



# Business Insight Through Real-time Analysis on Data in Motion

**Marcus Collins**

*Senior Analyst*

*Burton Group*

*Data Management Strategies*

[mcollins@burtongroup.com](mailto:mcollins@burtongroup.com)

[www.burtongroup.com](http://www.burtongroup.com)

Thursday – July 15, 2010



- **Thesis**

- Increasingly enterprises are realizing that the retrospective view business intelligence supports is no longer sufficient to meet the competitive demands of operating at “internet speed”. This is compounded by the increasingly sensor-enabled physical world resulting in a dramatic increase in the volume of business events received as digital feeds. The real-time analysis of streaming business events is becoming an important emerging technology. IT should understand the technology and work with the business to realize the benefits of adopting an event-driven paradigm.



# Agenda

- **Agenda**
  - Business Case and Organizational Relevance
  - Event Processing vs. Business Intelligence
  - Event Processing Process Architecture
  - Event Processing Technical Architecture
  - Standards
  - Recommendations
  - Conclusion




# Agenda

- **Agenda**
  - *Business Case and Organizational Relevance*
  - Event Processing vs. Business Intelligence
  - Event Processing Process Architecture
  - Event Processing Technical Architecture
  - Standards
  - Recommendations
  - Conclusion

- **Business Case**

- Technology innovation in sensors (RFID, GPS etc.)
- Increased instrumentation of key business processes
- Increased reliance on external information
- Three trends:
  - Increase in the volume of business events received as digital feeds
  - Business decisions requiring increasingly complex event scenarios
    - Algorithmic stock trading – combine stock trades with industry news
  - The value of business events degrade over time



“Sophisticated processing on large number of events in a short, well defined, time window”



- **Organizational Relevance**

- Real-time analysis is relevant to organizations:
  - Make faster decision in response to business events
  - Respond faster to business threats and opportunities
  - Reduce the cost of manually reacting to business events
  - Derive business value or operational efficiencies from an increasingly sensor-enabled and instrumented business environment
- Use cases
  - eCommerce – clickstream analysis, fraud detection
  - Supply chain/logistics – inventory movements
- Not exclusively appropriate to high volume and automated processing
  - Analysis of low volume events and retrospective analysis can still have significant business value

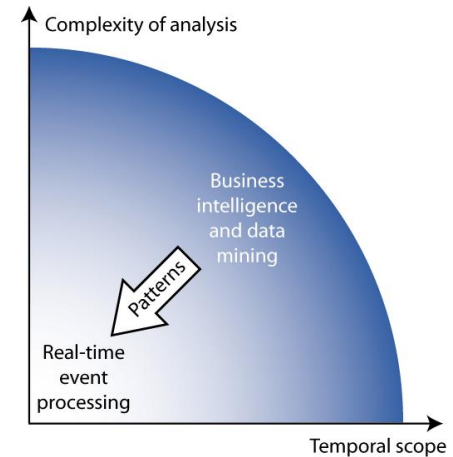


# Agenda

- **Agenda**
  - Business Case and Organizational Relevance
  - *Event Processing vs. Business Intelligence*
  - Event Processing Process Architecture
  - Event Processing Technical Architecture
  - Standards
  - Recommendations
  - Conclusion

- **Real-time Event Processing vs. Business Intelligence**

- Business Intelligence
  - Retrospective and longer term view – strategic focus
  - High degree of analysis complexity
- Real-time Event Processing
  - Immediate and short term – operational focus
  - Low/medium degree of analysis complexity
- Close Relationship
  - Discovery of processing or behavioral patterns through data mining
- Complex Event Processing (CEP)
  - Complex event capture (heterogeneous event streams)
  - Complex event analysis (business rules engine and feedback mechanisms)



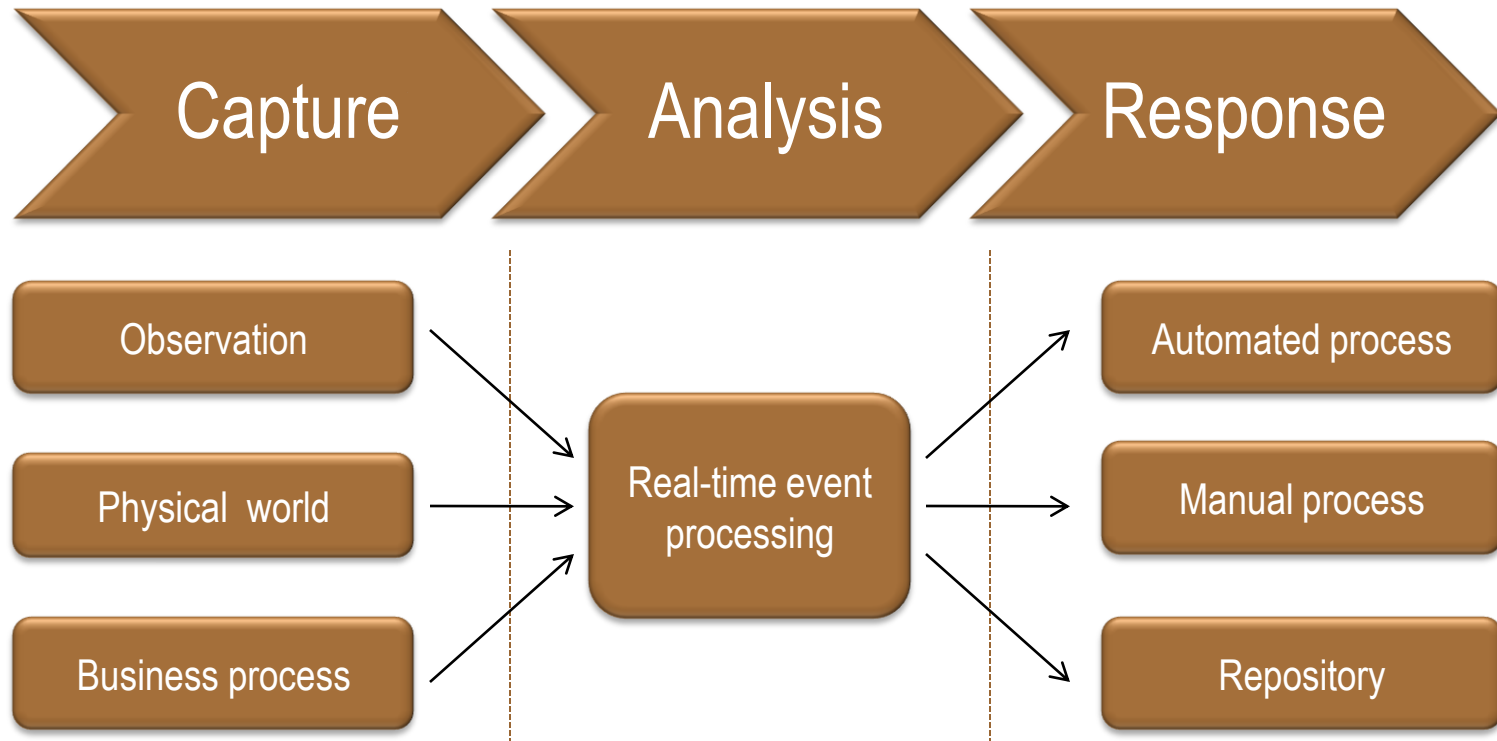




# Agenda

- Agenda
  - Business Case and Organizational Relevance
  - Event Processing vs. Business Intelligence
  - *Event Processing Process Architecture*
  - Event Processing Technical Architecture
  - Standards
  - Recommendations
  - Conclusion

- **Real-time Event Processing – Process Architecture**



- **Process Architecture**

- **Event Capture**

- **Observations**

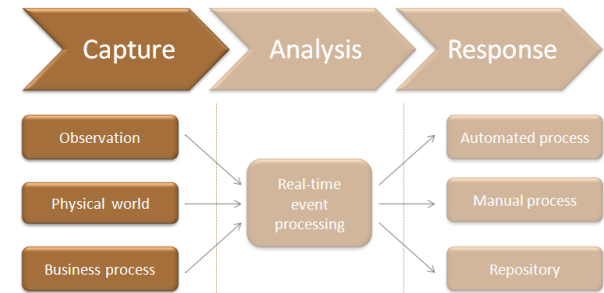
- Measurements and observations made by humans

- **Physical world sensors**

- RFID readers, GPS sensors, environmental sensors

- **Business processes**

- Capturing real-time changes to internal business processes e.g. sales, inventory movements etc.



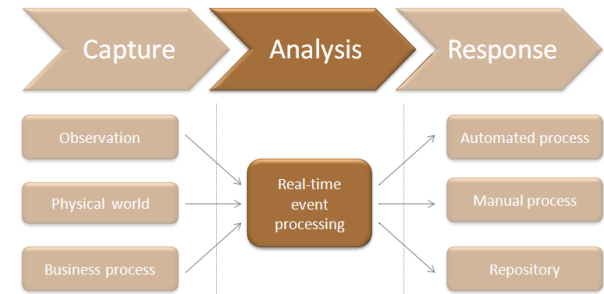
- **Process Architecture**

- **Event Analysis**

- Generation of *complex events*

- Processing event data

- Filtering (threshold / by exception)
- Event identification and classification
- Event correlation (out of sequence/missing events)
- Enrichment of sparse base events (context)
- Calculations (summarization, aggregation etc.)
- Application of business rules.



- **Process Architecture**

- **Event Response**

- **Automated processes**

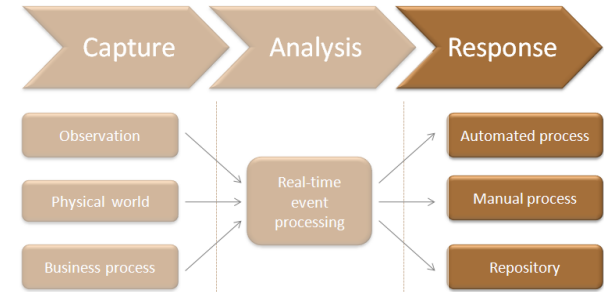
- Business rules and actions well understood
- Automate the response process

- **Manual processes**

- Enabling human based decisions

- **Repository**

- Event storage for retrospective analysis



- **Process Architecture**

- **Data Quality**

- **Complex to solve**

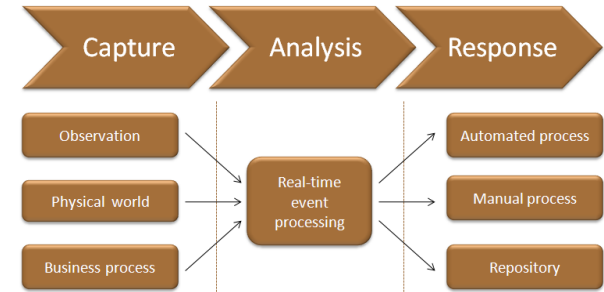
- Resolving the quality issues *in-flight* challenging
- Prevalence of externally sourced data streams

- **Important**

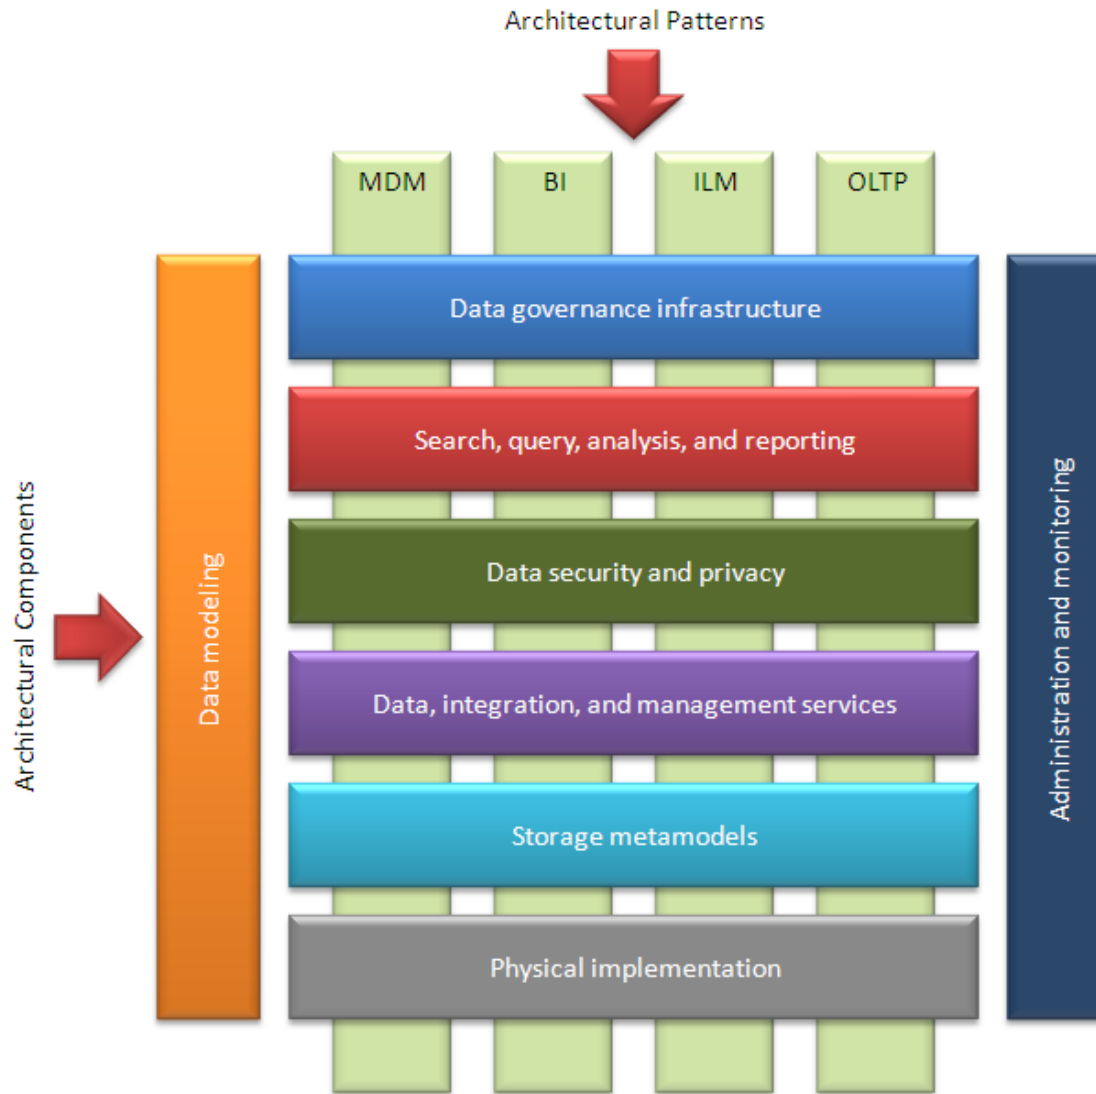
- Drives an enterprises operational decision making

- **Solutions**

- Correction – issues with legacy apps and external data streams
- Data profiling – retrospective data profiling, apply in Analysis
- Quality metrics – annotate complex events with quality metrics.

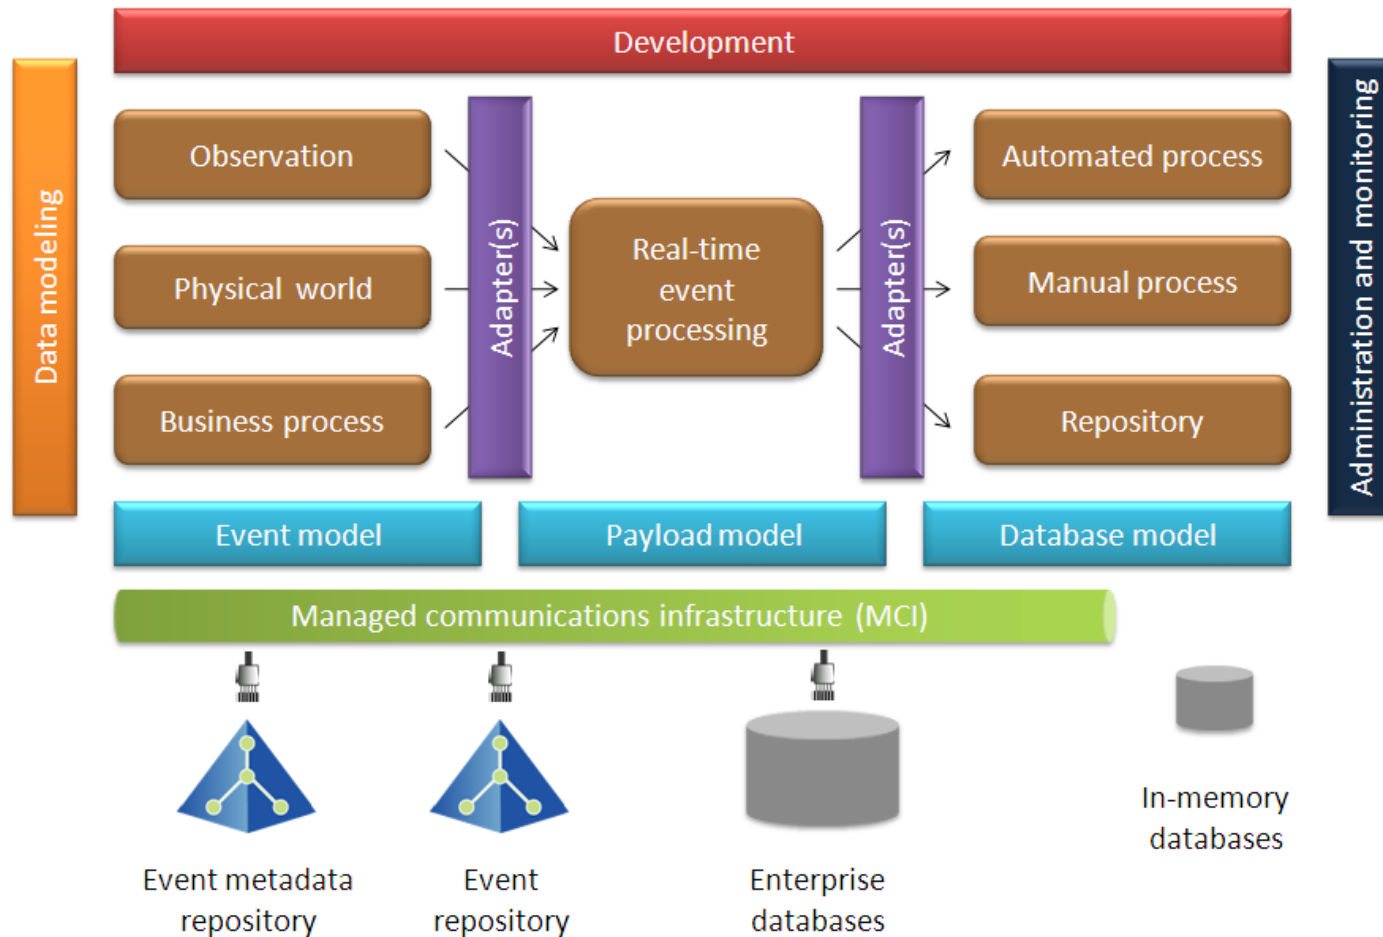


- Agenda
  - Business Case and Organizational Relevance
  - Event Processing vs. Business Intelligence
  - Event Processing Process Architecture
  - *Event Processing Technical Architecture*
  - Standards
  - Recommendations
  - Conclusion





- Real-time Event Processing – Technical Architecture



- **Technical Architecture**

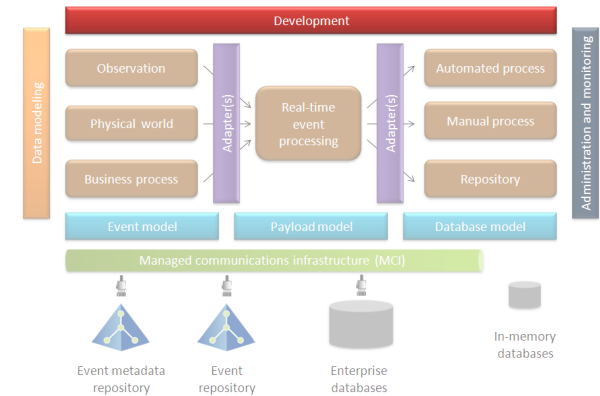
- **Development Environment**

- **Functionality**

- Adapter components
- Logic and rules of event processing

- **Collaborative and iterative environment**

- Value will be realized through collaboration between IT and business domain experts
- Pattern of events associated with significant business change are often only partially understood
- Business rules tend to change more frequently than does the business object with which they are associated.



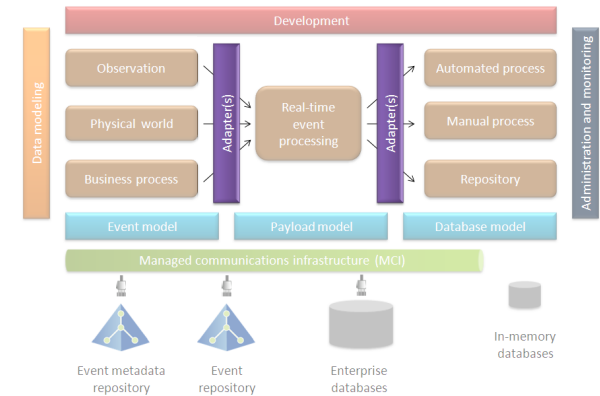
- **Technical Architecture**

- Adapter Architecture

- Integration point

- Pre-built adapters

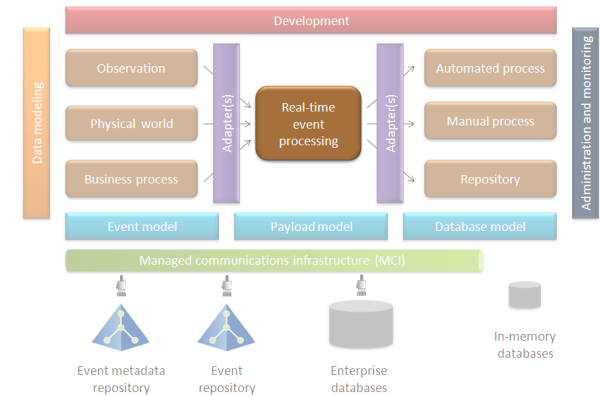
- Physical world sensors including RFID and GPS
- Messaging infrastructures including Java Message Service (JMS)
- Web feeds including Really Simple Syndication (RSS), Atom Syndication Format (Atom) etc
- Procedure invocation including Remote Procedure Calls (RPC), web services (SOAP and REST) etc
- Stock market data feeds from domestic and international exchanges
- Syndicated news feeds including Bloomberg, Reuters etc.



- **Technical Architecture**

- **Event Processing Engine**

- Functionality (earlier)
- Non-functional
  - Availability
  - Scalability
    - Scale-out (thread based architecture)
    - Scale-up (“divide and conquer” – application partition).



- **Technical Architecture**

- **In-memory Databases**

- **Required for low-latency**

- Not large cache in traditional db's

- Traditional db's:

- Minimized I/O overhead (buffer cache)

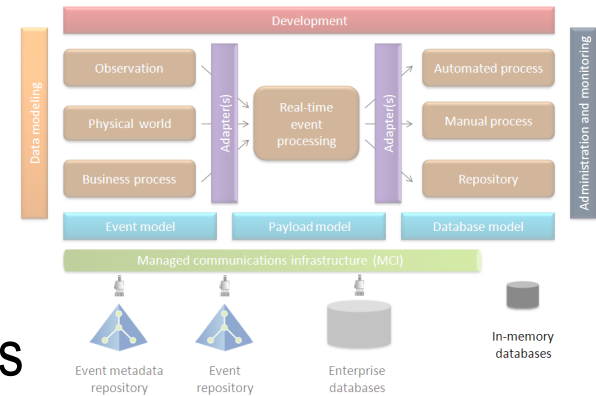
- Query optimizer and index structures minimize I/O

- In-memory db's:

- Absence of large block structures

- Indexes computationally cheap to scan

- Relax ACID transaction requirements.



- Agenda
  - Business Case and Organizational Relevance
  - Event Processing vs. Business Intelligence
  - Event Processing Process Architecture
  - Event Processing Technical Architecture
  - Standards
  - Recommendations
  - Conclusion

- **Standards**

- Enabler to technology adoption and vendor interoperability
- Lack of mature standards and common data formats
  - Standards within the adapter (integration) layer
  - Two initiatives to standardize event metadata and processing
    - Mitre Common Event Expression (CEE)
    - IBM Common Event Base specification



- Lack of standards should not delay evaluation and implementation
  - Considerable business value to be gained by real-time analysis and transitioning to an event based architecture.

- Agenda
  - Business Case and Organizational Relevance
  - Event Processing vs. Business Intelligence
  - Event Processing Process Architecture
  - Event Processing Technical Architecture
  - Standards
  - Recommendations
  - Conclusion



- **Recommendations**

- Understand the impact of event processing on business processes
  - Event processing is becoming increasingly important due to an increase in the volume of business events that are represented as digital feeds
  - The technical complexity means that the IT organization should lead the drive to understand the event-driven paradigm
  - The IT organization should work with the business community to explore areas within the enterprise where it could be of potential benefit.
- Engage business domain experts
  - The real value of adopting an event-driven paradigm will be realized through the involvement of business domain experts
  - The success of the event processing initiative will be measured in terms of the speed and accuracy of decision making and in the responsiveness of the enterprise to changes in the business environment.

- **Recommendations**

- Low volume events can have significant business value
  - The technical risk and cost of implementing a low volume event scenario will be less than a high-volume low-latency scenario
  - Ensure that the event rate and processing complexity is appropriate for the business process and focus on the low risk scenarios especially at the initial stages of the adoption.
- Retrospective analysis can have significant business value
  - Analyzing business events to determine the drivers of business change is of significant value to enterprises
  - This retrospective analysis is often a low risk approach as a pre-cursor to the adoption of a wider real-time event-driven initiative.

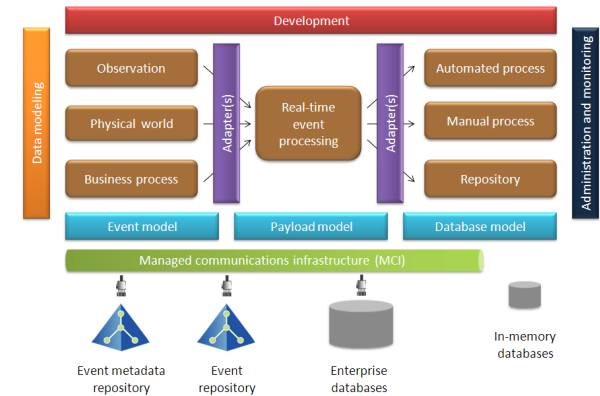
- **Recommendations**

- Start with a proof-of concept initiative
  - As with all emerging technology adoption initiatives – start small.
  - Focus on business units with:
    - a well understood business processes
    - where the adoption of an event-driven paradigm will yield measurable business value
  - As the initiative progresses and the business benefits can be shown to have been realized use this as a poster child to articulate the value of event-driven processing to the rest of the enterprise.
- Lack of standards should not limit adoption
  - Emerging technologies have few comprehensive standards for infrastructure and processing
  - This lack of standards should not delay organizations from implementing an event-driven paradigm if the business case is present.

- **Recommendations**

- Take an Architectural Approach

- The architecture of an event processing infrastructure is complex, consisting of many interlinked components that reside both within the enterprise and in the external environment
    - The adoption of an event-driven paradigm may include the task to instrument objects in the physical world i.e. RFID tags and GPS sensors
    - The complexity and the relative immaturity of the technology require the adoption of a reference architecture to reduce the technical risk.



- **Conclusion**

- The adoption of an event-driven paradigm and analysis of data information will enable enterprises to make faster and more accurate decision. Whilst the benefits are great, the complexity of the event-driven infrastructure and processing means that the implementation of such an initiative is fraught with technical risk. The IT organizations should be in the vanguard of this adoption and lead the drive to understand the event-driven paradigm and then work with the business community to explore areas within the enterprise where it could be of potential benefit by driving increased value in key business processes.

