisco	CA	94107
San Francisco	CA	94103
San Francisco	CA	94109
San Francisco	CA	94104



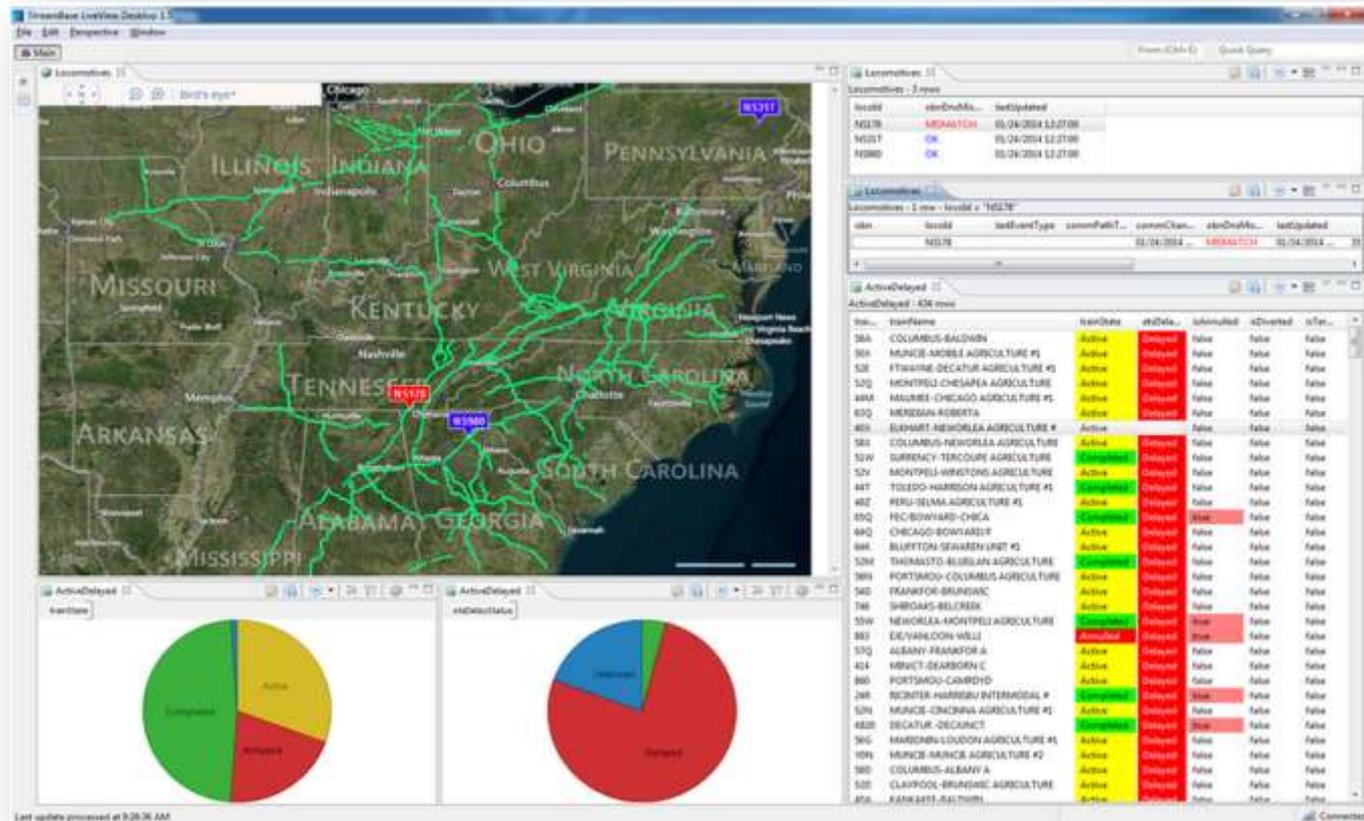
TIBC™ | Fast Geo-Spatial Data

The screenshot displays the Streambase Desktop 1.5 interface. The main window shows a map of San Francisco with various taxi data points and routes overlaid. The interface includes several panels:

- Task 11:** A table titled "Taxi - 32 rows - where rowid < 300" with columns: total_passenger, in_state_id, address, city, state, zip. The data is color-coded by status: green for "Deliver", red for "Deliver", and blue for "Pickup".
- Task 12:** A table titled "Total Pickups - where rowid = 100" with a "label" column. It shows a horizontal bar chart with six bars of different colors (blue, green, grey, teal, light blue, dark grey).
- Task 13:** A table titled "Has Passengers - where rowid < 100" with a "has_passenger" column. It shows a pie chart with a large green section and a smaller red section.
- Task 14:** A video player showing a street-level view of a road.

total_passenger	in_state_id	address	city	state	zip
		2570 Geary Boulevard	San Francisco	CA	94122
		1900-1908 El Camino Real	San Francisco	CA	94123
		1001-1008 Franklin Street	San Francisco	CA	94109
		85-88 Townsend Street	San Francisco	CA	94107
		1051-1059 Divisadero Street	San Francisco	CA	94115
		1823 Union Street	San Francisco	CA	94123
	Deliver	2300 18th Street - 4203	San Francisco	CA	94103
	Pickup	1680 Market Street	San Francisco	CA	94102
	Pickup	1-99 Trinity Place	San Francisco	CA	94104
		276-290 Fulton Street	San Francisco	CA	94107
		2480-2488 Larkin Street	San Francisco	CA	94109
		403-444 Battery Street	San Francisco	CA	94111
	Pickup	101 3rd Street	San Francisco	CA	94105
		3256-3306 10th Street	San Francisco	CA	94103
		3264-3284 Cross Street	San Francisco	CA	94133
		344-348 9th Street	San Francisco	CA	94103
	Deliver	218 Monterey Boulevard	San Francisco	CA	94127
		178-178 The Embarcadero	San Francisco	CA	94105
		1063-1090 Hyde Street	San Francisco	CA	94109
	Deliver	55-89 Sutter Street	San Francisco	CA	94104

- **Open Client Support**
 - For custom GUI development
 - Allows open partner development
 - Ad-hoc query
 - Notifications
- **HTML5**
 - D3, jQuery, ExtJS, Google Charts, Bing, AngularJS
- **.NET**
 - For custom .NET development
- **Java**
 - For custom Java GUI development



Operational Intelligence

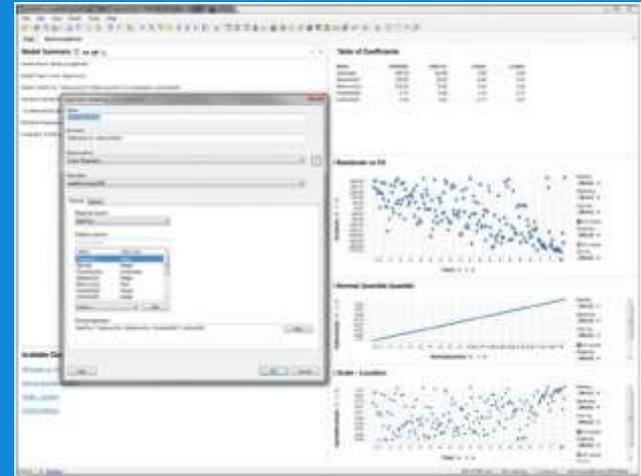
Human decisions in real time
informed by up to date information



The Challenge: The right
information, ready to be used by
the right people at the right time.

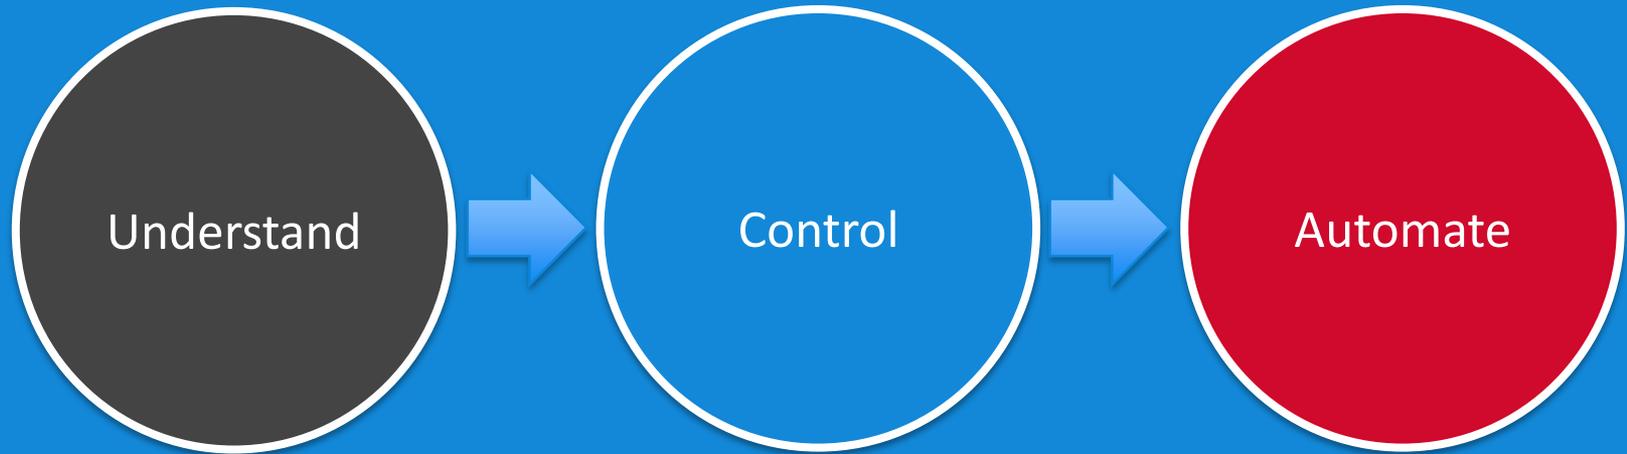
Algorithmic Operations

Automated action based on analytic
models of history combined with
live context and business rules



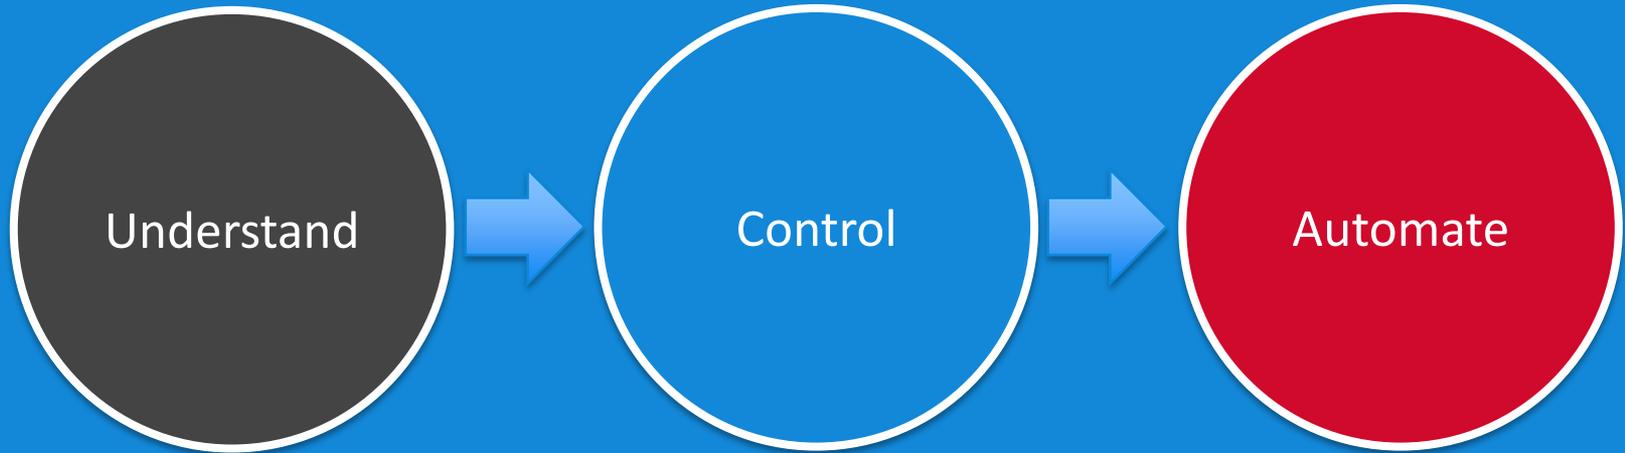
The Challenge: Empowering business
stakeholders to understand and
control the selection of models

To automate your business with analytics,
you must give the business
visibility and control



Questions?

How will you automate with analytics?
How will you give visibility and control?



More info at tibco.com
Or at our booth here

Follow me @tibbetts
Follow up with tibbetts@tibco.com

