Large-Scale Distributed Stream Processing with IBM InfoSphere Streams

INNOVATE. OPTIMIZE. PERFORM.

Salishan Lodge, Gleneden Beach, OR, April 29, 2010

Henrique Andrade, Research Staff Member
Runtime Co-Architect, InfoSphere Streams
IBM Research

Something Meaningful is Happening…

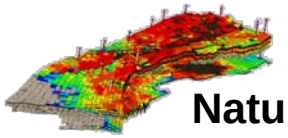

Natural Systems
- Seismic monitoring
- Wildfire management
- Water management

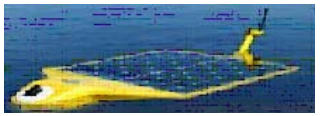

Transportation
- Intelligent traffic management

Manufacturing
- Process control for microchip fabrication

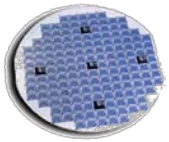

Stock market
- Impact of weather on securities prices
- Analyze market data at ultra-low latencies

Law Enforcement
- Real-time multimodal surveillance

Fraud prevention
- Detecting multi-party fraud
- Real time fraud prevention

Radio Astronomy
- Detection of transient events

Health & Life Sciences
- Neonatal ICU monitoring
- Epidemic early warning system
- Remote healthcare monitoring

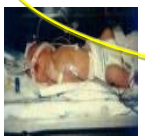

Telecom
- Processing of Call Detail records
- Real-time services, billing, advertizing
- Business intelligence
- Churn Analysis, Fraud Detection

Challenges for New Real-time Intelligence

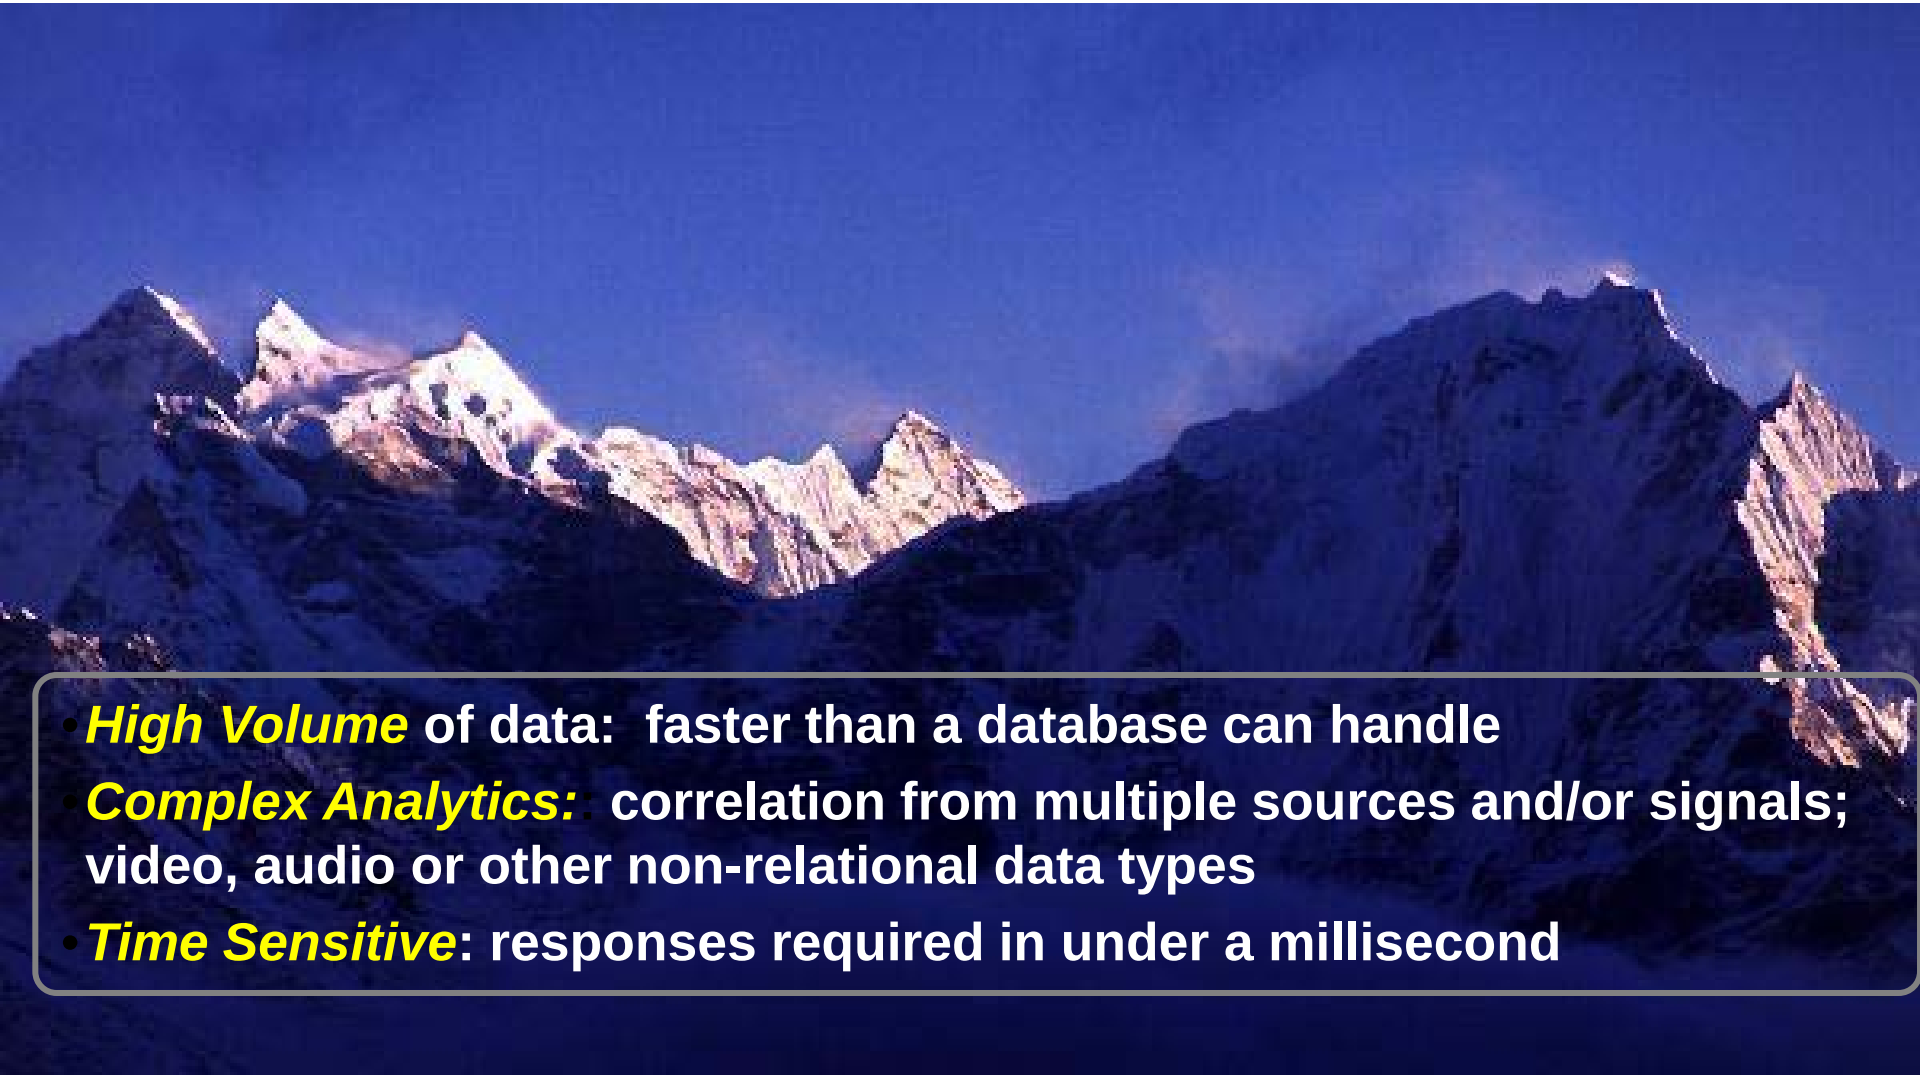

- ***High Volume*** **of data: faster than a database can handle**
- ***Complex Analytics:*** **correlation from multiple sources and/or signals; video, audio or other non-relational data types**
- ***Time Sensitive*****: responses required in under a millisecond**

InfoSphere Streams Value Proposition

"What is <u>happening now</u> that requires <u>continuous analysis and rapid response</u>?"

Enable real-time analytic applications

Continuous ingestion

Continuous analysis

Continuous results

Extreme event volume

Millions of events per second

Extreme speeds

<< 1 ms latency for "true-real time" response for some workloads

Extreme analysis

Simple and complex data types processed with simple filters and rules up through arbitrarily complex processing

Extreme flexibility

New analysis can be deployed on-demand

Stream computing represents a critical paradigm shift

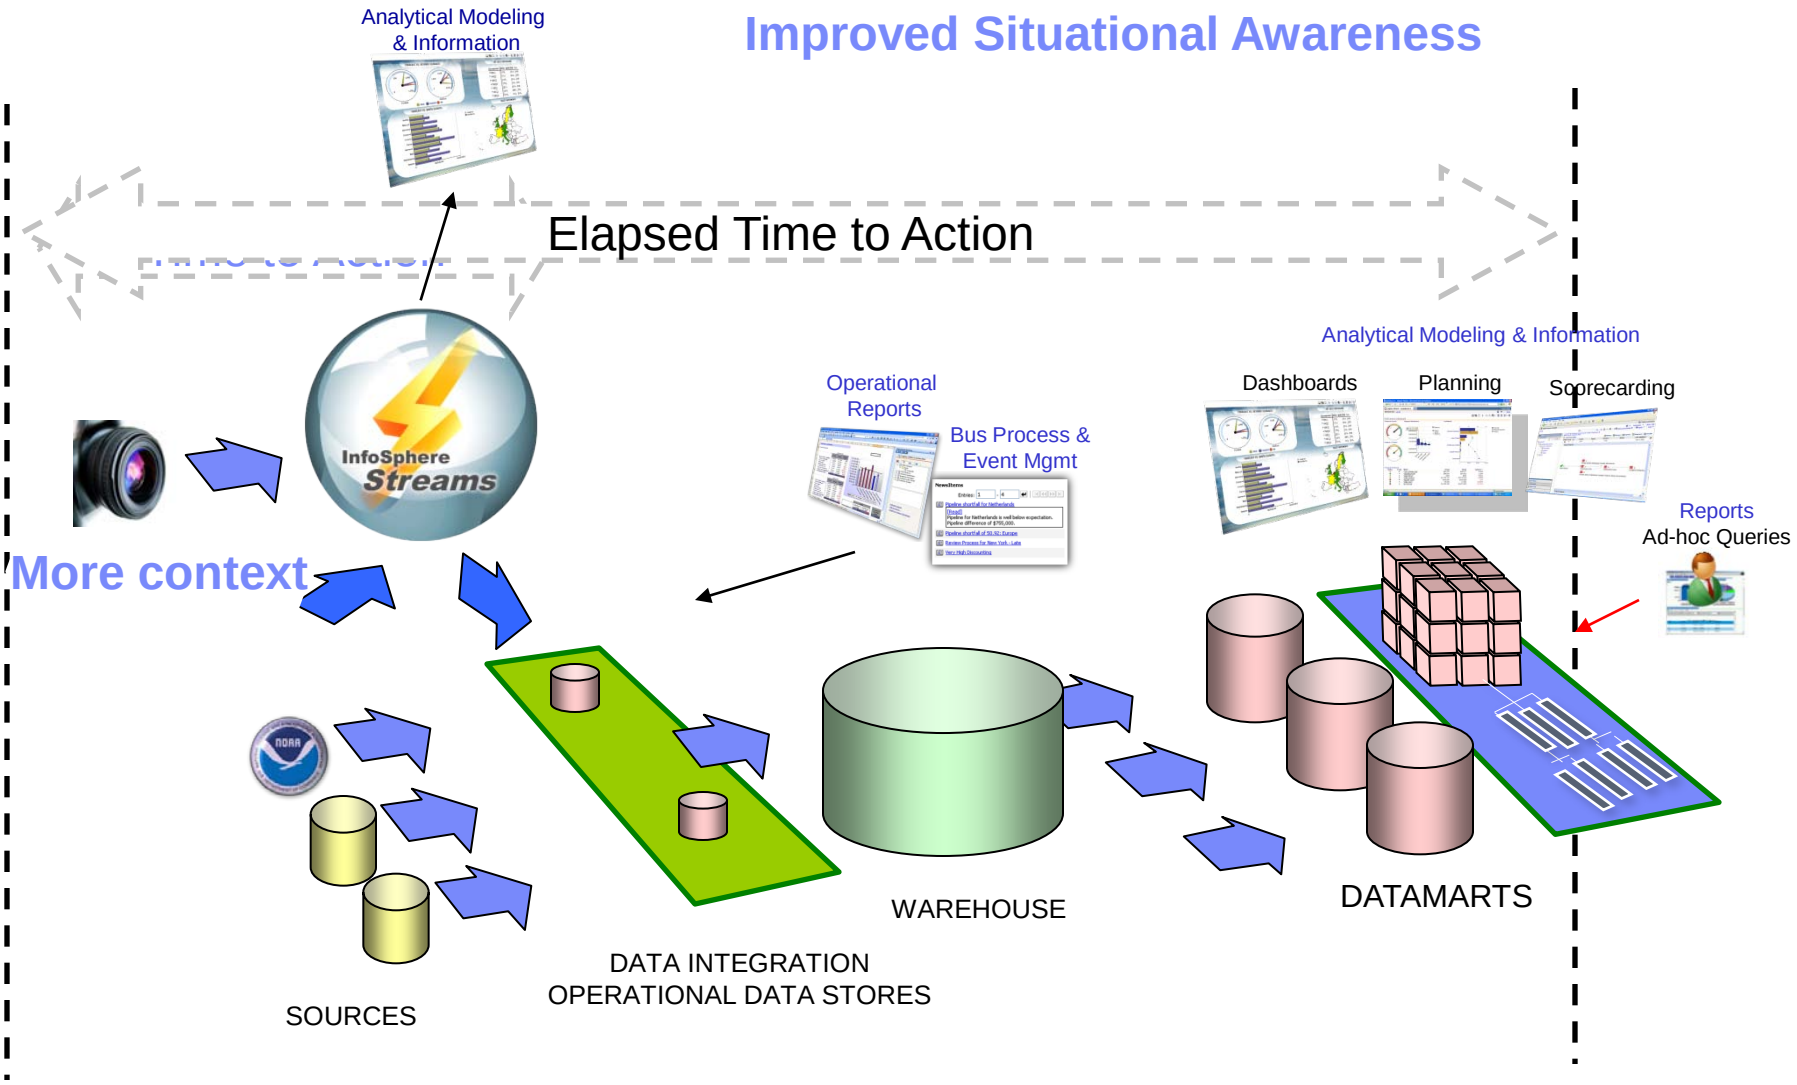

Stream Computing Illustrated

Continuous Ingestion → **Continuous Complex Analysis in Microseconds**

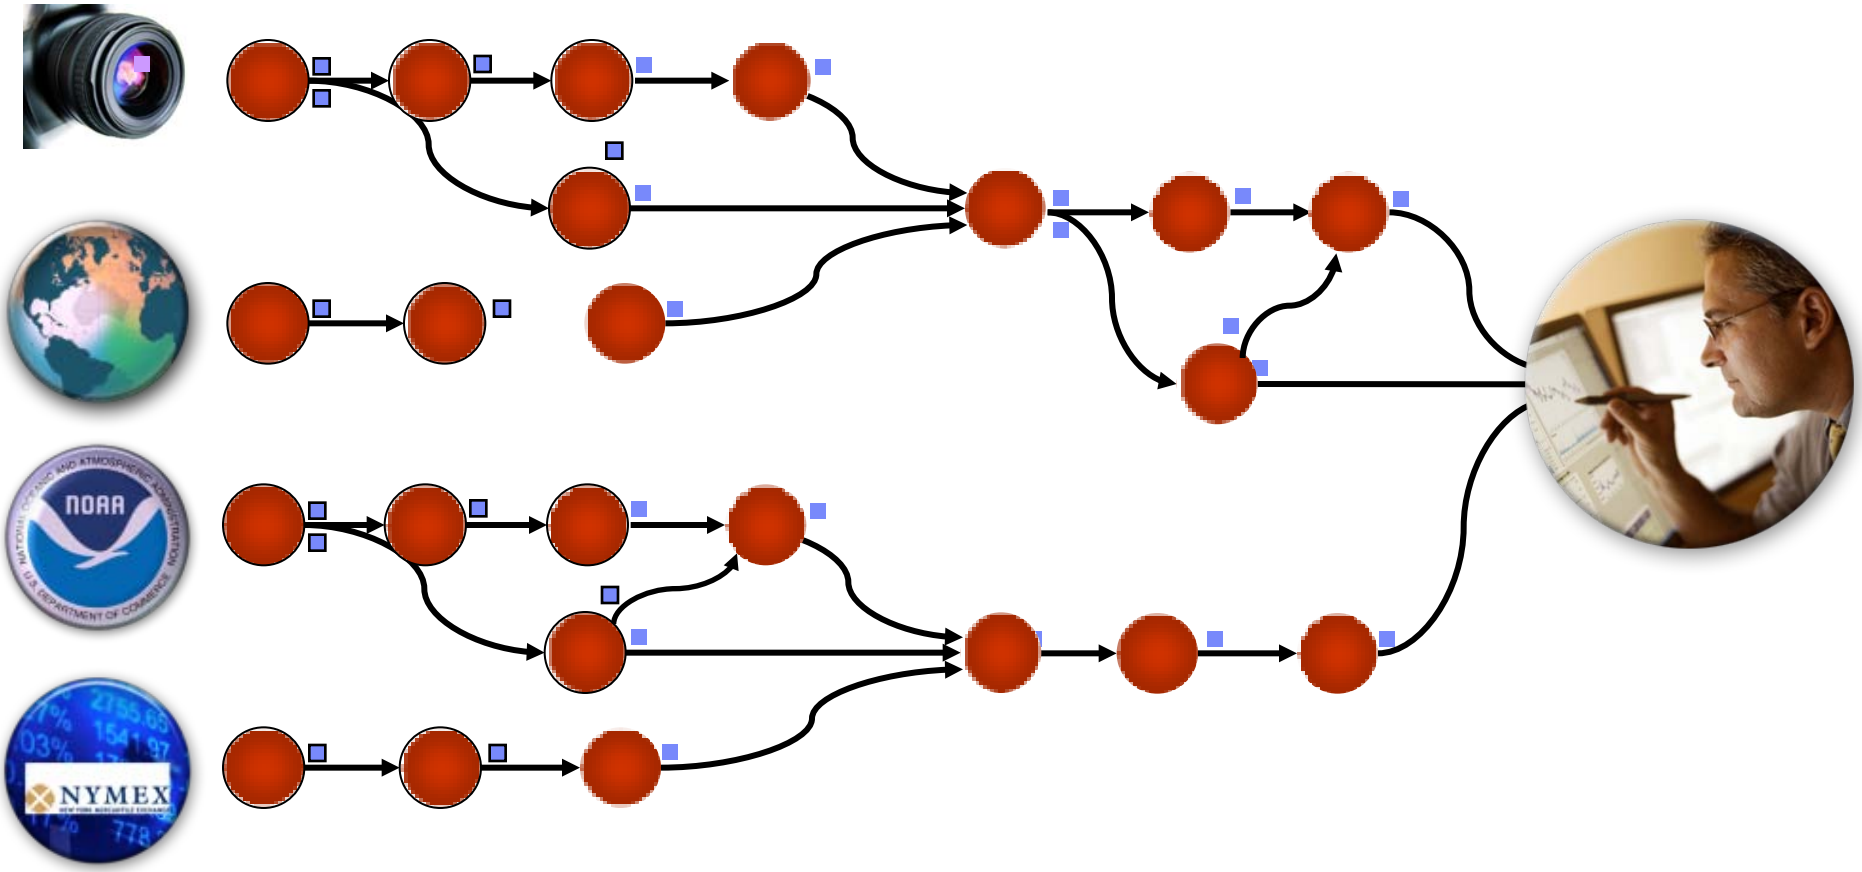

Stream computing handles extreme volumes and diverse information and events

Large spectrum of events/data

Structured data ***Unstructured data***

-High usefulness density
-Simple analytics
-Well defined event
-High speed (million events per sec)
-Very low latency

-Low usefulness density
-Complex analytics
-Event needs to be detected -High volume (TB/sec)
-Low latency

High Performance in 3 Critical Dimensions:

Throughput, Latency, and Scale

- **Scalability**
  - A single Streams instance can exploit more than 1000 cores
    - 125 servers are supported per instance
- **Baseline throughput (number of messages/second)**
  - 1.0 32-bit = 421,538 msgs/sec
  - 1.0 64-bit = 417,622 msgs/sec
- **Sustained data rate of 1 million messages / second**
  - 6 node system can process at 1.3 million msgs/sec
- **Latency for a two operator application can be much less than 1 millisecond**
  - 1.0 32-bit = 116 microseconds
  - 1.0 64-bit = 107 microseconds
  - *Significant additional improvements with LLM/Infiniband tech preview*

The InfoSphere Streams 1.0 Platform

Development Environment

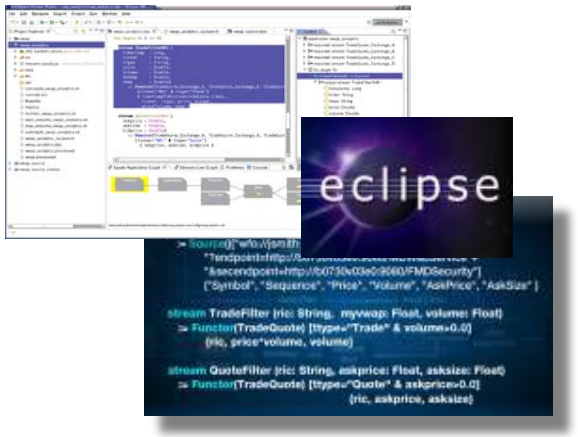

Streams Studio
Eclipse IDE for SPADE

Runtime Environment

Scalable stream processing runtime

Tools and Technology Integration

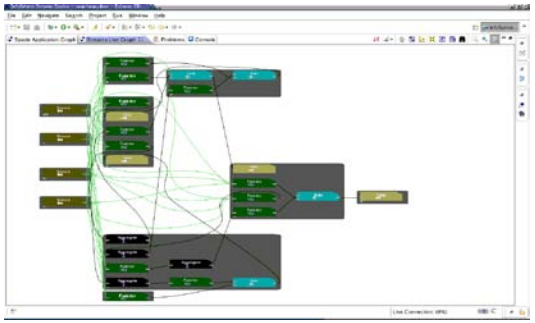

WebSphere
Front Office 3.0

Streamsight,
Built-in Stream Relational Analytics,
Adapters

Streams Studio Integrated Development Environment

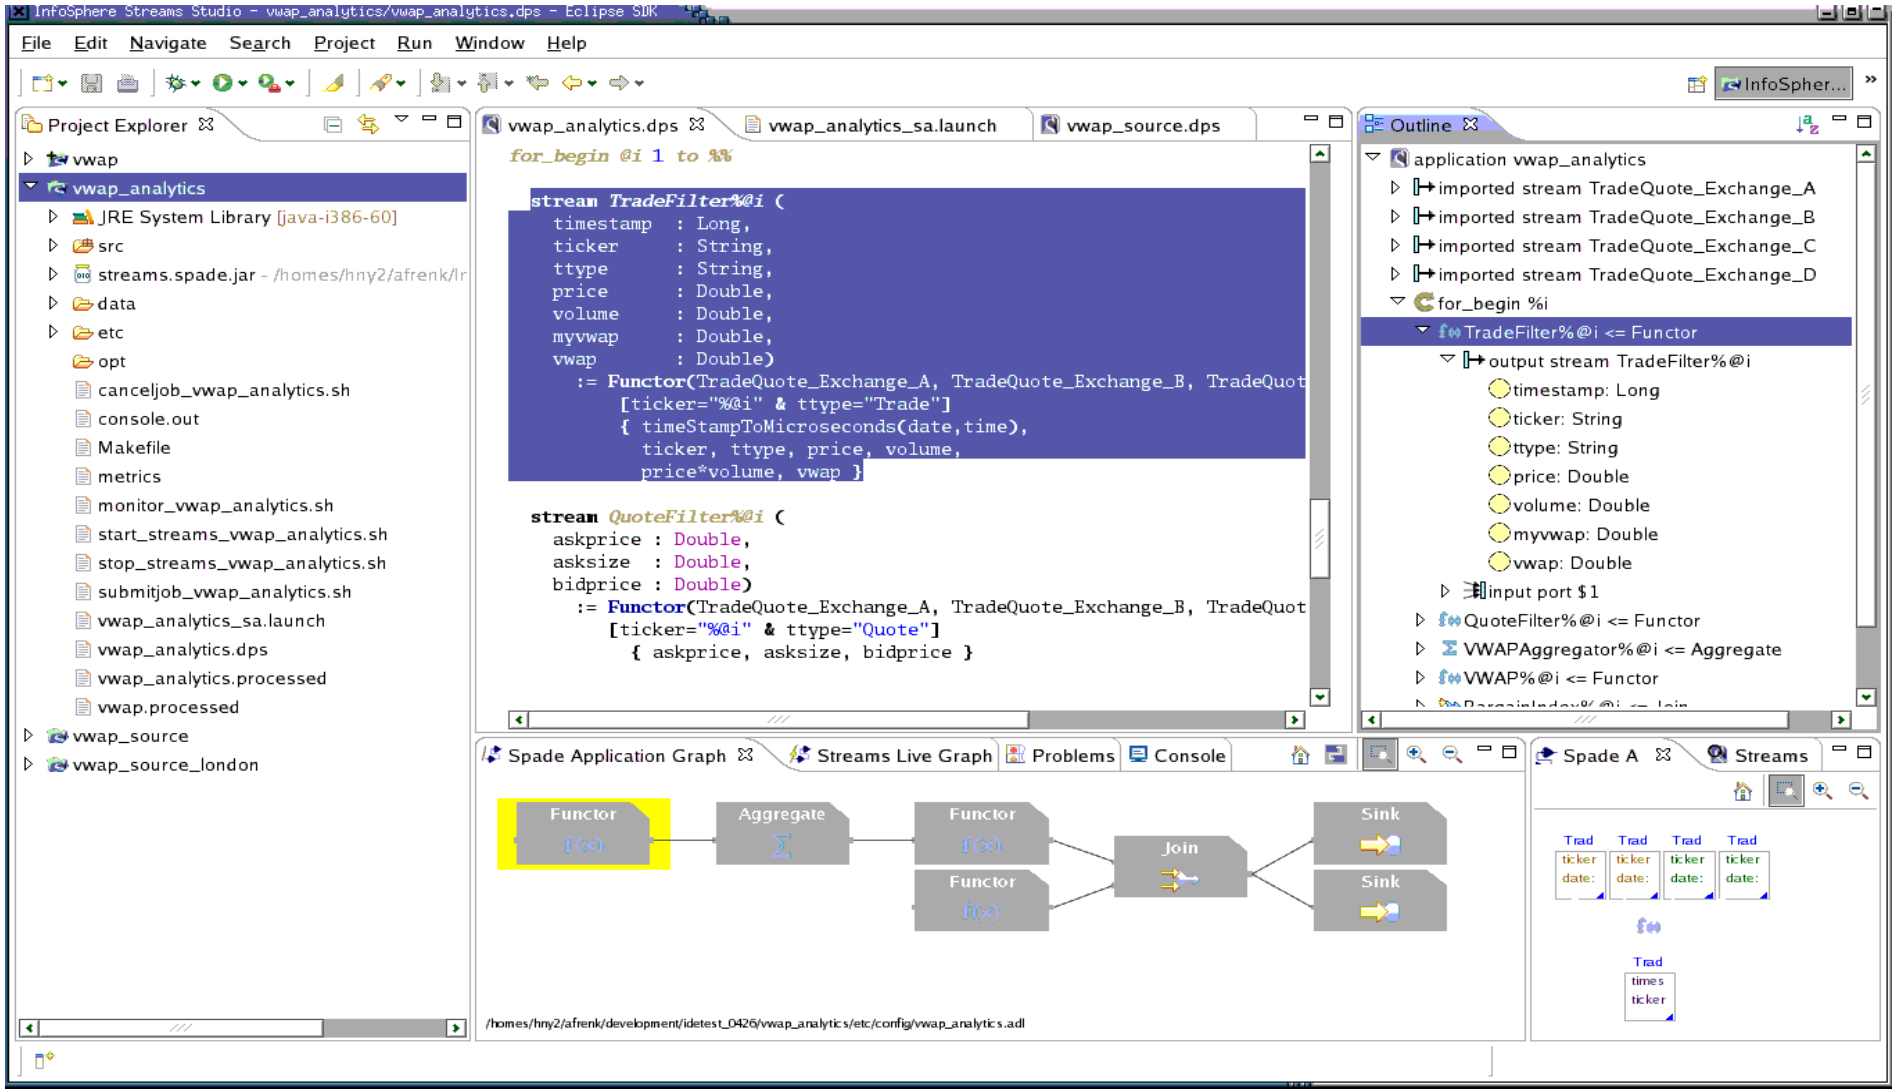

Monitor applications, track and debug data flow

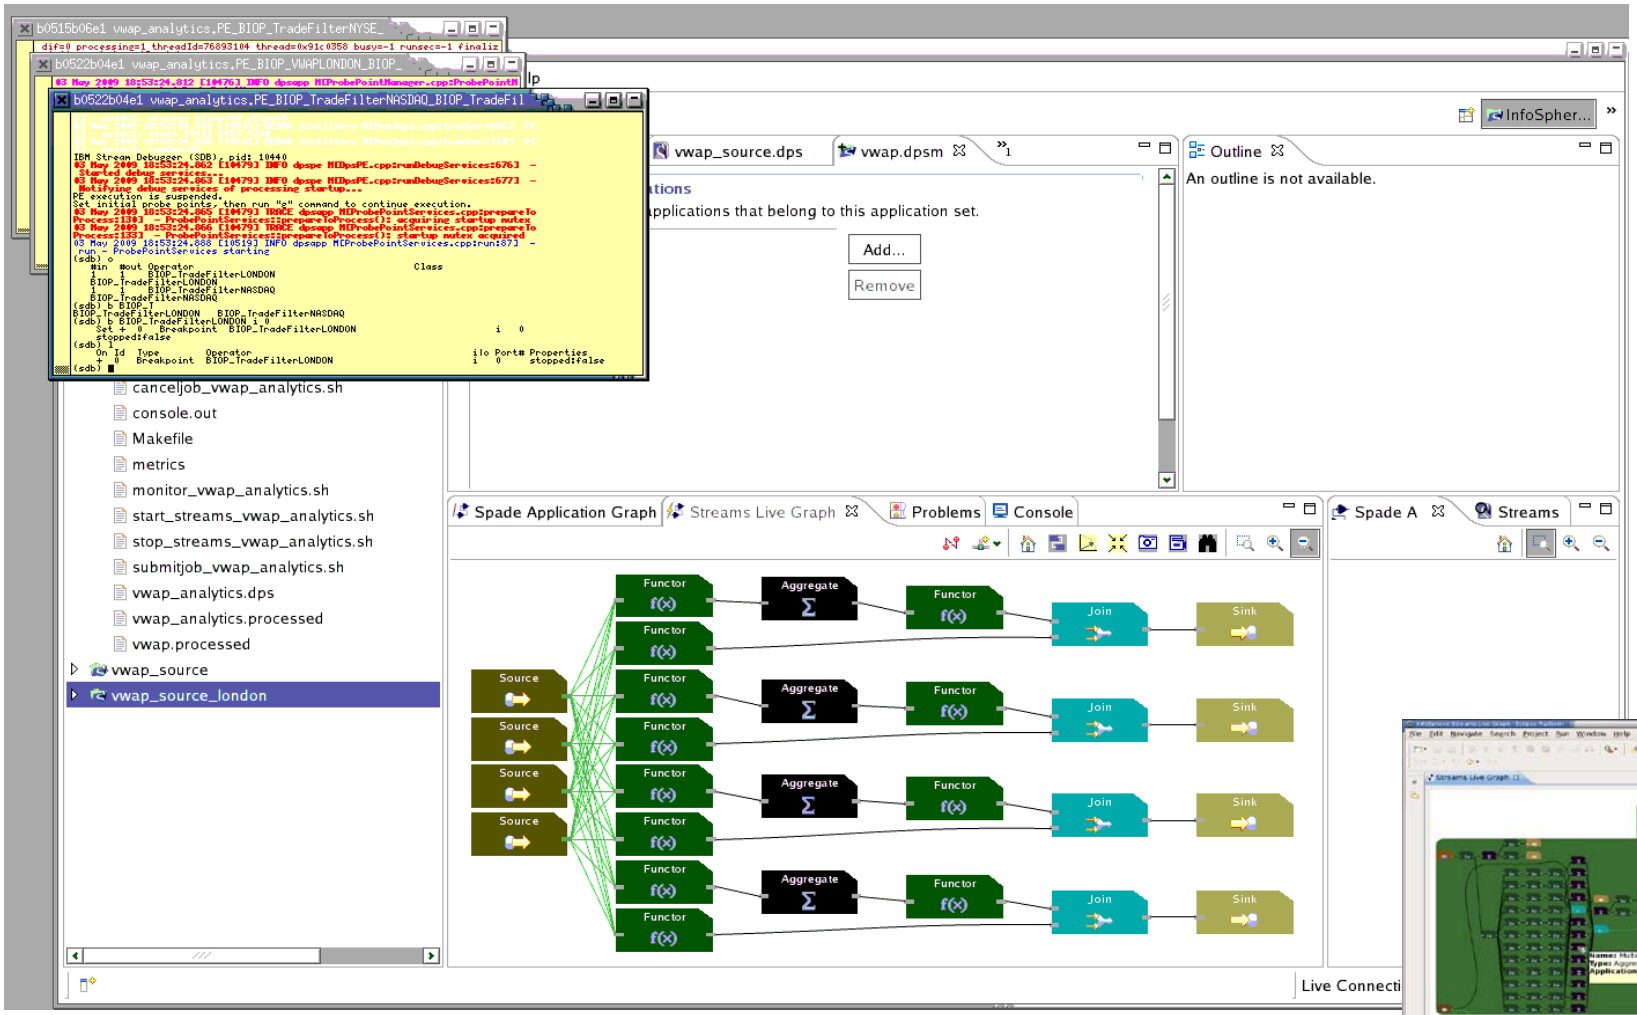

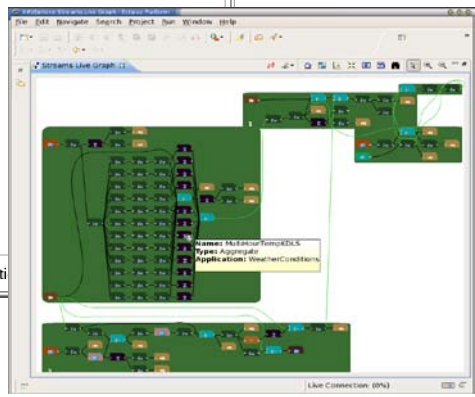

Administration – Application Management

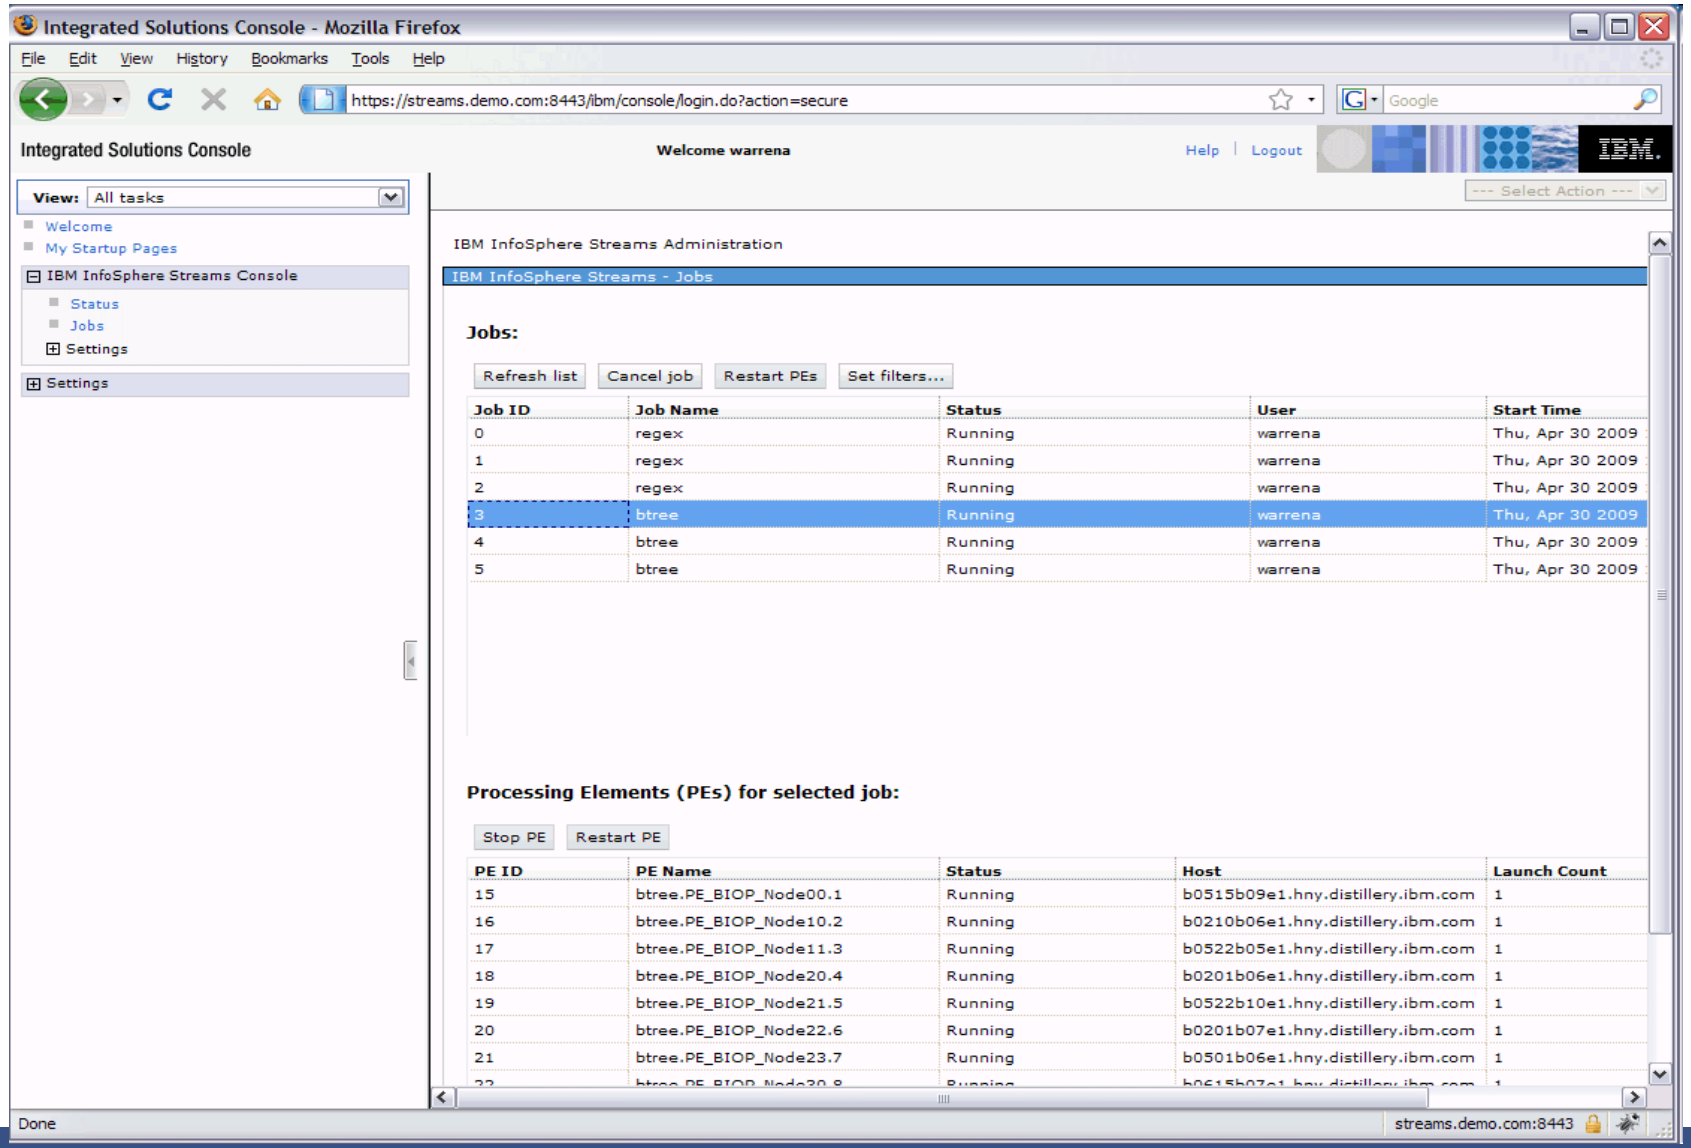

Over fifty samples for faster time to value

Simple samples:

- Regex
- Tumbling window
- Aggregator
- RSS feeds
- VWAP
- AlgoTrading

Hypothetical Commodity Purchasing Application:

- Your company's business relies upon a critical commodity.
- There are multiple suppliers of this commodity.
- Your goal is to always buy from the optimal supplier.

Considerations:

- Quality/condition of the commodity
- Minimization of risk

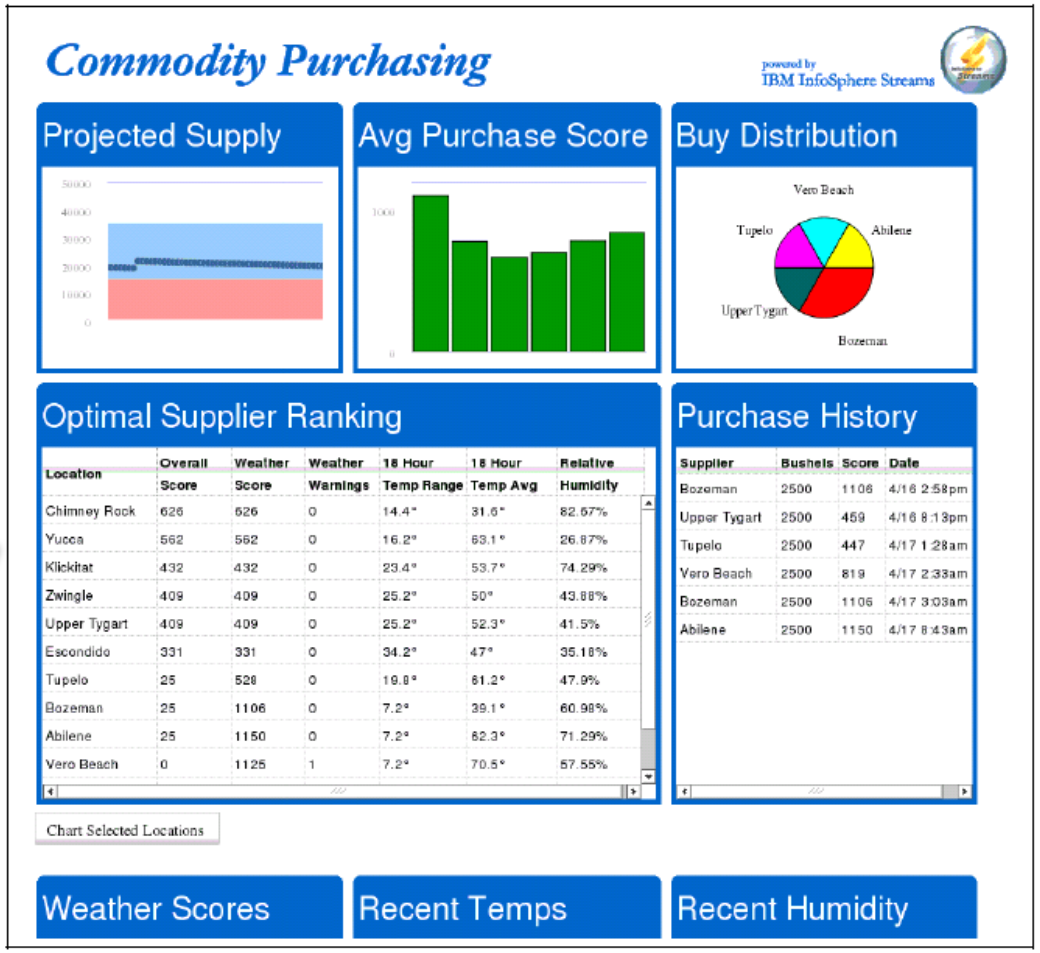

Streams Mining Toolkit
(Currently in Beta)

- **Streams Mining Toolkit enables real-time predictions and discoveries as data arrives**
  - Standard representation for models (PMML)
    - PMML models can be generated by InfoSphere Warehouse
    - Also generated by other compliant tools (e.g., SAS, SPSS)
  - Variety of techniques
    - Supervised: data with known results is analyzed to build a profile, or model to predict results for new data
    - Unsupervised: data is analyzed to discover interesting similarities or groupings
- **Predictions can be made seconds, hours, days, … before actual results can be known**

Streams Mining Toolkit

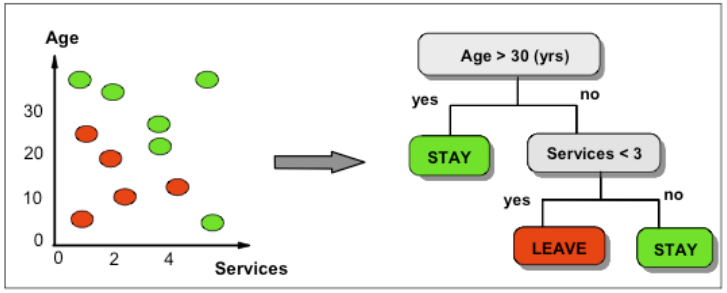

- **Classification**
  - Predicts whether a record belongs to a certain class
    - Which type of vehicle part is most likely to fail?
    - Is this employee likely to leave?
  - Algorithms: Decision Trees, Naïve Bayes

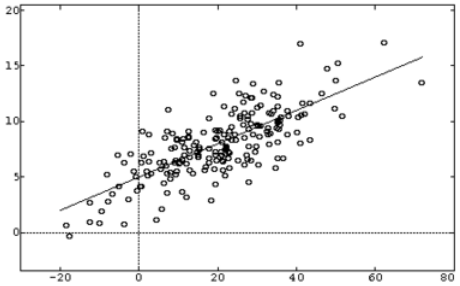

- **Regression**
  - Predicts the quantity or probability of an outcome
    - What is the likelihood of heart attack, given age, weight, …?
    - What is the expected profit a customer will generate?
    - What is the forecasted price of a stock?
  - Algorithms: Logistic, Linear, Polynomial, Transform

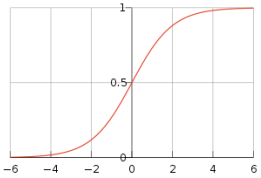

Streams Mining Toolkit

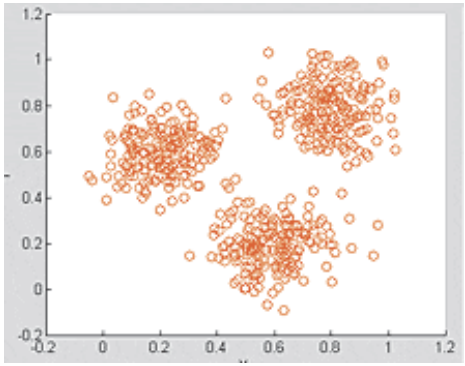

- **Clustering**
  - Identifies groups with common characteristics, or properties of similar groups
    - What are behavior-based properties of various types of servers (e.g., database, application, …)
    - Which healthcare providers may be submitting fraudulent claims?
  - Algorithms: Demographic, Kohonen
- **Incremental Learning**
  - Learns model incrementally, as data arrives
    - Is the data being received drifting from the model?
    - Should I use a model based on more recent events?
  - Algorithms: Incremental decision tree learner, …

Stream Computing in Telecom Sector

Customer Application Requirements
Revenue Assurance, Business Intelligence, Fraud Management

- ***Real-time Call Detail Record Processing***
  - *Mediation, standardization*
  - *Incremental KPI computation*
  - *Preprocessing for BI applications*
- ***Cost-Effective Analytic Solutions***
  - *Churn Prediction and Prevention*
- ***Real-time Complex Analytics***
  - *Context driven multimodal advertizing*
  - *Fraud detection and prevention*
- ***Real-time Dashboarding***
  - *Complex KPI monitoring*

Integrated Solution
From data-in-motion to data-at-rest, … and back

- ***IBM InfoSphere Streams***
- ***IBM InfoSphere Warehouse***
- ***IBM Cognos Now!***
- ***IBM Research: SNAzzy***

Benefits

- ***Proactive business intelligence***
- ***Excellent scaling***
- ***Cost-Effective Solution***
- ***Significantly reduced time to react***

Solution Architecture

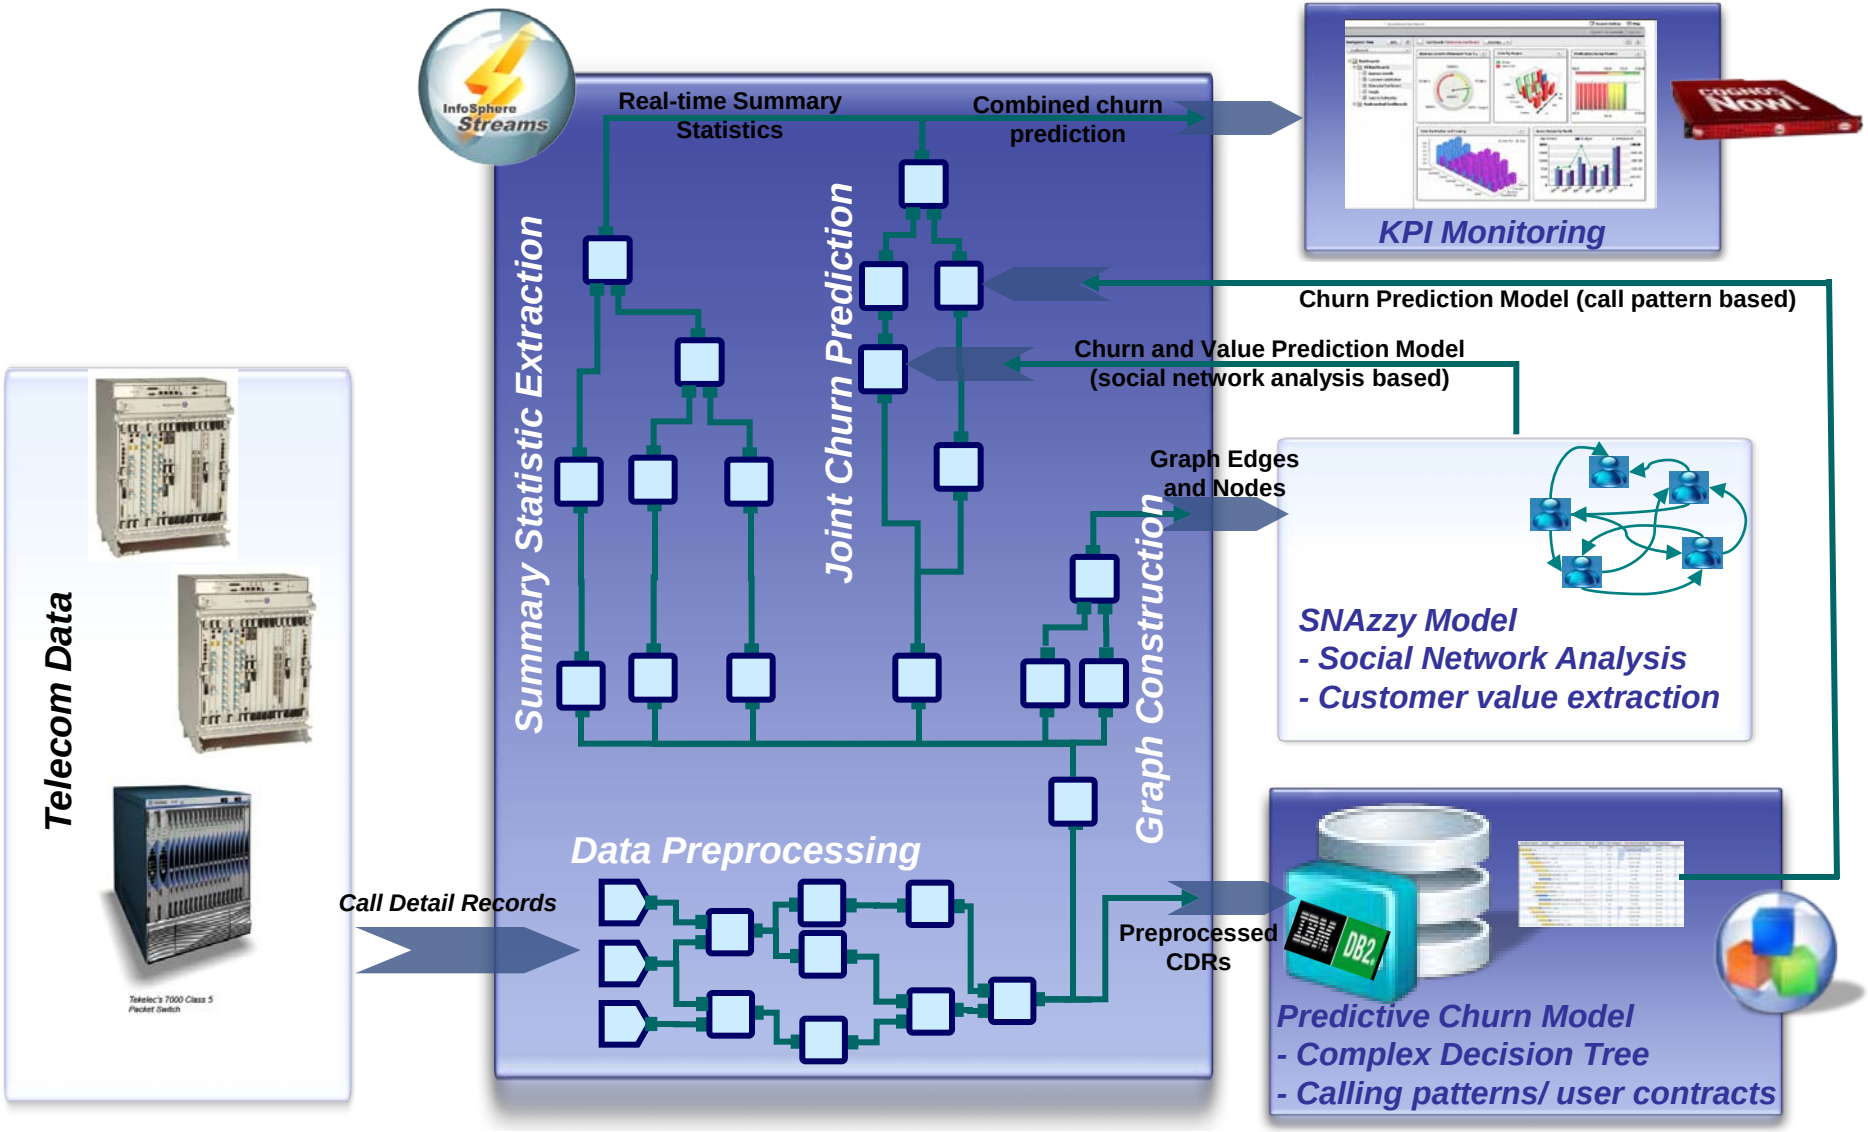

Streams Spatial Extender Toolkit (In Progress, Research)

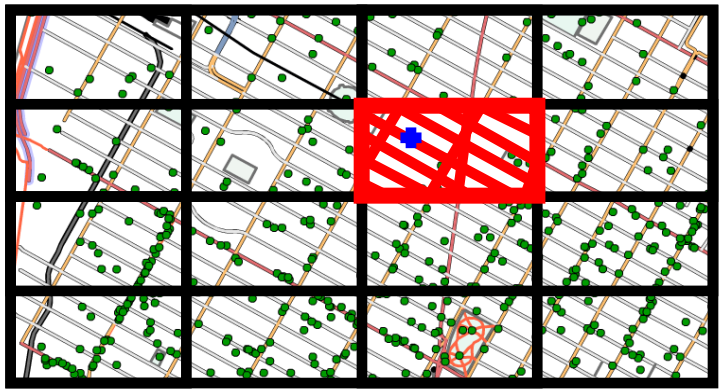

- **Brings geospatial information processing capabilities to InfoSphere Streams**
- **Functions and operators to analyze information and relationships about geographic features**
  - Computing distances between shapes (points and polygons)
  - Shape containment
  - Finding closest shapes for a given location.
- **Computations keep the data in the unprojected form**
  - i.e., assumes figure of the Earth is an oblate spheroid

Intelligent Transportation

- **Multimodal Data Streams**
  - GPS
  - Counts, speeds, travel times
  - Public Transport
  - Pollution measurements
  - Weather Conditions

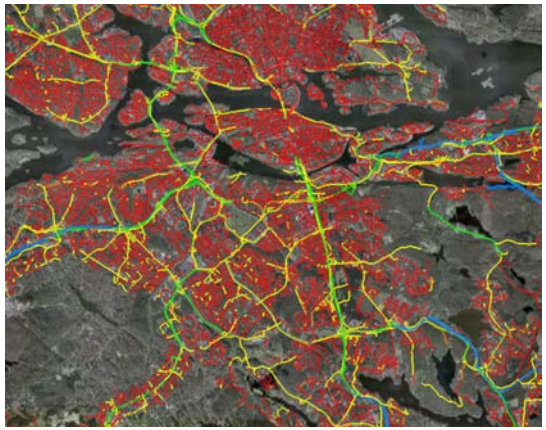

- **Archiving of cleansed data**
- **Real Time Traffic Monitoring**
- **Real Time Traffic Information**
- **(Multimodal) Travel Planner**

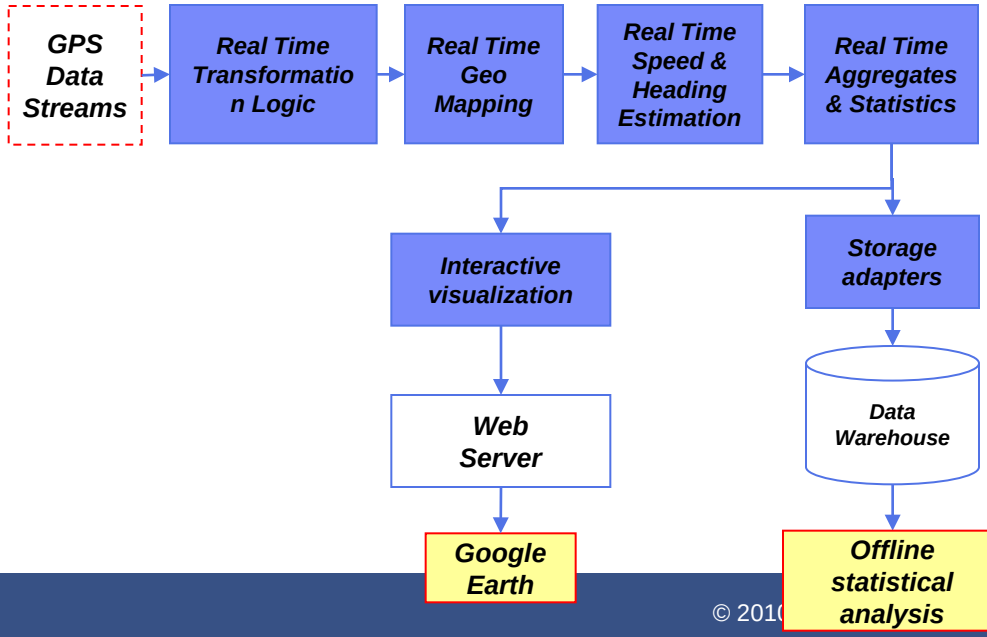

Only 4 x86 Blade servers to process 250,000 GPS probes per second

Rules on Streams (In progress, Research)

Stage 1

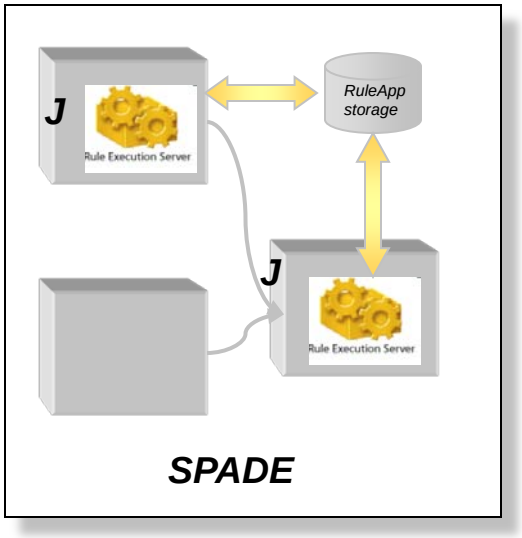

JRules UBOP starts a JVM, which contains rule execution server.

Uses Java UBOP

Schema mapping processed at run-time.

Stage 2

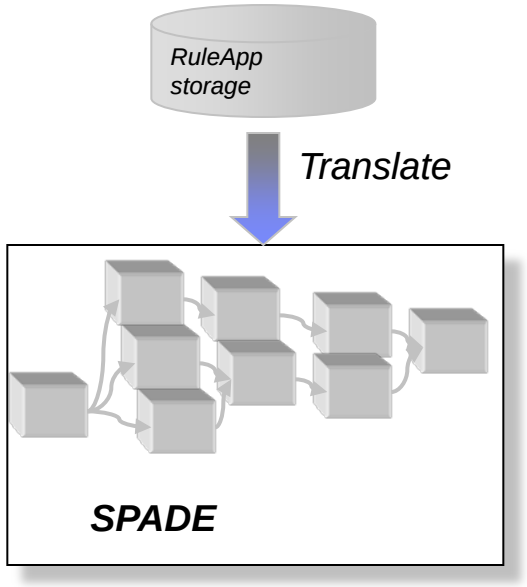

JRules automatically translated to a SPADE program

Does not require Java

Transformations included in generated application

Market Surveillance Solution

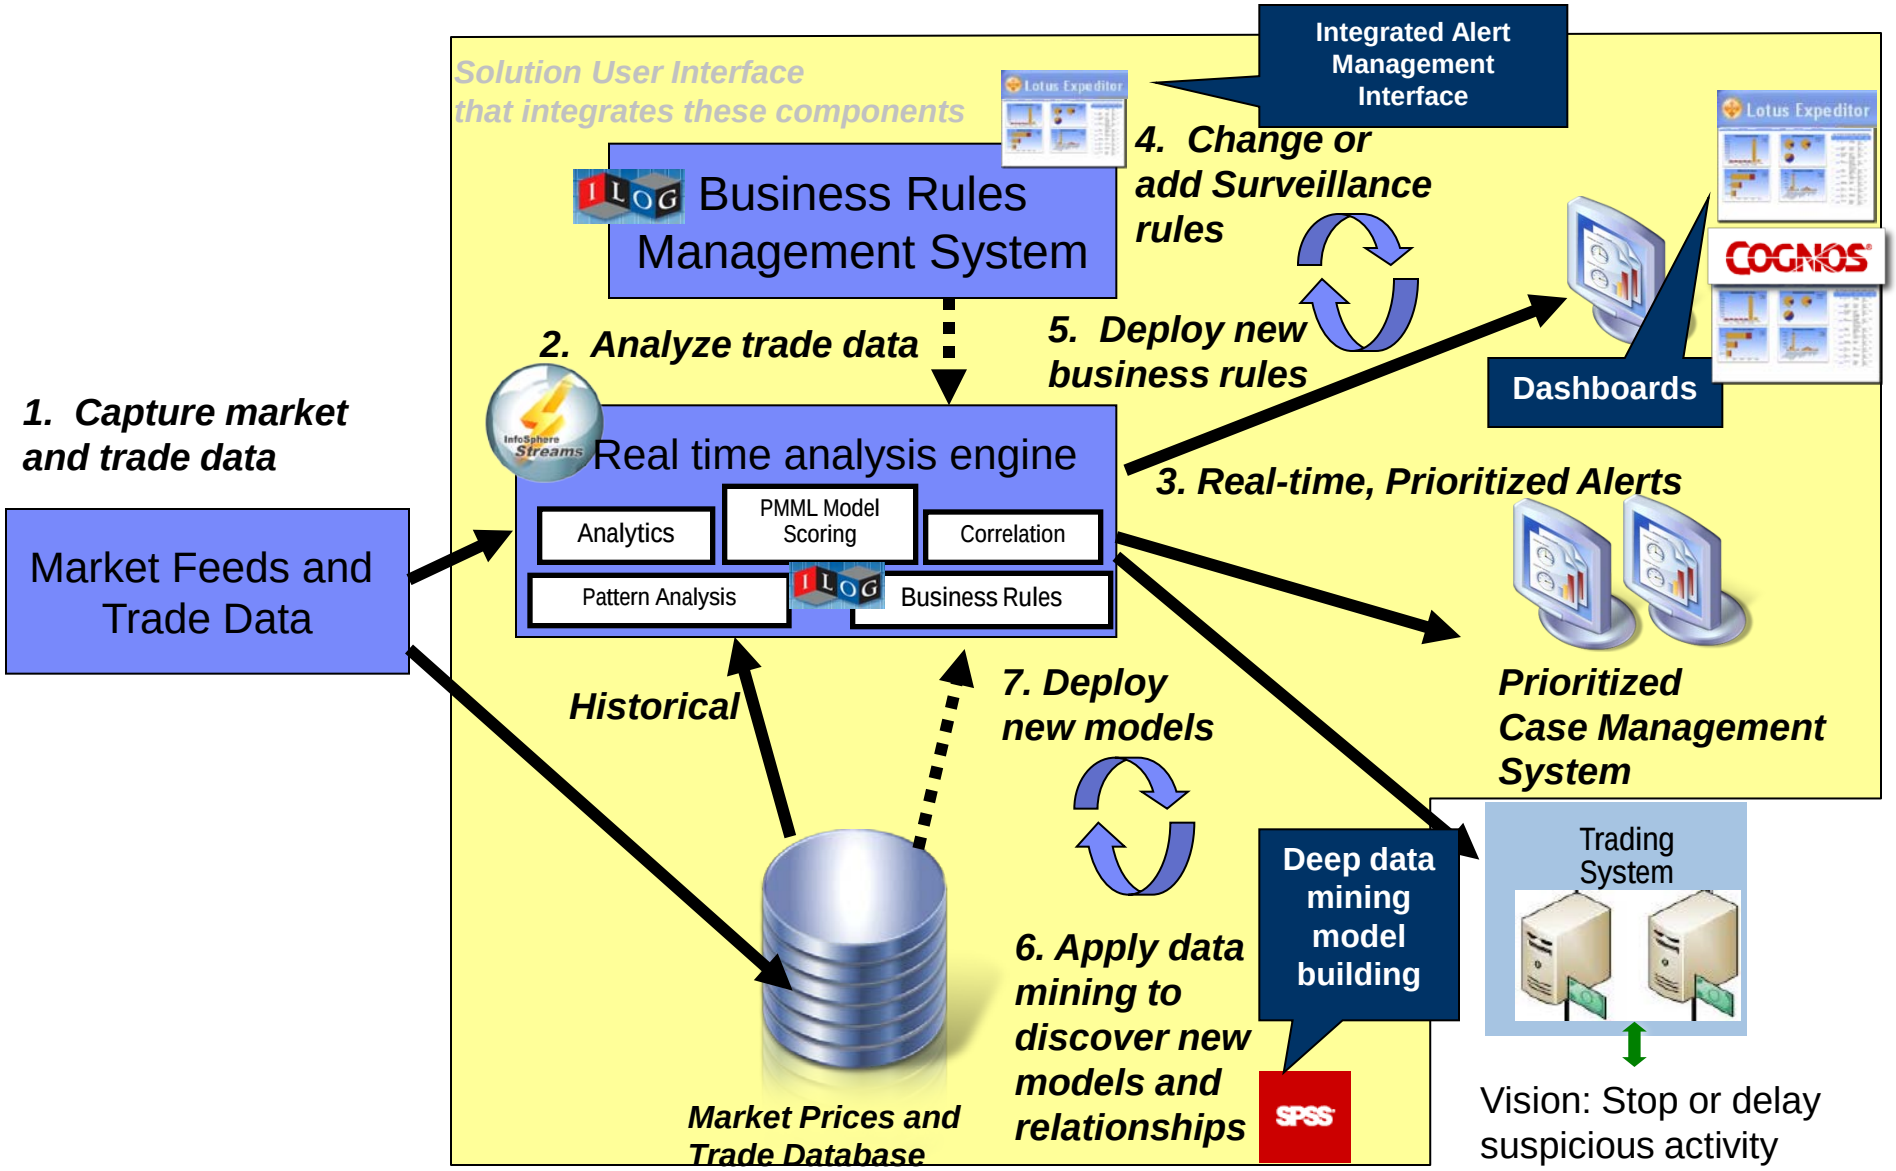

Streams Graph Toolkit (In Progress, Research)

- **Various operations over graph structures**
  - interconnections of nodes and edges
- **Key Components**
  - Graph Aggregate will compute aggregations over moving windows of streams of graph edges
  - Continuous/spot queries over distributed graphs
  - Temporal Graph captures temporal relations (edges) between entities (nodes)
    - Nodes: IP addresses, Event names, Subnets, Email IDs, …
    - Edges: Entity ‘B’ *typically* occurs *soon after* entity ‘A’
    - Graph characterized by proximity gap (k-value) which specifies the proximity relation

Time Series Analysis (In Progress, Research)

Data in motion, data at rest

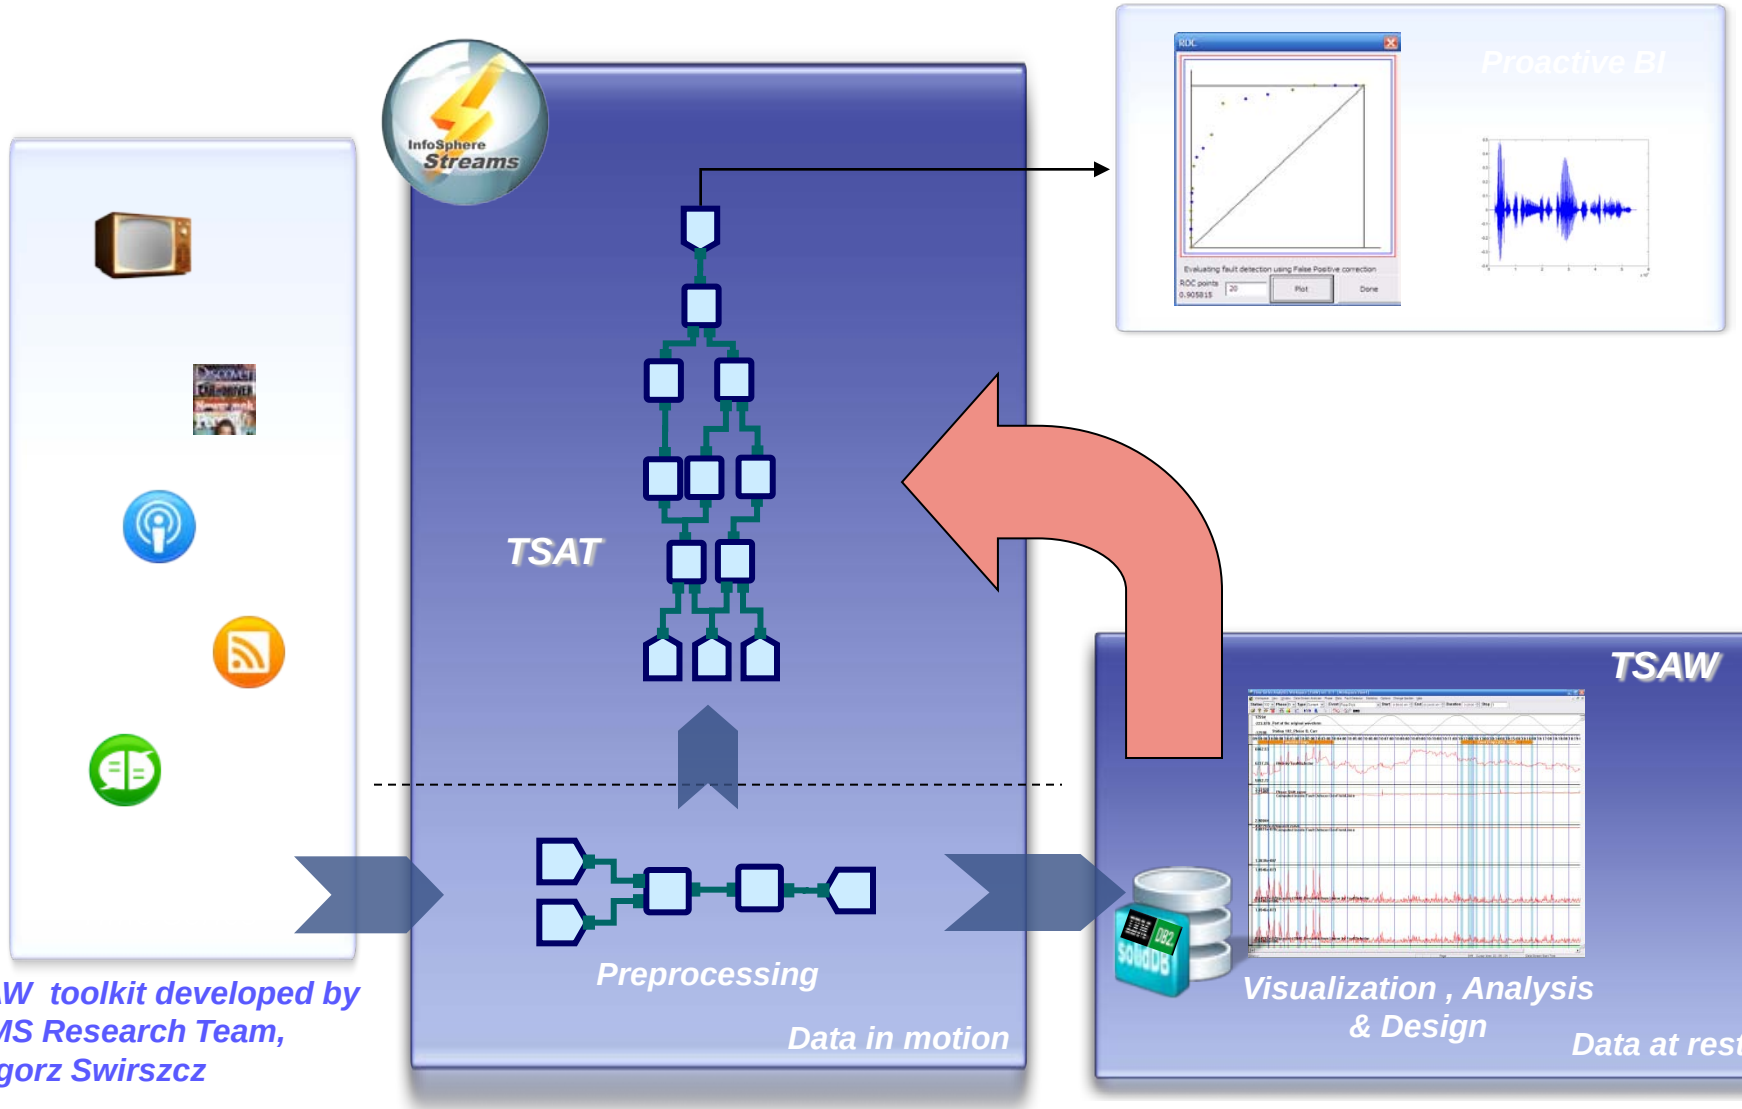

Predictive Analytics using InfoSphere Streams in a neo natal ICU has potential to detect life threatening conditions up to 24hrs earlier

- **Real Time analytics and correlations on physiological data streams**
  - *Blood pressure, Temperature, EKG, Blood oxygen saturation etc.,*
- **Early detection of the onset of potentially life threatening conditions**
  - *Potentially up to 24 hours earlier than current medical practices*
  - *Early intervention leads to lower patient morbidity and better long term outcomes*
- **Technology also enables physicians to verify new clinical hypotheses**

UOIT
CHALLENGE INNOVATE CONNECT

Summary

- **Streams enables**
  - Continuous, scalable data analysis
    - Single Streams instances of up to 125 servers
  - Wide variety of analytics over a variety of data types
  - Full data analytic lifecycle management
- **Extends and complements rules-based approaches**
  - Data mining approach enables models to be data driven, higher accuracy, and automatically updated
    - Alleviates full dependence on expert formulated rules
  - Higher accuracy and prioritized alerts reduces analyst workload