

# The New Generation of Data Management

## In-Memory Technologies in Action

Stephan Grotz  
*Terracotta Solution Architect*

Dr. Jürgen Krämer  
*VP CEP & Analytics*

10.01.2013

# Agenda

- 1) Motivation
- 2) Introduction to Terracotta
- 3) Introduction to CEP
- 4) Introduction to Nirvana
- 5) In-Memory Data Management Platform
- 6) Q&A

# Motivation

**Get There Faster.™**

# Data explosion compounds challenges



## Data is getting bigger

Rapid Growth  
of Global Data  
from 2009-2020<sup>1</sup>

**From  
1 to 35  
Zetabytes**

Global mobile data  
traffic will surpass<sup>3</sup>

**10 exabytes in  
2016**

**70%**  
of the data  
generated by  
individuals<sup>1</sup>

The number of  
mobile-connected  
devices will exceed  
the world's  
population in 2012<sup>3</sup>

**7 Billion**

RFID Market to see  
some serious growth<sup>2</sup>

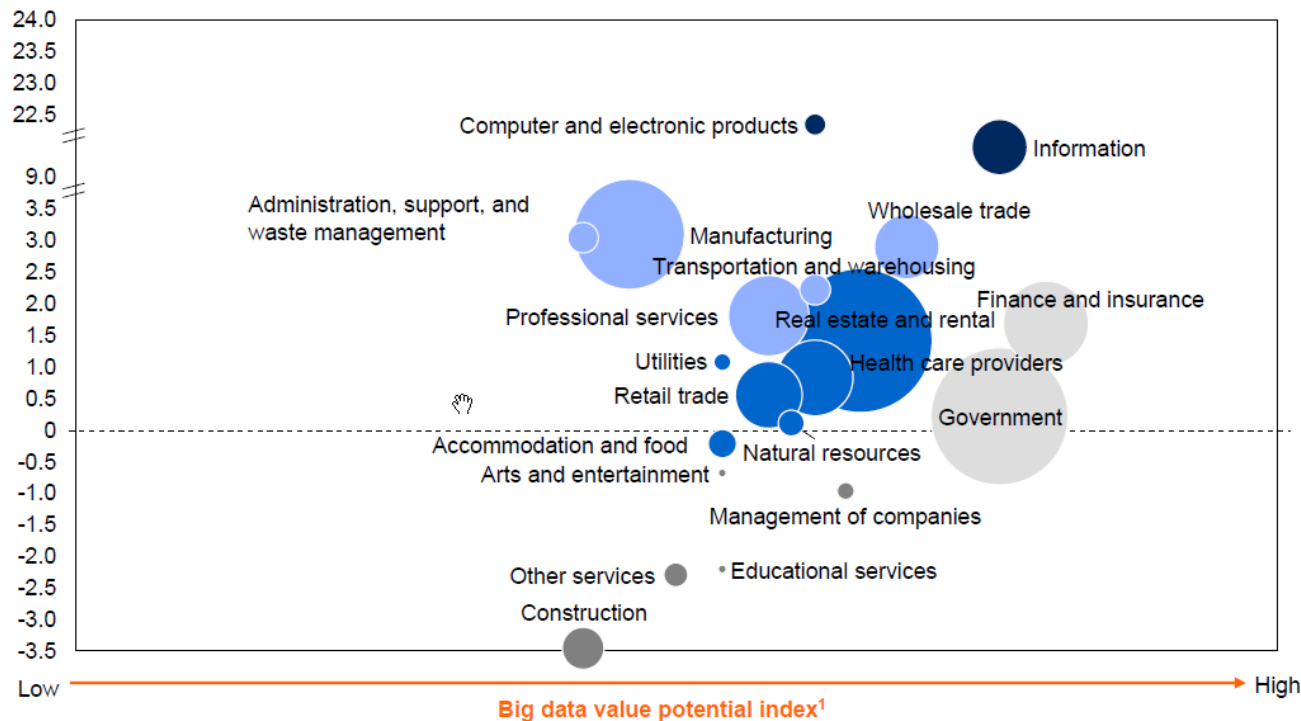
**125 Billion  
RFID tags in 2020**

Every minute in the  
Internet<sup>4</sup>

**100.000 Twitter  
tweets**

**240.000 shared  
Facebook content**

# Big Data can generate significant financial value across sectors



## Verticals (Gartner)

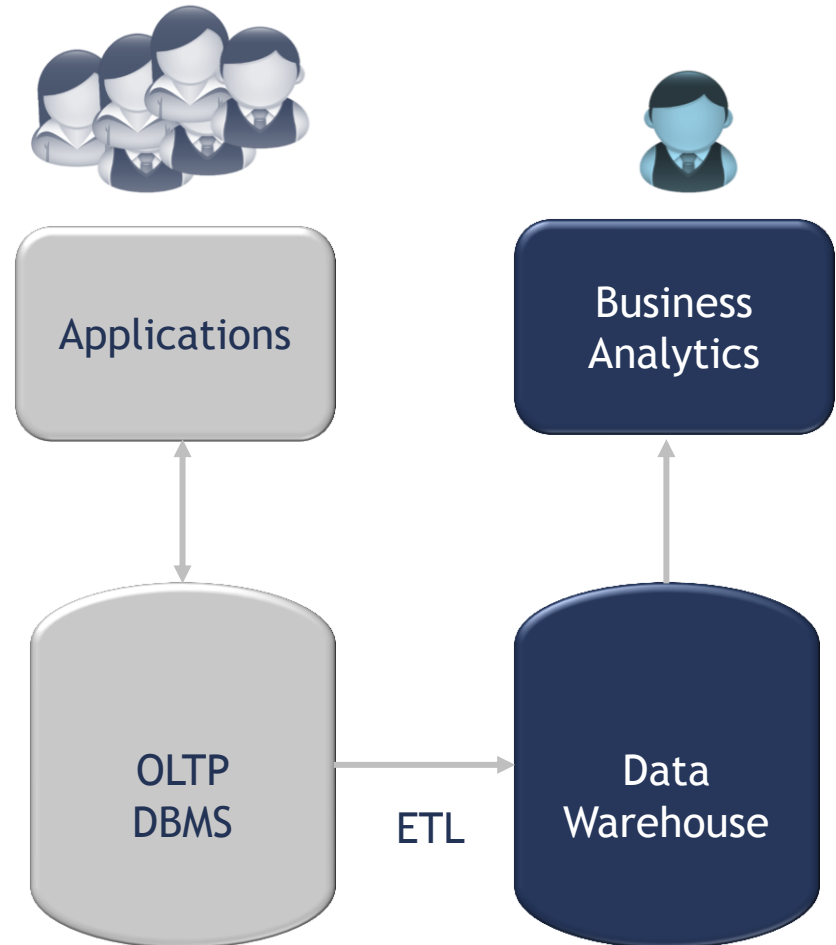
- Securities Trading
- Banking
- Travel and Transportation
- Telecom
- Media and Entertainment
- Online Gaming
- High Tech Manufacturing
- eCommerce/Retail
- ...

McKinsey: Potential increase with Big Data in US retailers' operating margin<sup>1</sup> **60%**

(1) SOURCE:US Bureau of Labor Statistics; McKinsey Global Institute analysis

# Classical data management architecture

- Operational and analytical environments are separated
- Mainly based on structured data
- Data exchange and transformation with latency
- Different ways end-user interacts with data



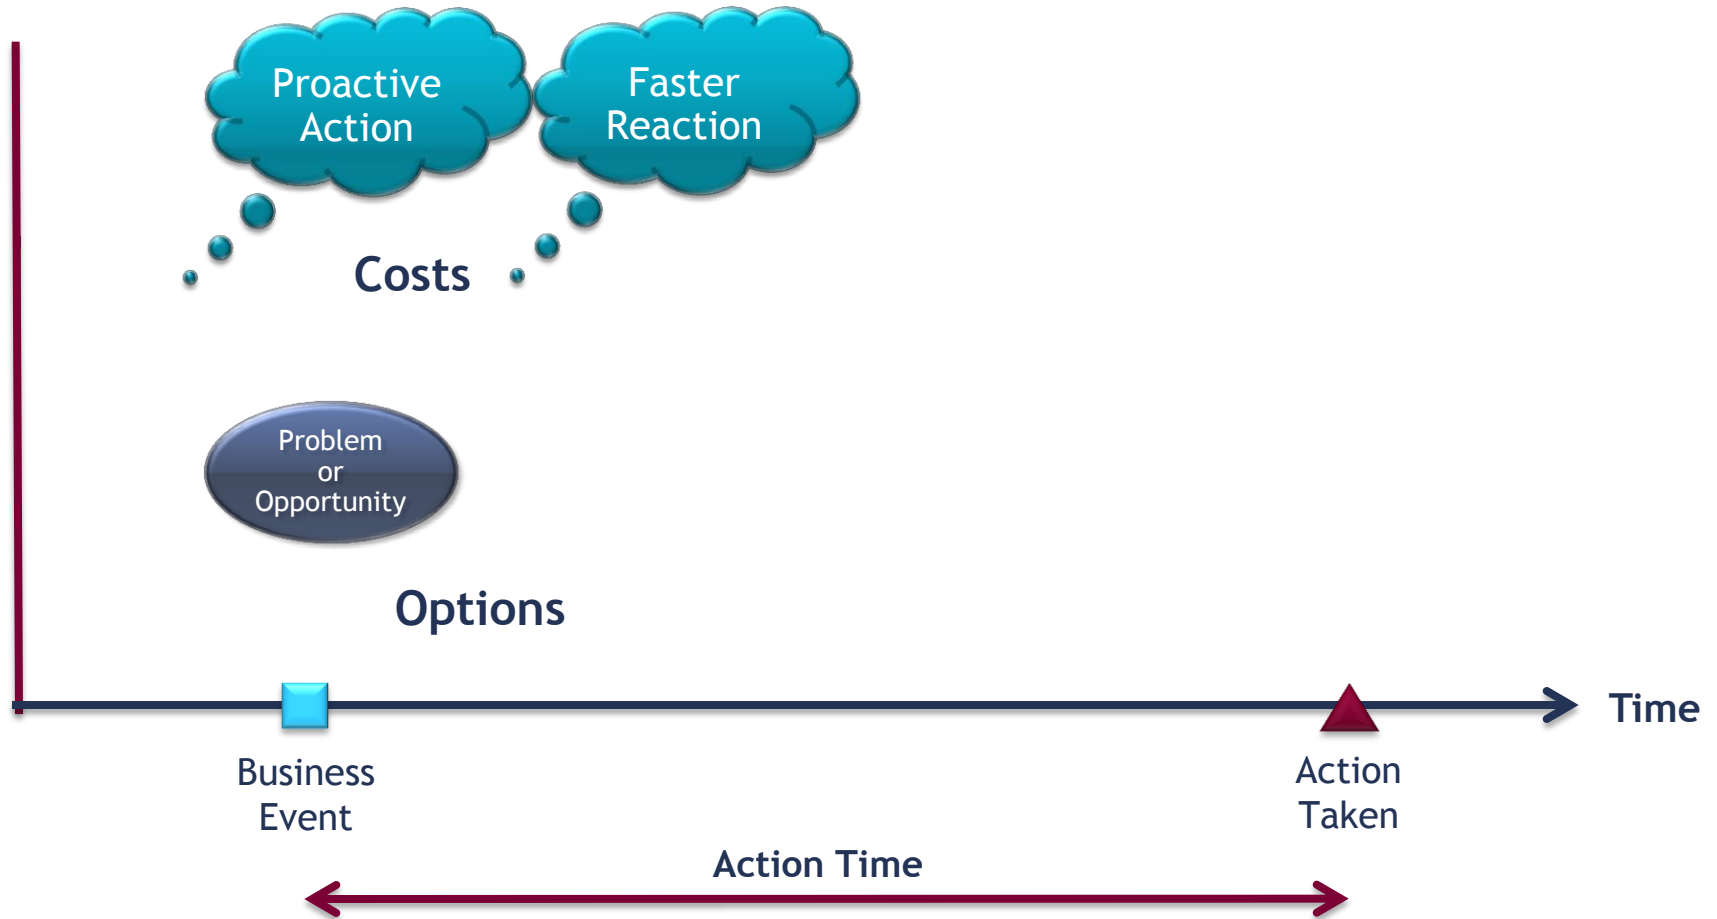
# Pace of Business Drives Need for Real-time Analytics

## Business is Moving to Real-Time Transactions

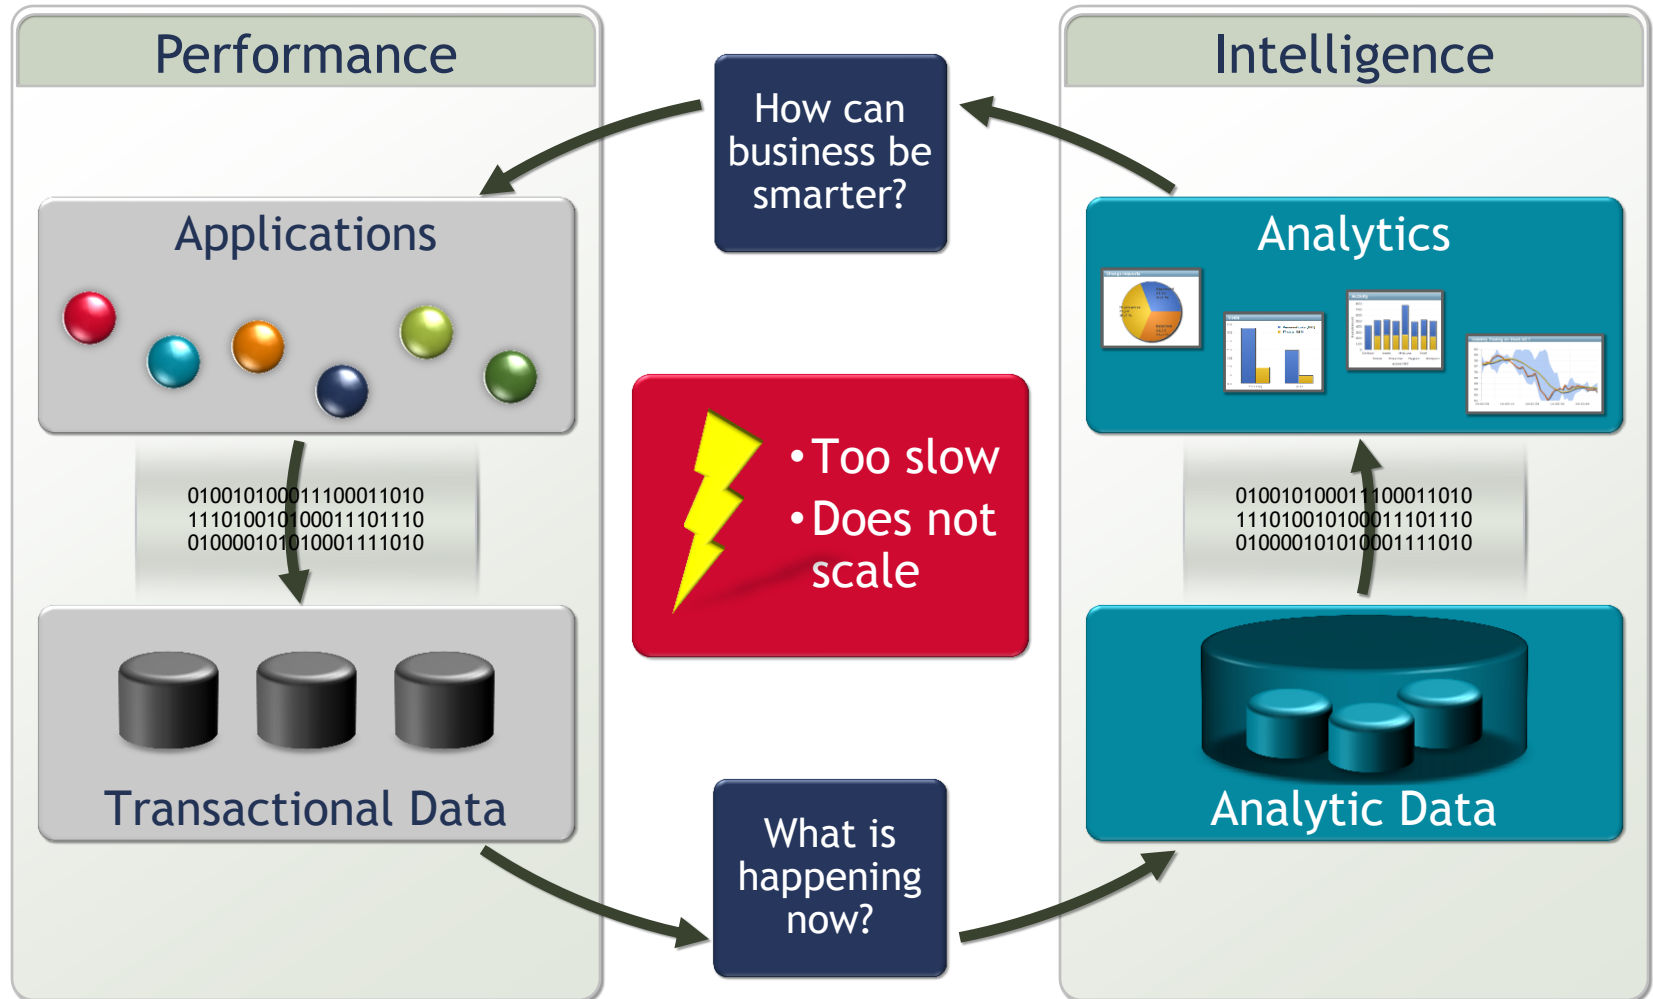




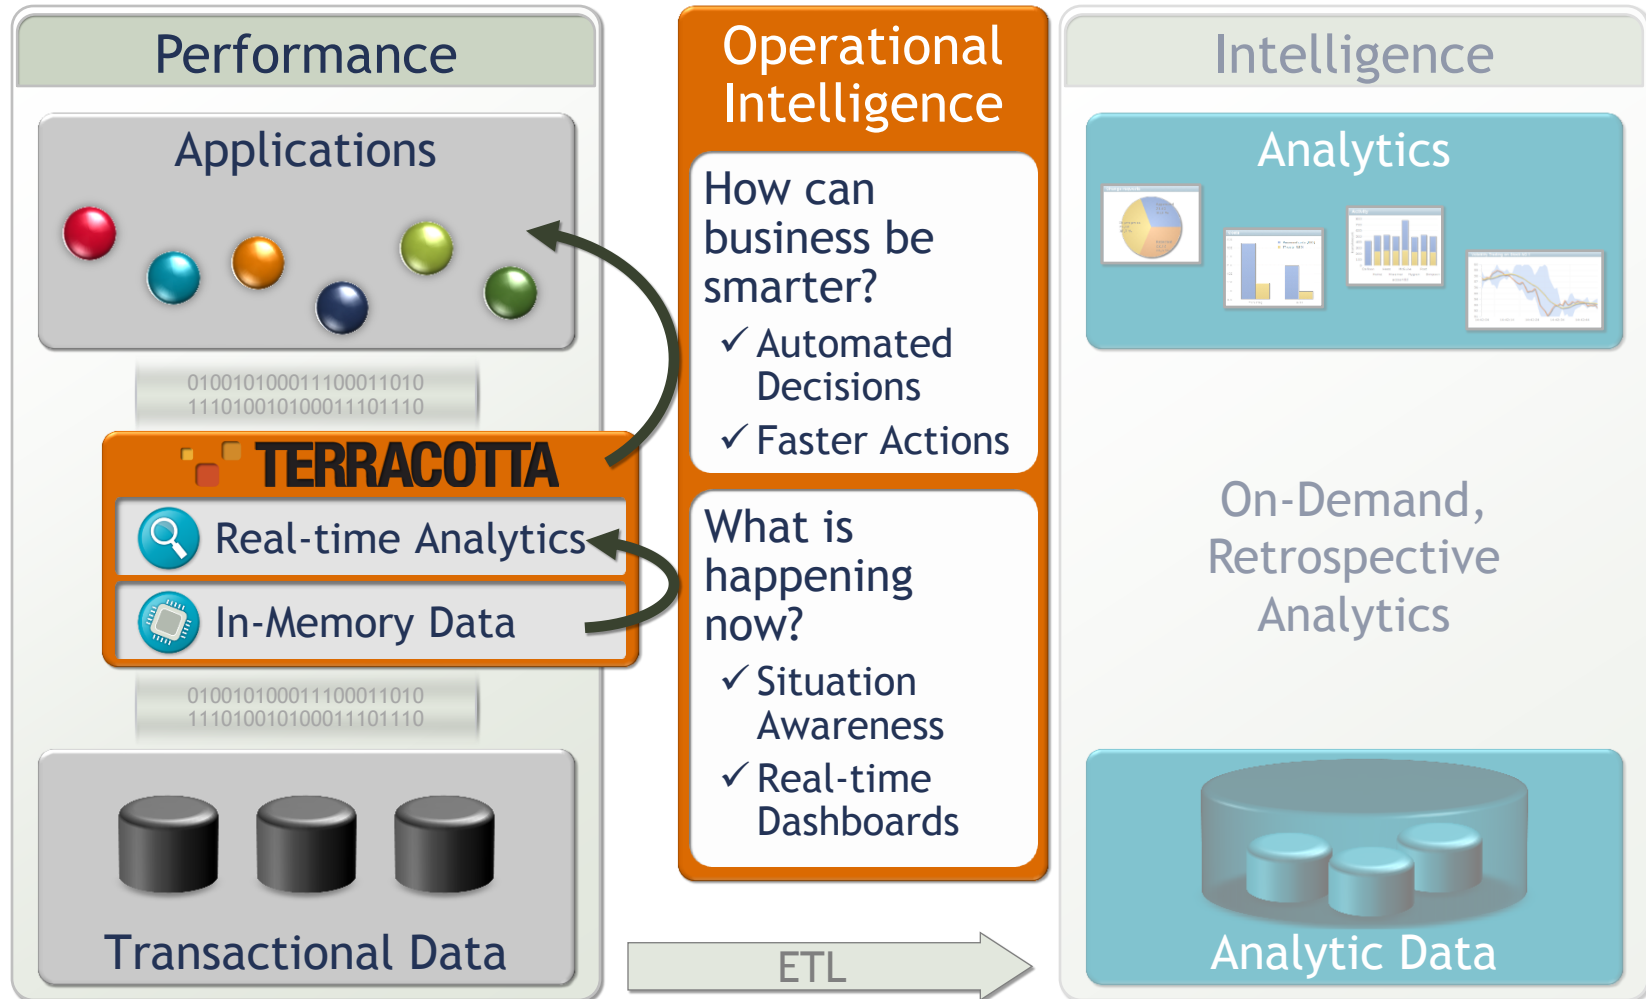
# Business Value of Faster Actions



# Real-time data management goals



# In-Memory data management for the real-time business



# Introduction to Terracotta

**Get There Faster.™**

# TERRACOTTA

- Provider of in-memory data management solutions for the enterprise
- Behind some of the most widely deployed software for scalability and performance
- 70% of Global 10000 use Terracotta
- Wholly-owned subsidiary of Software AG

500,000+ Enterprise Deployments  
Majority of Global 2000



# Big Memory: Deliver Real-Time, High Volume Data Access

## In-Memory

How do we maximize our use of cheap memory?

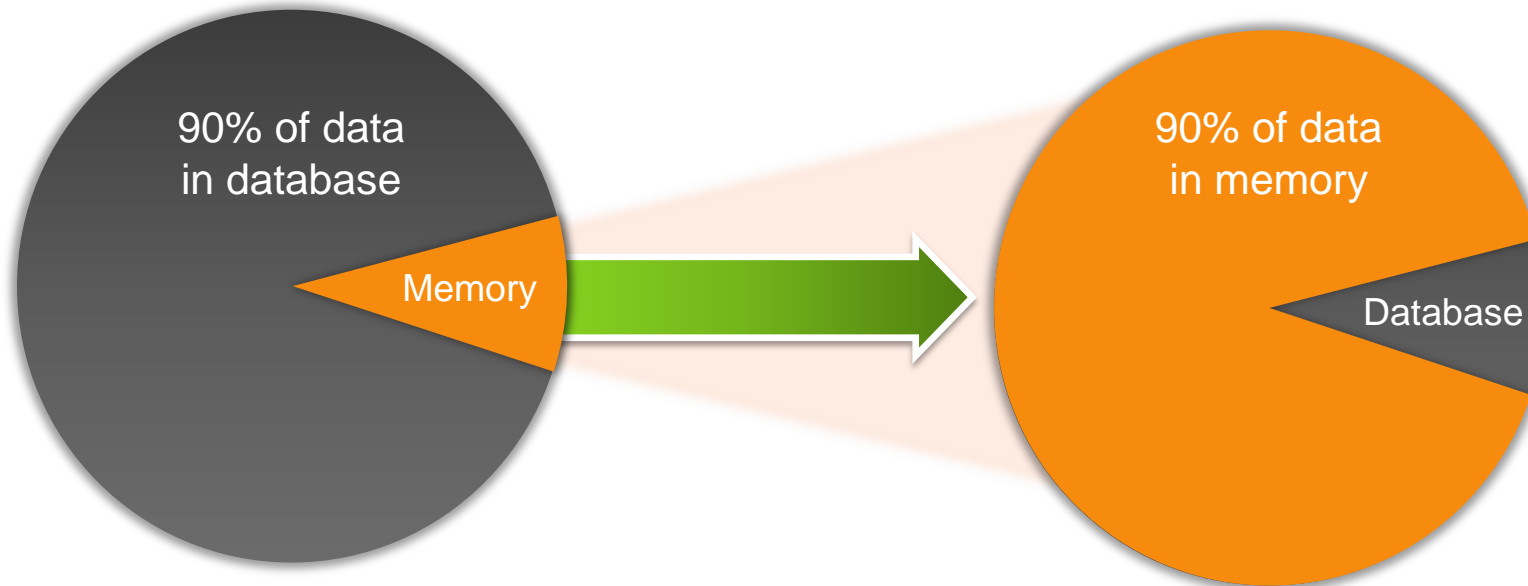
Steep drop in price of RAM

## Big Data

How do we unlock the all value in our data?

Explosion in volume of business data

# BigMemory's Innovation



**Milliseconds**  
App Response Time



**Microseconds**  
App Response Time

# Terracotta BigMemory Delivers Scale & Speed

Real-time access to  
massive amounts of business data

**GO BIG.**

**MORE**

Data

Users

Customers

Transactions

**GO FAST.**

**QUICKER**

Processing

Analysis

Services

Decisions



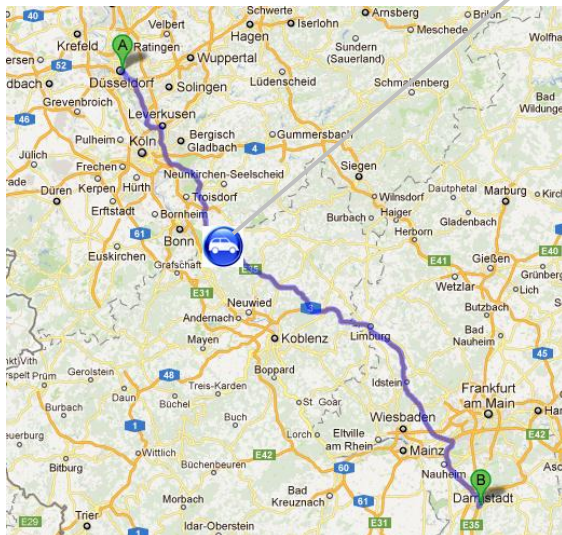
# Introduction to Complex Event Processing (CEP)

**Get There Faster.™**

# What is an Event?

## Example: Vehicle

VehicleId: 23  
 Timestamp: 11:22  
 Geolocation:  
 +50° 43' 33.69", +7° 14' 0.89"



## Example: SmartMeter

MeterId: 8754862  
 Timestamp from: 16:00  
 Timestamp to: 16:15  
 Consumption: 4.67 kWh  
 Max. Power: 1.12 kW



## Example: Stock Price

Share: SOW  
 Timestamp: 14:57 Uhr  
 Price: €26.57



## Event Streams

Stream of position reports from a truck

GPS GPS GPS GPS GPS GPS GPS GPS GPS

Stream of stock prices

Price Price Price Price Price Price Price Price Price

Stream of sensor readings from a smart meter

Value Value Value Value Value Value Value Value Value

## Increase in Data Velocity

Stream of position reports from a truck

GPS GPS GPS GPS GPS GPS GPS GPS GPS

Stream of stock prices

Price Price Price Price Price Price Price Price Price

Stream of sensor readings from a smart meter

Value Value Value Value Value Value Value Value Value

# Increase of Data Volumes


Stream of position reports from a truck



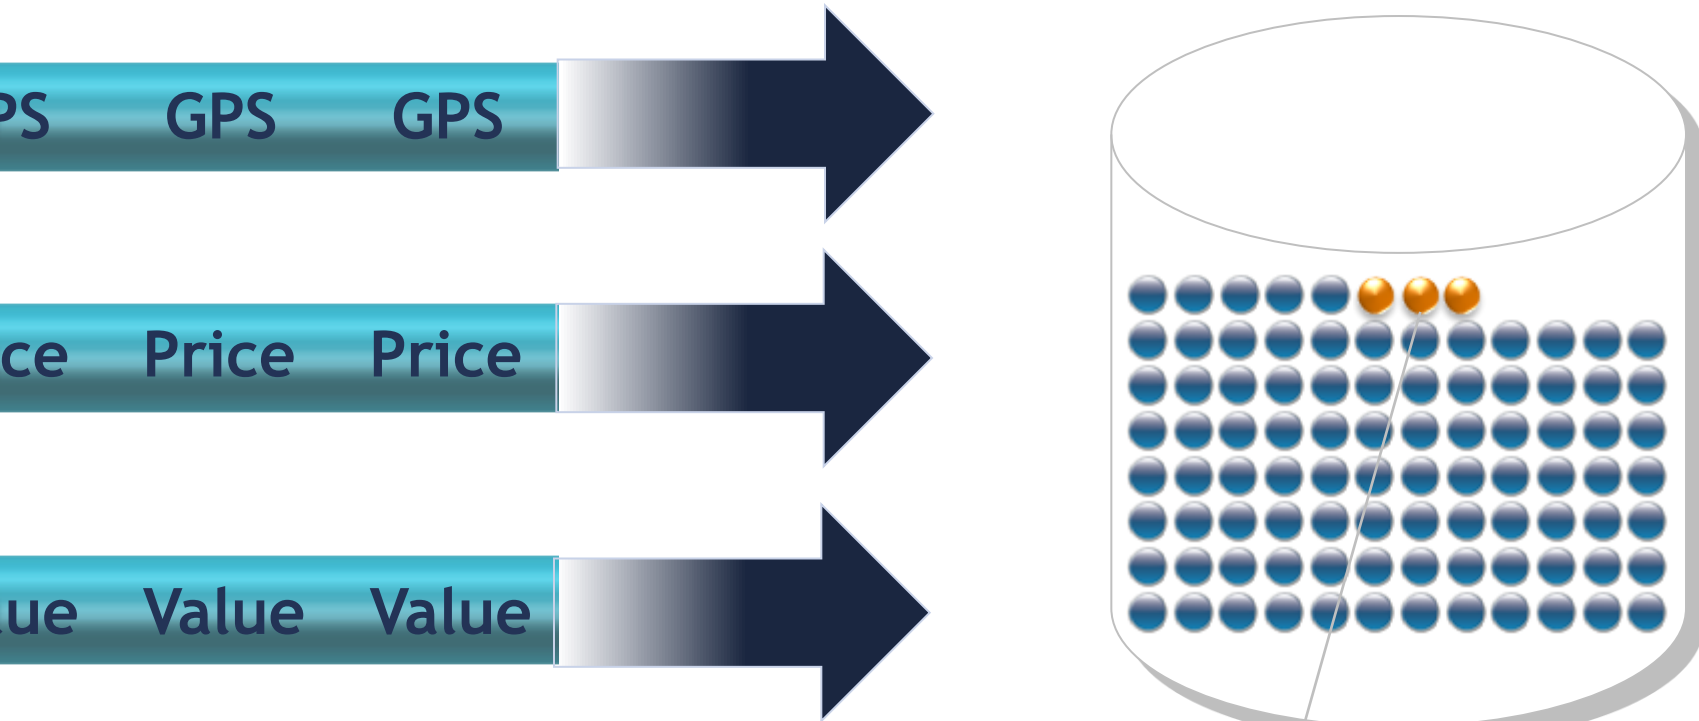
Stream of stock prices



Stream of sensor readings from a smart meter



# Need for Real-time Analytics



→ Newest data is more relevant!  
 → Database becomes bottleneck!

# In-Memory, In-Flight Analytics

## Complex Event Processing (CEP)

Definition: Continuous analytics to derive meaningful business events from different event streams or other event sources like databases in real-time to gain situation awareness and trigger immediate actions.

- Event-driven, incremental processing
- High efficiency and scalability
- Enrich events with context data
- Detect patterns with time/location parameters

GPS GPS GPS

Price Price Price

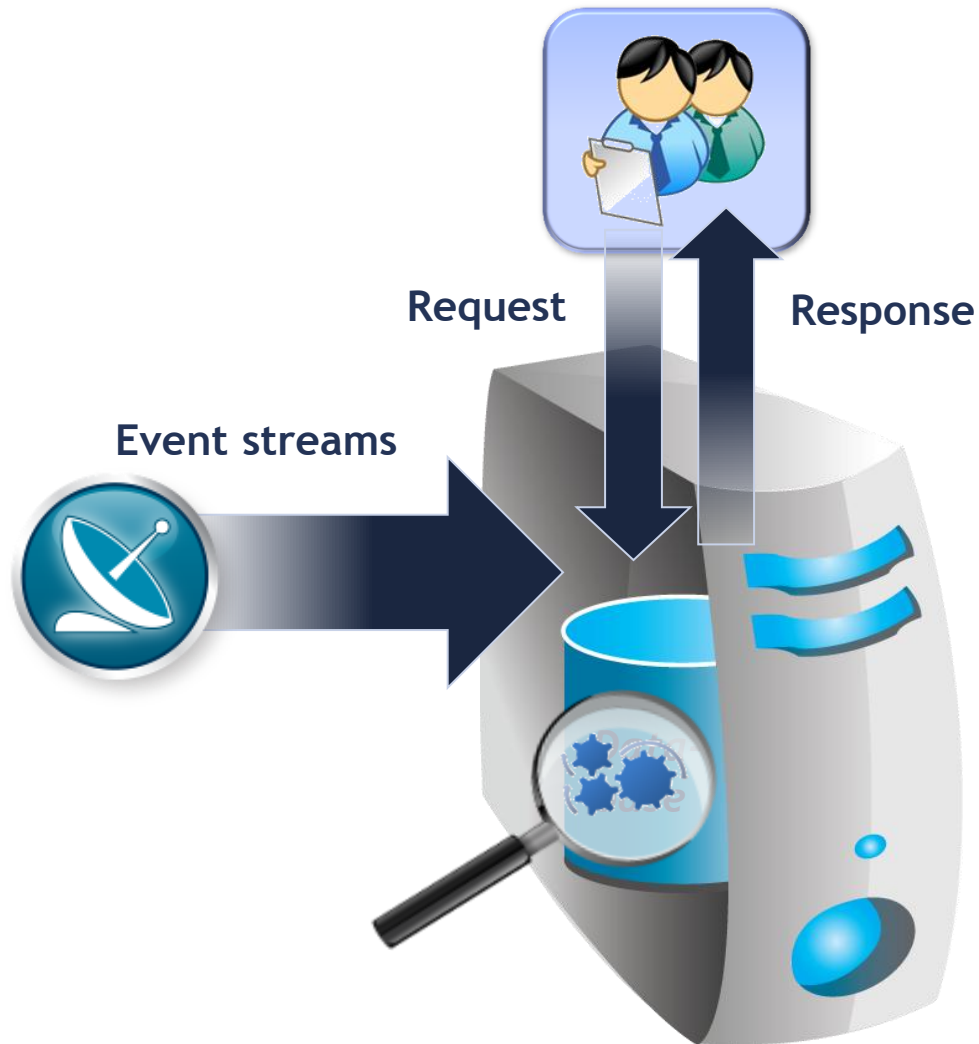
Value Value Value

GPS GPS

Price Price

Value Value

## Traditional „Store-and-Analyze“



## Two phases

1. Store data
2. Pull-based analysis

## Problems

- Data store grows permanently
- ➔ Expensive search & analysis
- Not designed for continuous query evaluation
- ➔ Workarounds entail high load



# Event Stream Processing

## Two phases

1. Store business logic
2. Push-based analysis



Register business logic  
(patterns, KPIs, ...)

Input streams

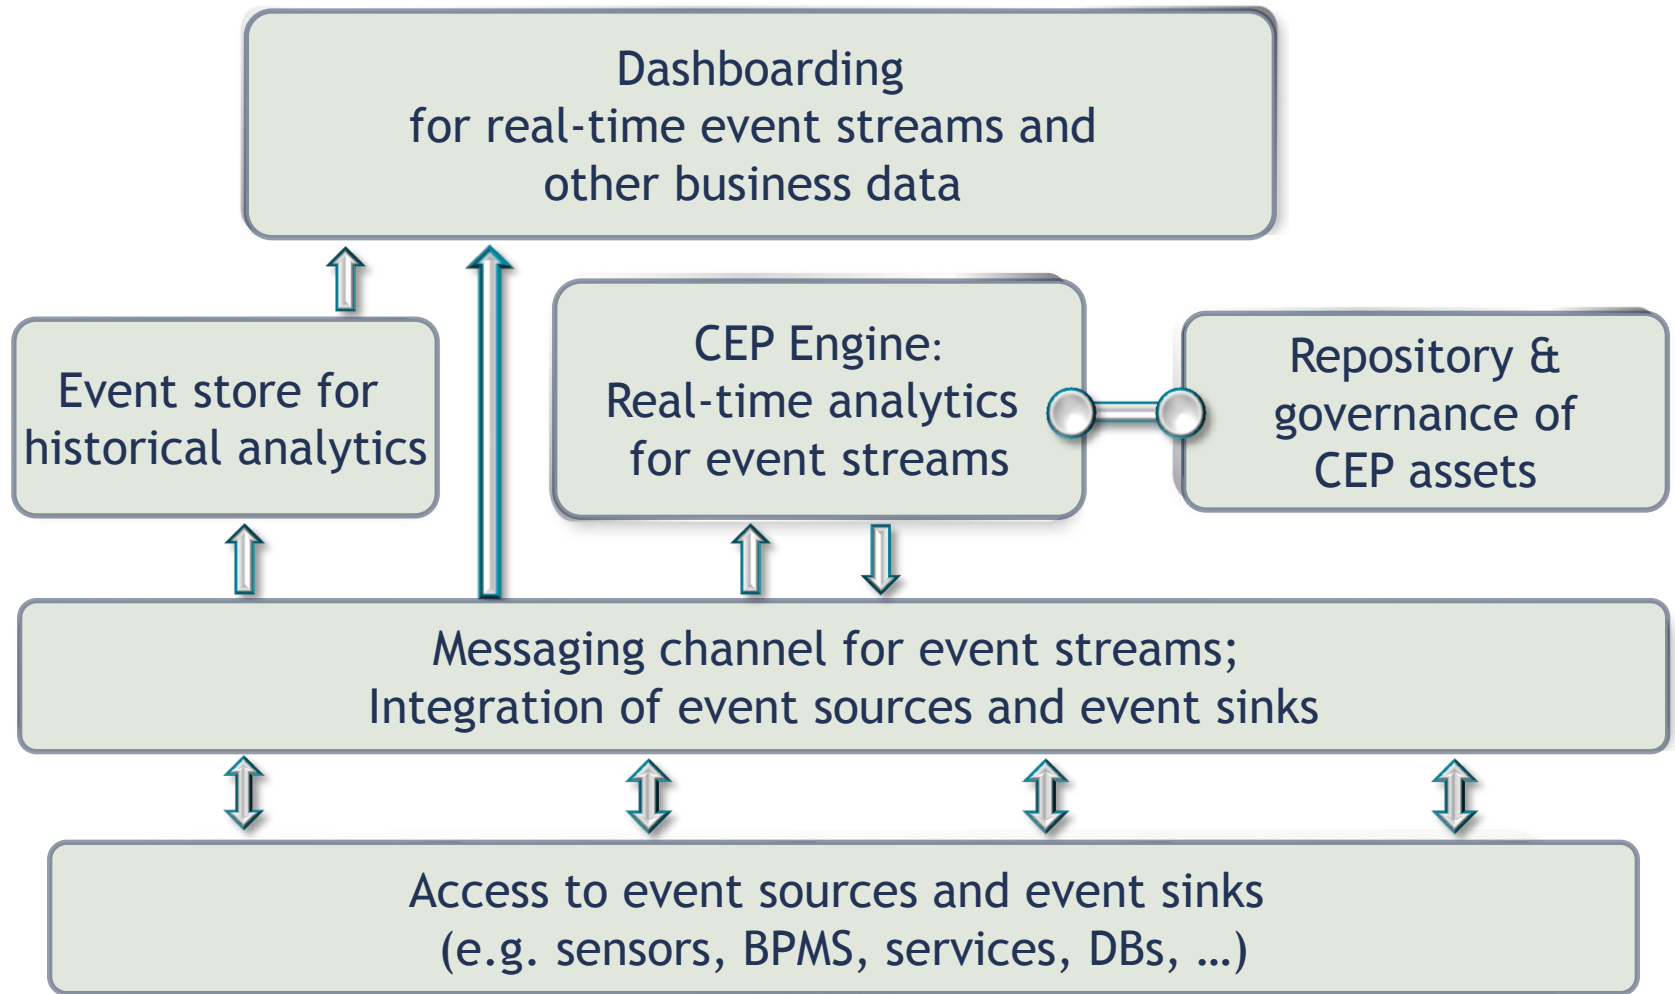
Complex Event Processing Engine

Output streams

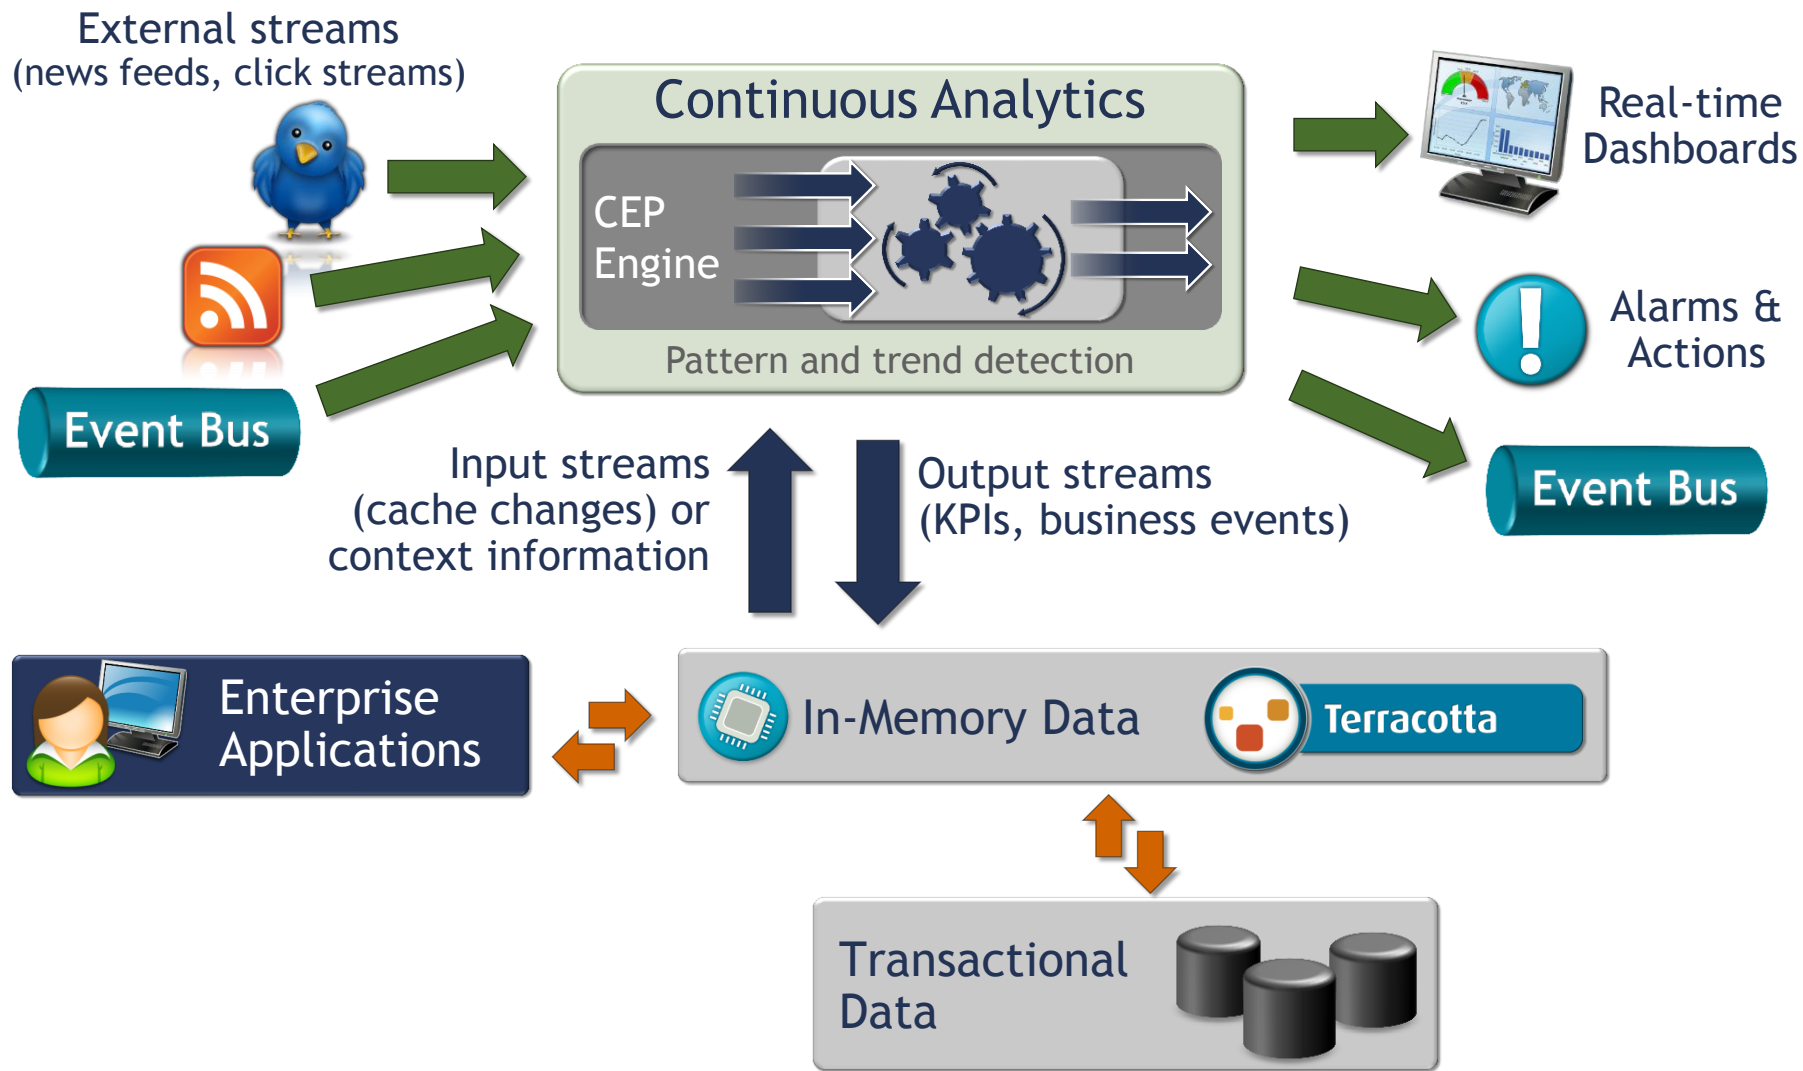


**In-memory, on-the-fly analysis!**

# Software AG's CEP Platform: webMethods Business Events



# Real-time analytics for Big Data in motion



# Basic idea of continuous cache queries

## EHCACHE

key		value	
id	customer	region	revenue
1	AB AG	EMEA	1.000.000 €
4	CD GmbH	NA	2.000.000 €
6	EF Ltd	EMEA	30.000 €
8	AB AG	EMEA	3.770.000 €

Continuous queries:


- revenue ≥ 1.000.000 €
- sum(revenue), avg(revenue)  
group by region
- sum(revenue), avg(revenue)



key		value	
id	customer	region	revenue
1	AB AG	EMEA	1.000.000 €
4	CD GmbH	NA	2.000.000 €
8	AB AG	EMEA	3.770.000 €

key		value	
region	sum	avg	
EMEA	4.800.000 €	1.600.000 €	
NA	2.000.000 €	2.000.000 €	

key		value	
agg	sum	avg	
AVG	6.800.000 €	1.700.000 €	



Ad hoc query result			
id	customer	region	revenue
1	AB AG	EMEA	1.000.000 €
4	CD GmbH	NA	2.000.000 €
8	AB AG	EMEA	3.770.000 €

Ad hoc query:  
revenue ≥ 1.000.000 €

## Continuous Cache Queries

- Use Cache change events
- Maintain result cache
  - Always up to date
  - Access without additional costs/delay

# Technical demo

The screenshot displays the CacheMonitor application interface, which is divided into two main sections: TERRACOTTA (Ehcache) on the left and RTM (CEP-Engine) on the right.

**TERRACOTTA (Ehcache) Section:**

- Data-Generator:** A vertical label on the far left indicates the source of the data.
- sourceCache:** A table showing a list of key-value pairs. The keys are IDs (e.g., id\_170634) and the values are names and ages (e.g., DOROTHY JONES (f, 72)).
- FilterQuery:** A table showing the result of a filter query. The values are counts (e.g., 74, 37, 68).
- AggregationQuery:** A table showing the result of an aggregation query. The value is a list of aggregates: [11568, 573684, 20, 79].
- Annotations:** The text "original cache" is overlaid on the sourceCache table, and "ad-hoc filter query" is overlaid on the FilterQuery table.

**RTM (CEP-Engine) Section:**

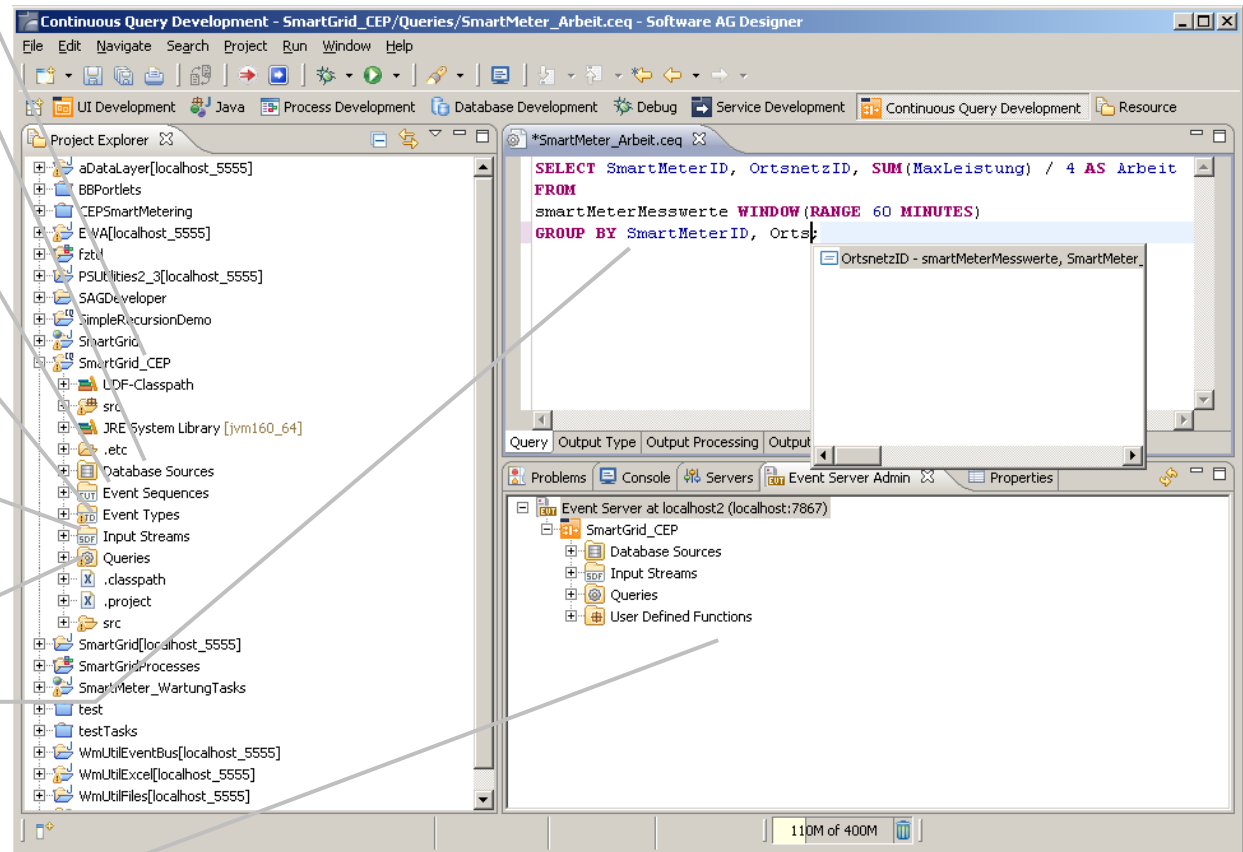
- Continuous\_Search\_Query1:** A table showing a continuous search query result. The values are lists of key-value pairs (e.g., [id\_170634, DOROTHY JONES (f, 72), ...]).
- Continuous\_SQL\_Query1:** A table showing a continuous SQL query result. The values are lists of key-value pairs (e.g., [id\_170634, [id\_170634, DOROTHY JONES (f, 72), 72]).
- Continuous\_Search\_Query2:** A table showing a continuous search aggregation query result. The values are lists of key-value pairs (e.g., [MAX(age), 79]).
- Continuous\_SQL\_Query2:** A table showing a continuous SQL aggregation query result. The values are lists of key-value pairs (e.g., [MAX(age), 79]).
- Annotations:** The text "continuous filter query" is overlaid on Continuous\_Search\_Query1, and "continuous filter query (SQL)" is overlaid on Continuous\_SQL\_Query1. Similarly, "continuous aggregation query" is overlaid on Continuous\_Search\_Query2, and "continuous aggregation query (SQL)" is overlaid on Continuous\_SQL\_Query2.

The interface includes a top toolbar with navigation buttons and a status bar at the bottom of each window showing query execution details like "Run query" time and "Results".

# CEP Projects in Software AG Designer

Single Eclipse-based development environment for all webMethods products

- Project definition
- Database integration
- Test data
- Event Type definition
- Event Streams
- Business logic
- Editor for SQL-style queries
- Testing & Debugging



# Business logic via SQL

## Ease of use

- Standardized, well-established query language
- Same query language for event streams and historic data (relations)
- ➔ Lots of developers are familiar with SQL
- No custom coding, no proprietary languages
- Implicit handling of time concepts
- Update of business logic at runtime
- ➔ Higher reusability, flexibility, and quality
- ➔ Reduced development times



## Deterministic and sound semantics (based on application time)

- Snapshot-compliant with databases
- Enables query optimization for continuous queries


# Reproducible and traceable results

## Issues for well-defined query semantics

- Notion of time
- Windowing constructs
- Schema stability
- Continuous output
- Stream ordering
- Determinism
- Optimizations

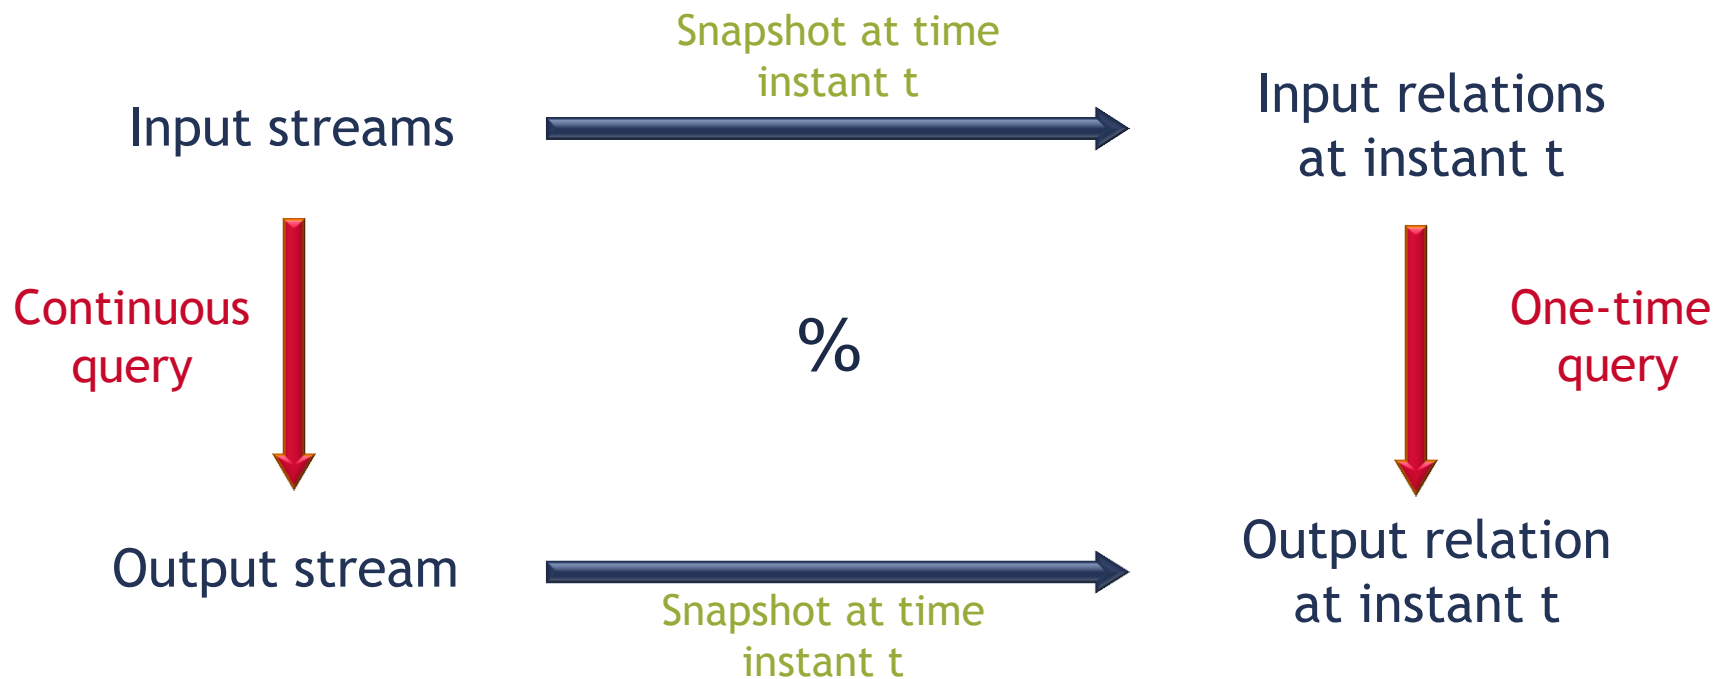
Semantics must not depend on system internals!

Key Definition? ?





## Semantic compliance with databases



Exact specification of query results for any point in time

- DBS would produce identical results if applied to every single time instant

All conventional transformation rules applicable due to snapshot reducibility

- ➔ Powerful query optimizations applicable

## Event processing - design patterns

### Basic

- Filtering
- Transformation
- Merge
- Correlation
- Aggregation
- Grouping
- Enrichment

**EXACT  
ANSWERS**

### Advanced

- Duplicate elimination
- Difference
- Pattern matching
- State machines
- Non-event detection
- Set memberships
- User-defined functions
- Recursion
- Database lookups & writes

### Temporal

- Time-based windows
- Count-based windows
- Partitioned windows
- Control over window advance
- Control over time granularity
- Temporal patterns

### Stream Mining

- Dynamic forecasting
- Outlier detection
- Density estimation
- Cluster detection

**APPROXIMATE  
ANSWERS**

## Event engine - design patterns

### Operator-level

- ✓ Push-based implementations
  - Time-interval algebra
  - Positive-negative algebra
  - Punctuations
- ✓ Heartbeats
- ✓ In-memory caching & indexing
- ✓ Pub/Sub mechanism
- ✓ Pull-based implementations for request-respond queries
- ✓ User-defined functions & operators

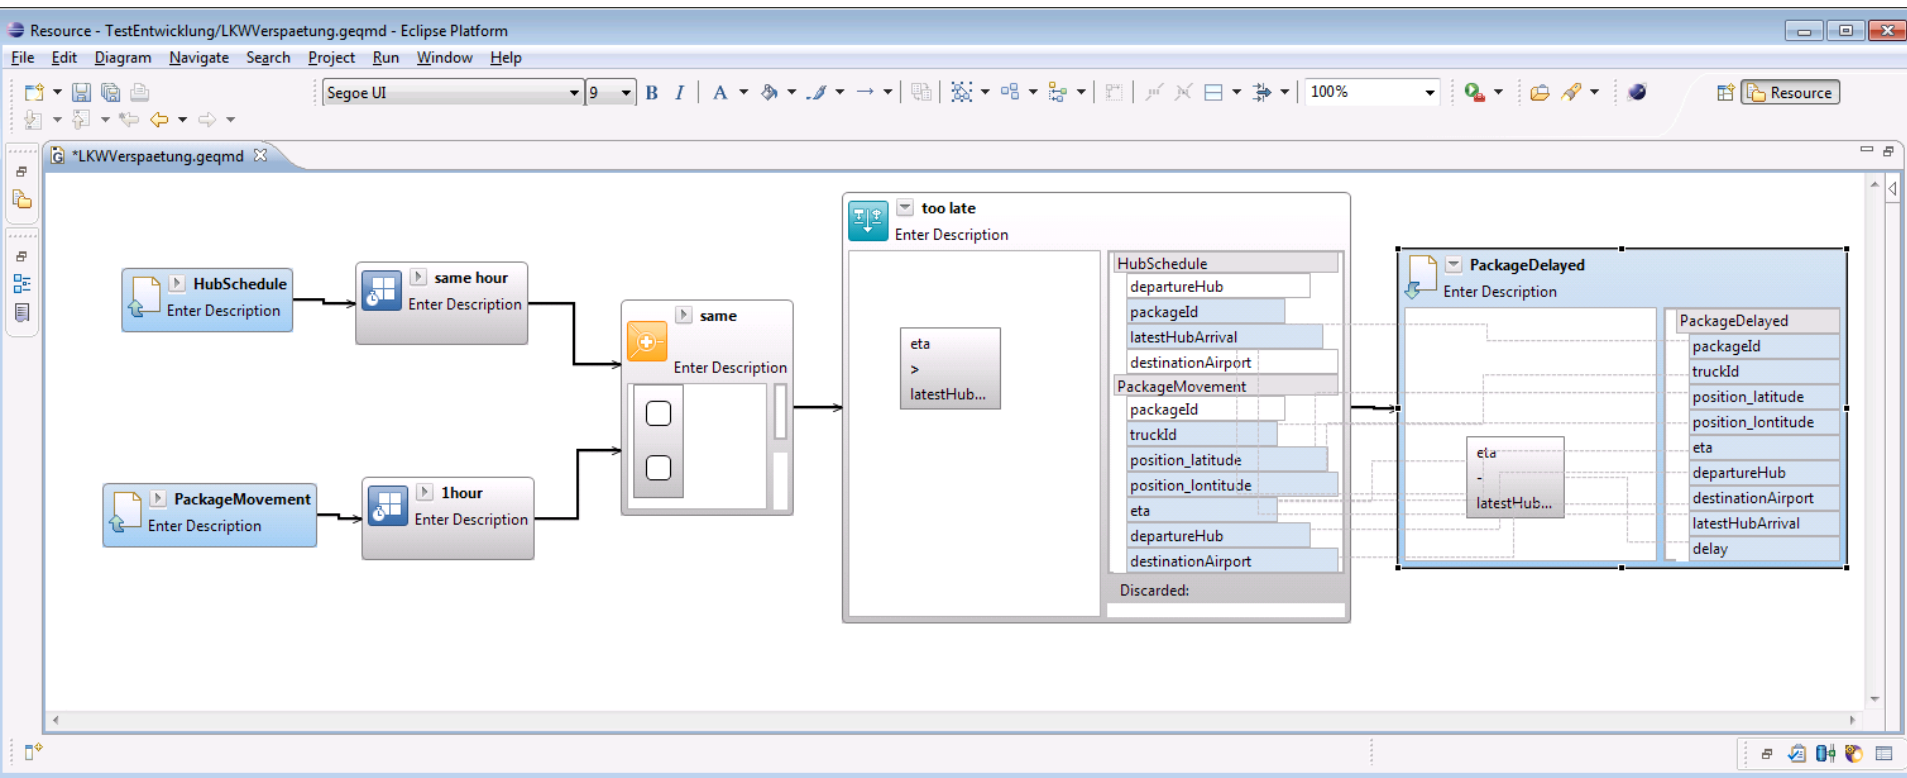
### Query-level

- ✓ Dynamic registration
- ✓ Semantic subscriptions
- ✓ Authentication & authorization

### System-level

- ✓ Application time
- ✓ Simulation & replay
- ✓ Query batches
- ✓ Metadata management
  - Schema information
  - Runtime statistics
- ✓ Multi-threading
- ✓ High availability
- ✓ Query optimization
  - Subquery sharing
  - Cost model
- ✓ Distributed execution
- ✓ Fine-granular testing & debugging
- ✓ webMethods Deployer
- ✓ SAG Designer

# Graphical Query Modeler in next release



# Mashup Dashboards for Live Insights

Software AG  
as data source

webMethods

Adabas

Natural



Public data



Desktop data

Department data

Enterprise data

Spreadsheets stored locally



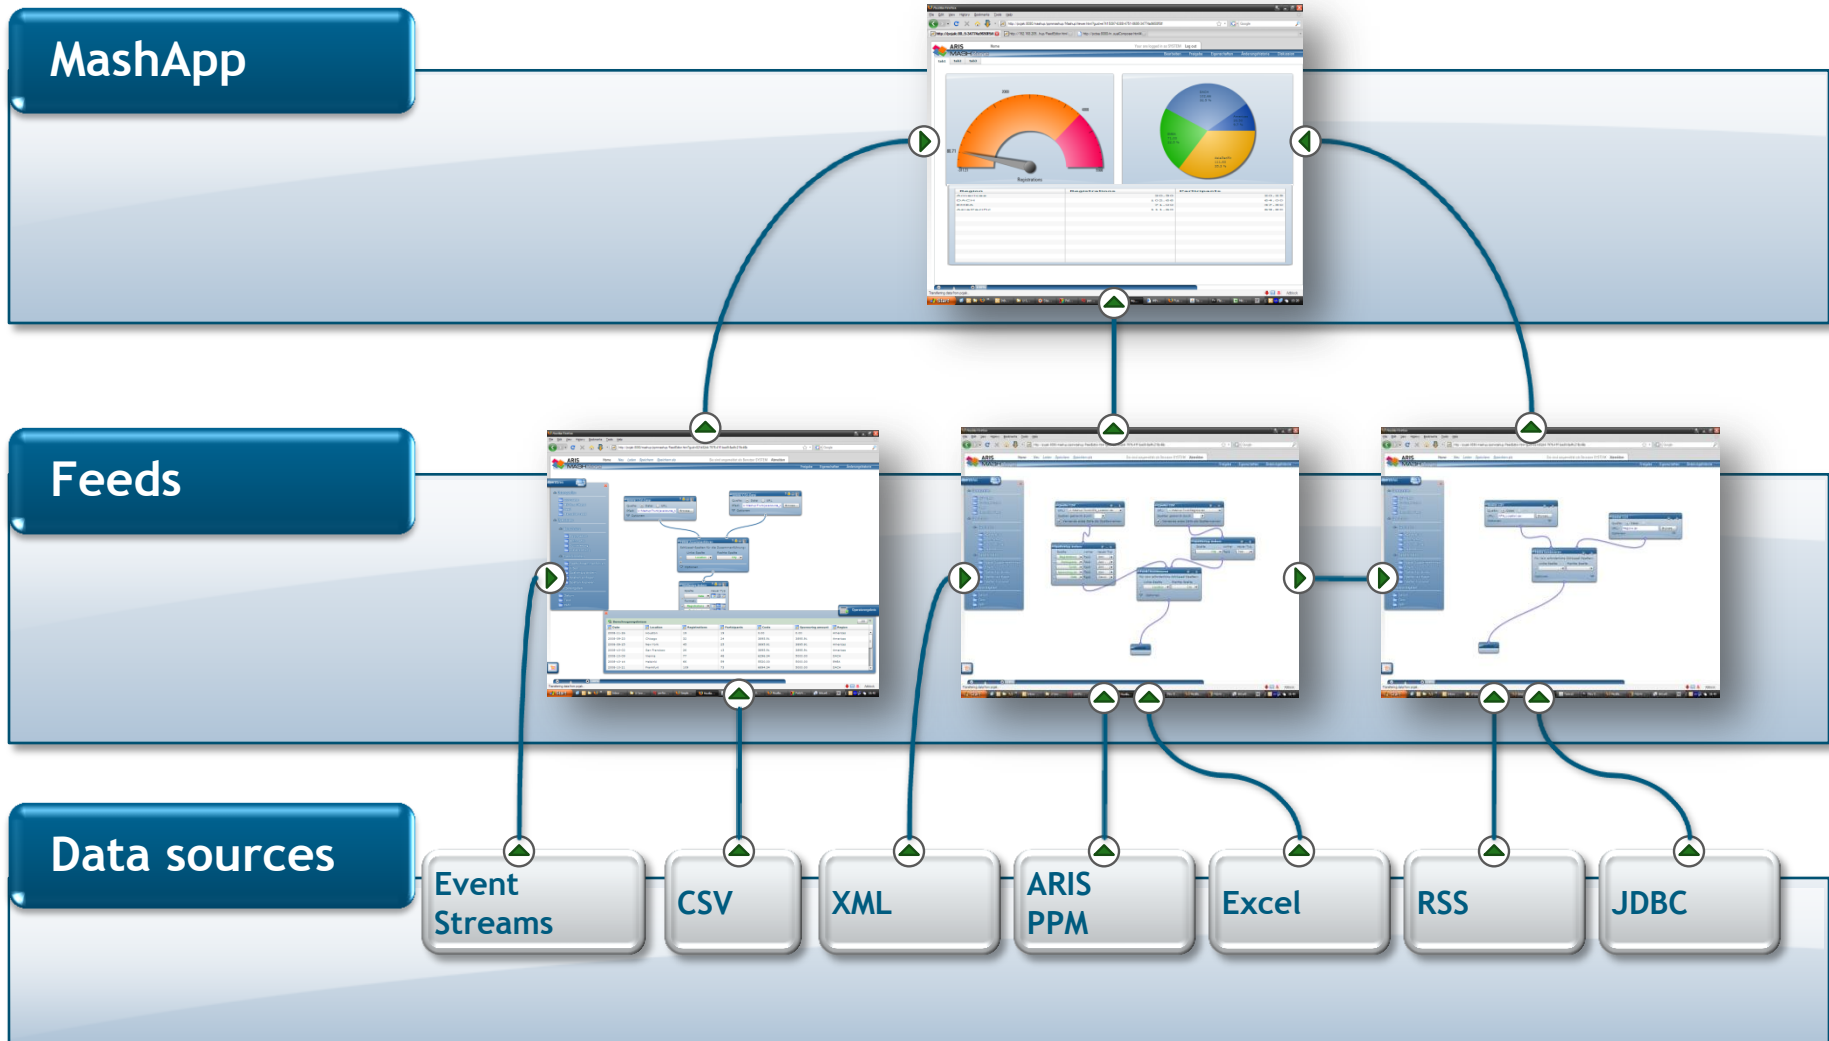
Queries and reports from ERP, DWH, etc.



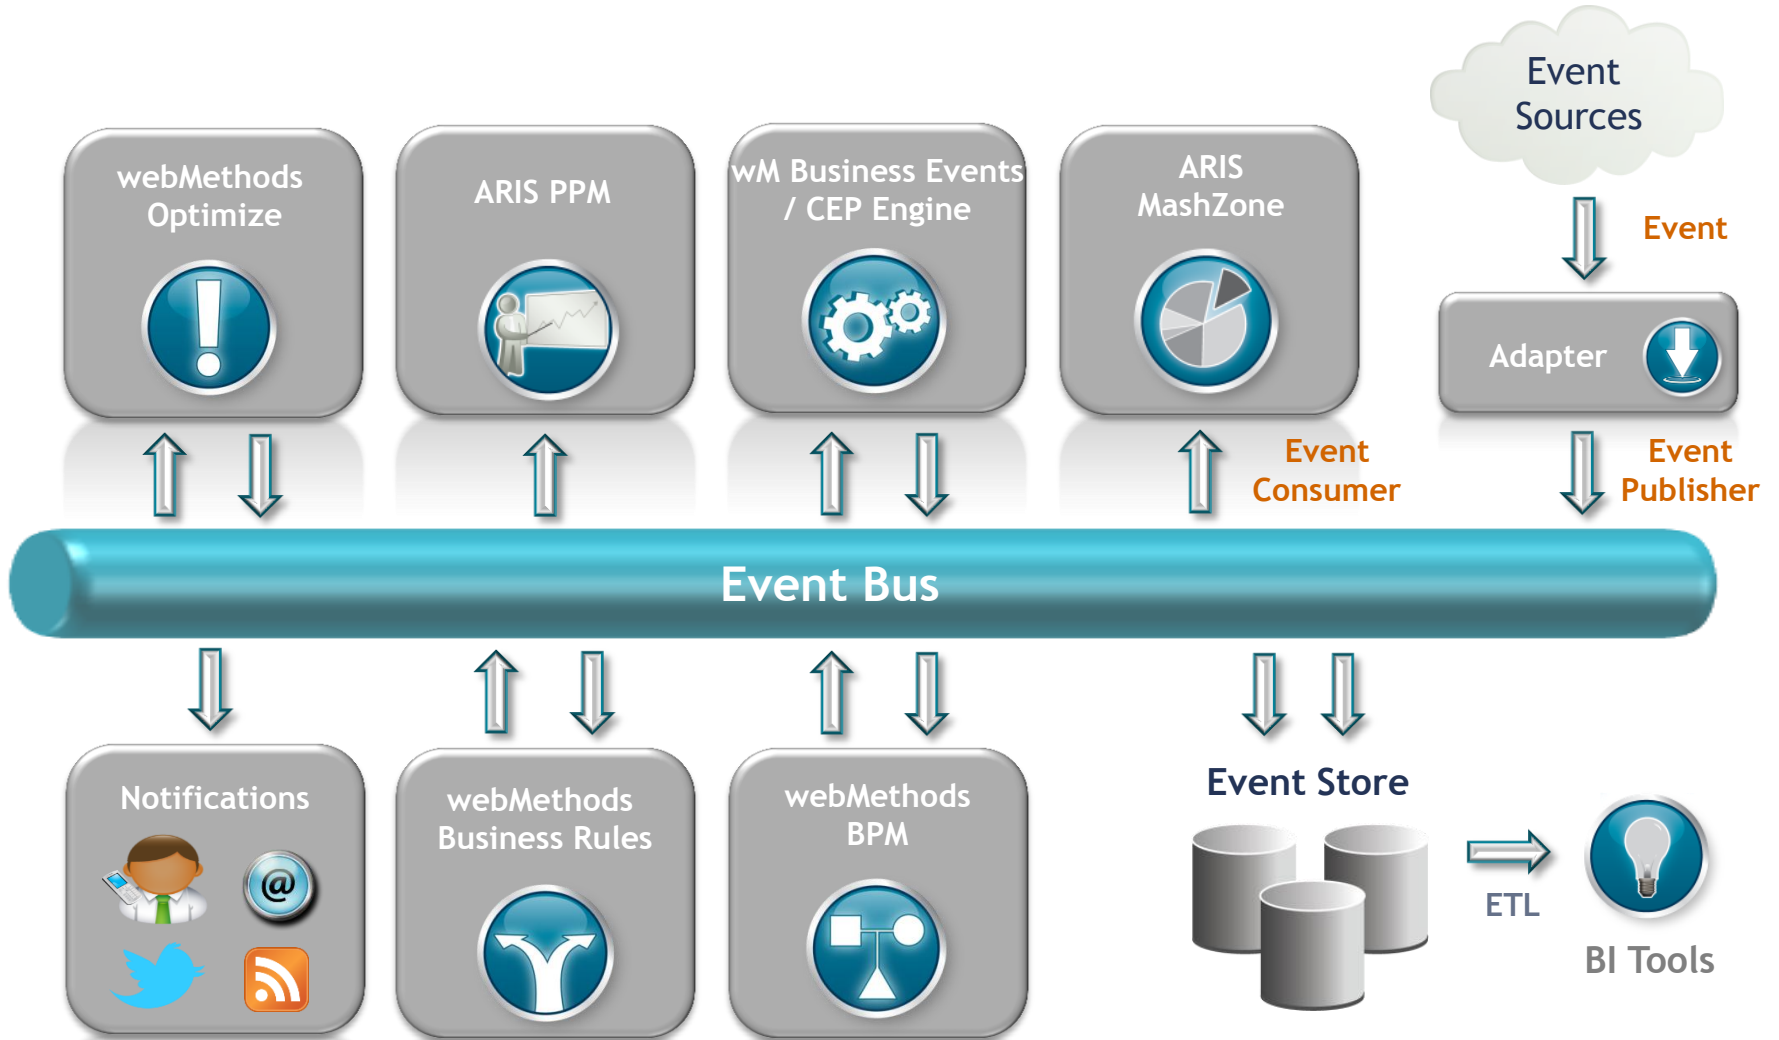
CRM reports, forecasts



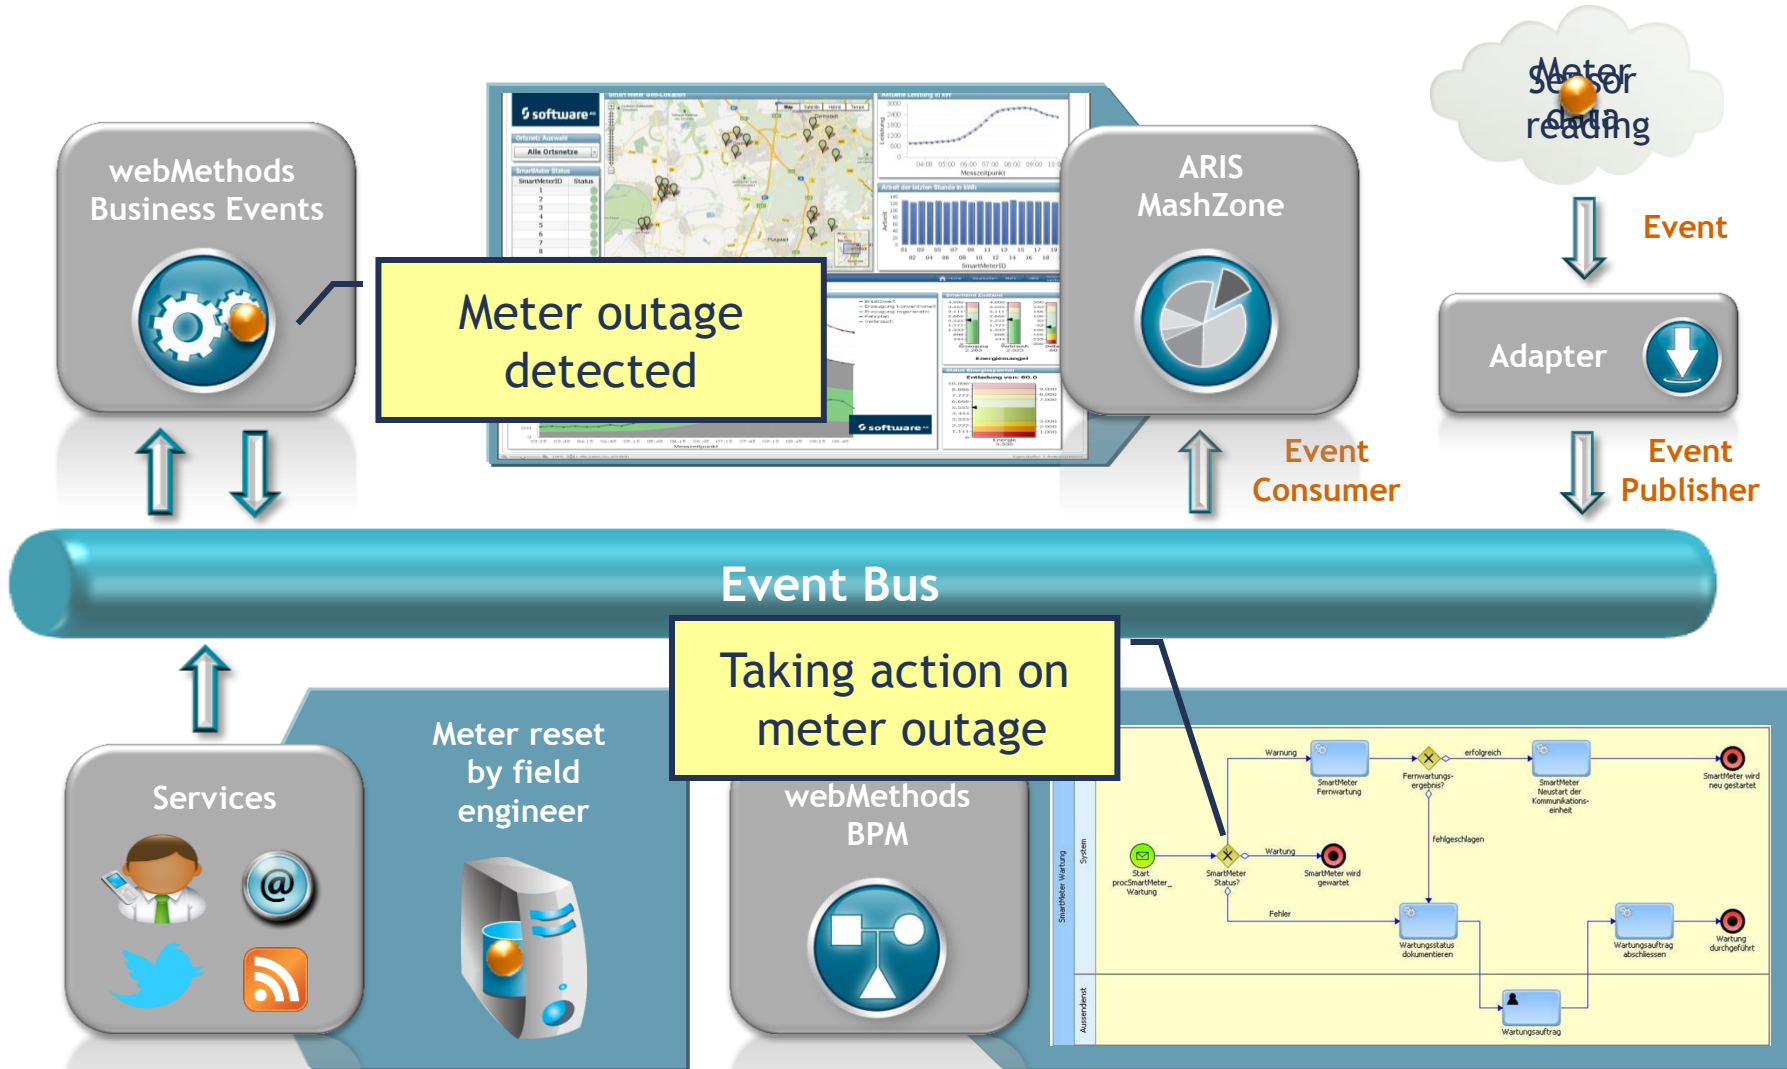
# ARIS MashZone Creates Transparency



# Event Driven Architecture (EDA)



# Transform Real-time Insights into Immediate Actions






# Intelligent Business Operations: Next Best Action

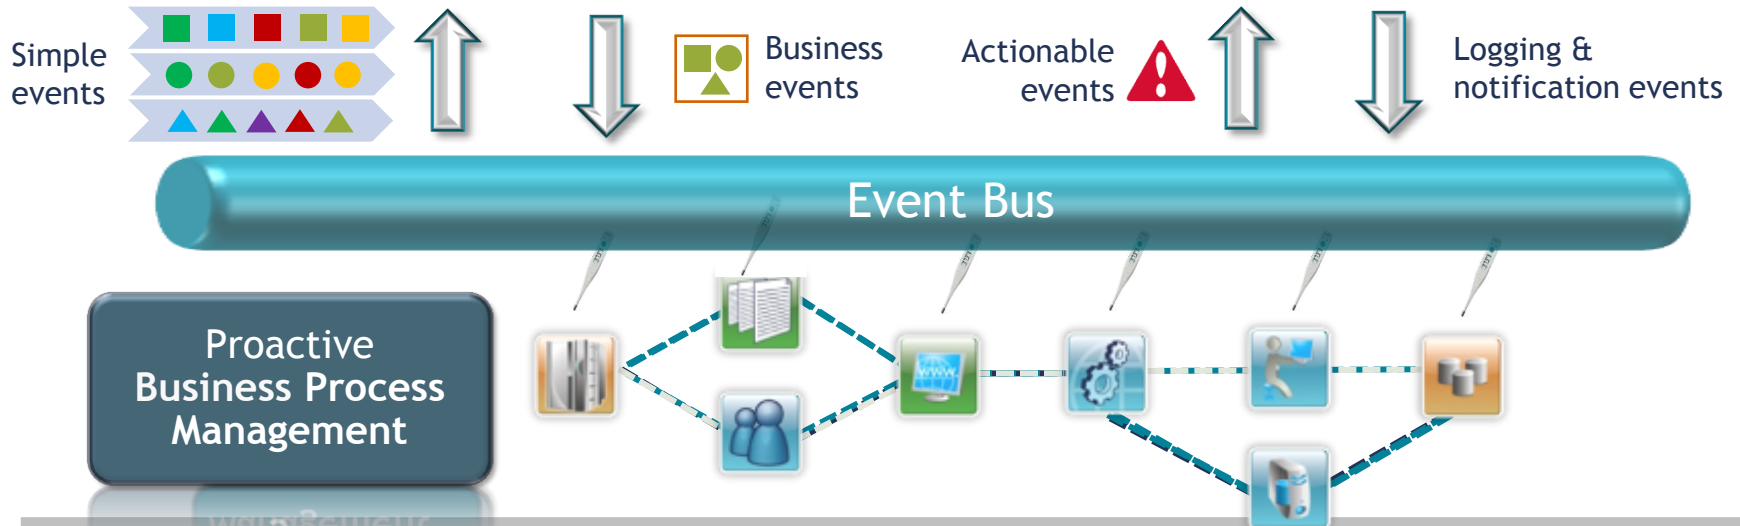
## 1) Situational Awareness

### webMethods Business Events

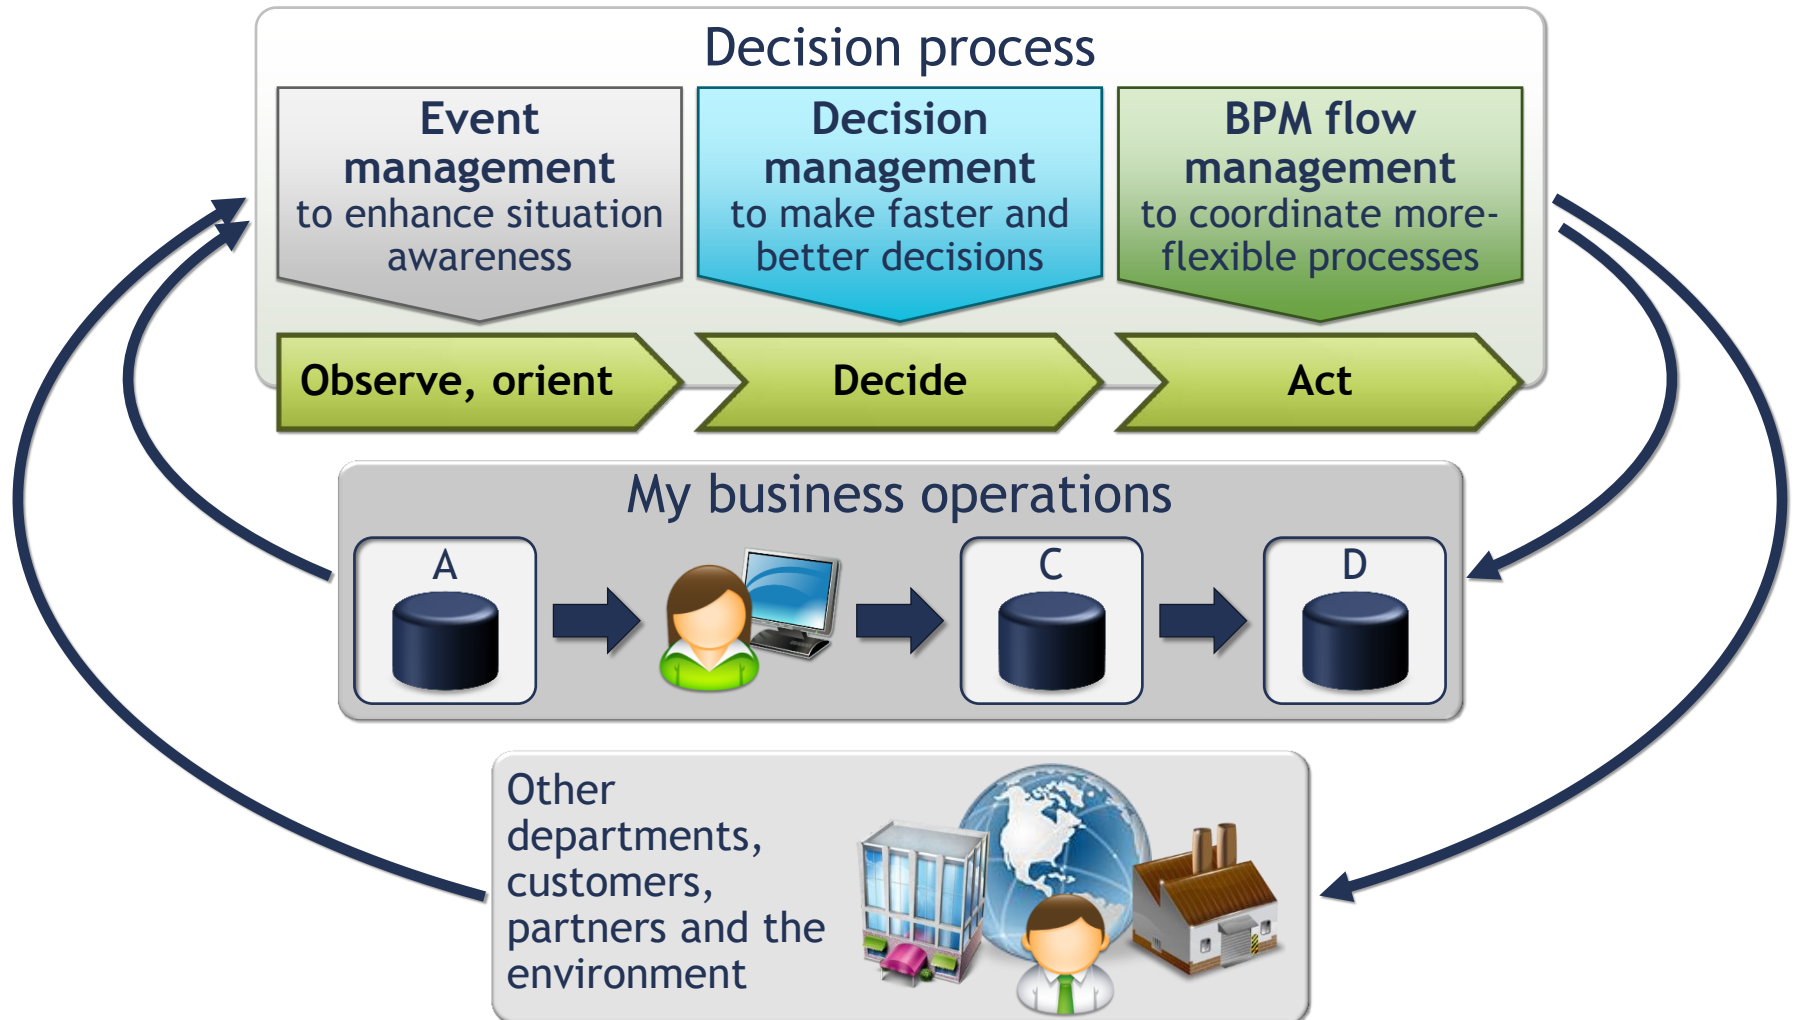


## 2) Immediate Actions

### webMethods Business Rules



# Intelligent Business Operations (IBO)



Source: W. Roy Schulte: *Apply Three Disciplines to Make Business Operations More Intelligent*, Gartner Inc., March 2012

## Business Impact of CEP

- Improves the quality of decision making by presenting information that would otherwise be overlooked.
- Enables faster response to threats and opportunities.
- Helps shield business people from data overload by implementing a management-by exception policy.
- Reduces the cost of manually processing the growing volume of event data in business.

**CEP is one of the pillars of the emerging "big data" movement.**

*Gartner Hype Cycle for Data Management 2011*

Source: Gartner Inc.: W. Roy Schulte; Zarko Sumic: „Hype Cycle for Data Management”, 2011.

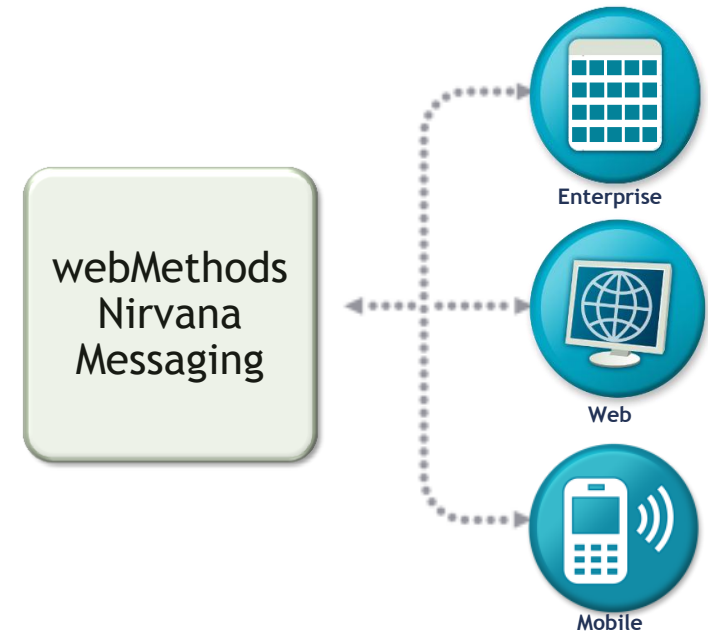
# High-Speed Messaging

**Get There Faster.™**

# webMethods Nirvana Messaging

## High Speed Universal Messaging

- A single end-to-end solution for the delivery of real-time data
  - Across multiple delivery channels (Enterprise, Web, Mobile)
  - Between multiple languages and platforms
- Supporting an unrivalled selection of languages, messaging standards and protocols using common APIs
- Proven track record in the most demanding markets
  - Our customers deliver real-time data to tens of thousands of connected users every day



# Benefits of Nirvana Universal Messaging

## SIMPLICITY

Common APIs for all delivery channels  
One product - multiple use cases

## SCALABILITY & PERFORMANCE

No bridging between components  
Data travelling over fewer hops

## FLEXIBILITY

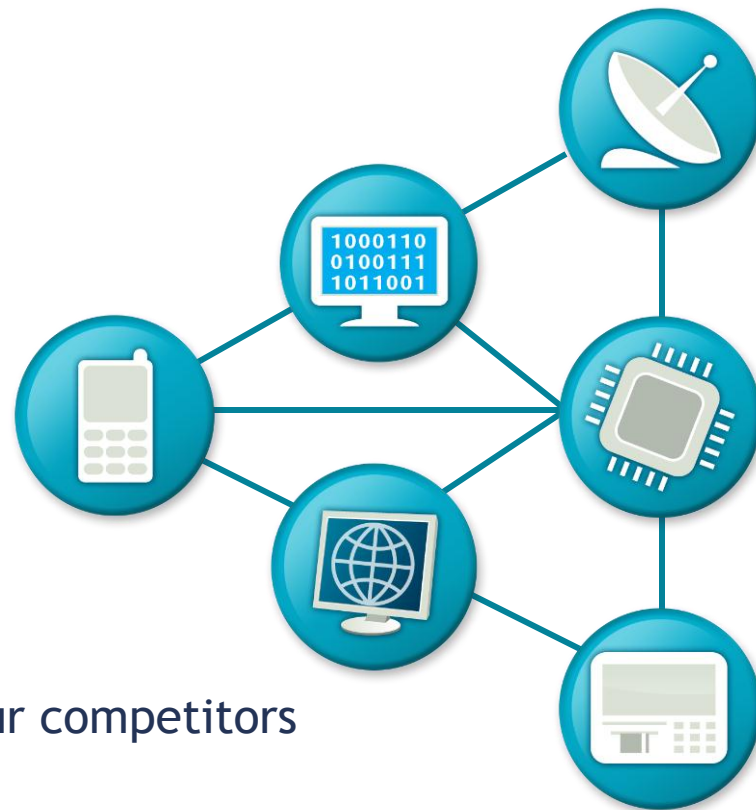
Support for all your delivery channels

## FUTURE PROOFING

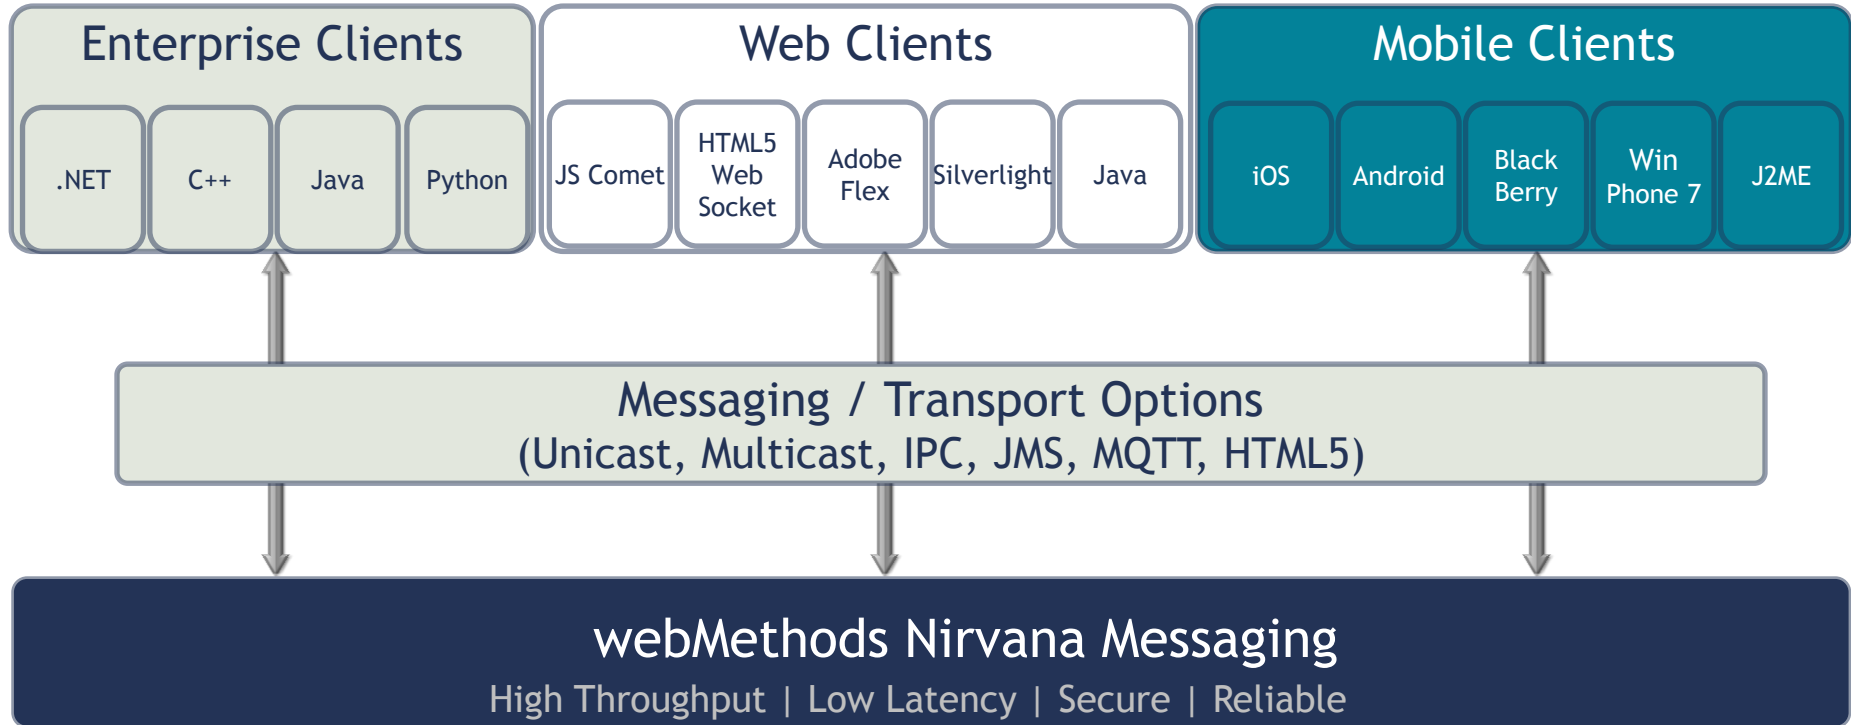
Ready for new delivery channels before your competitors

## VALUE

Less vendors, less cost - hardware, software, operational savings



# Universal Messaging Today



# webMethods Nirvana Messaging Customers

## Customer Successes

- Global customer base
- Powers 2 of the top 3 market-leading Foreign Exchange platforms
- Increasing reach beyond financial services

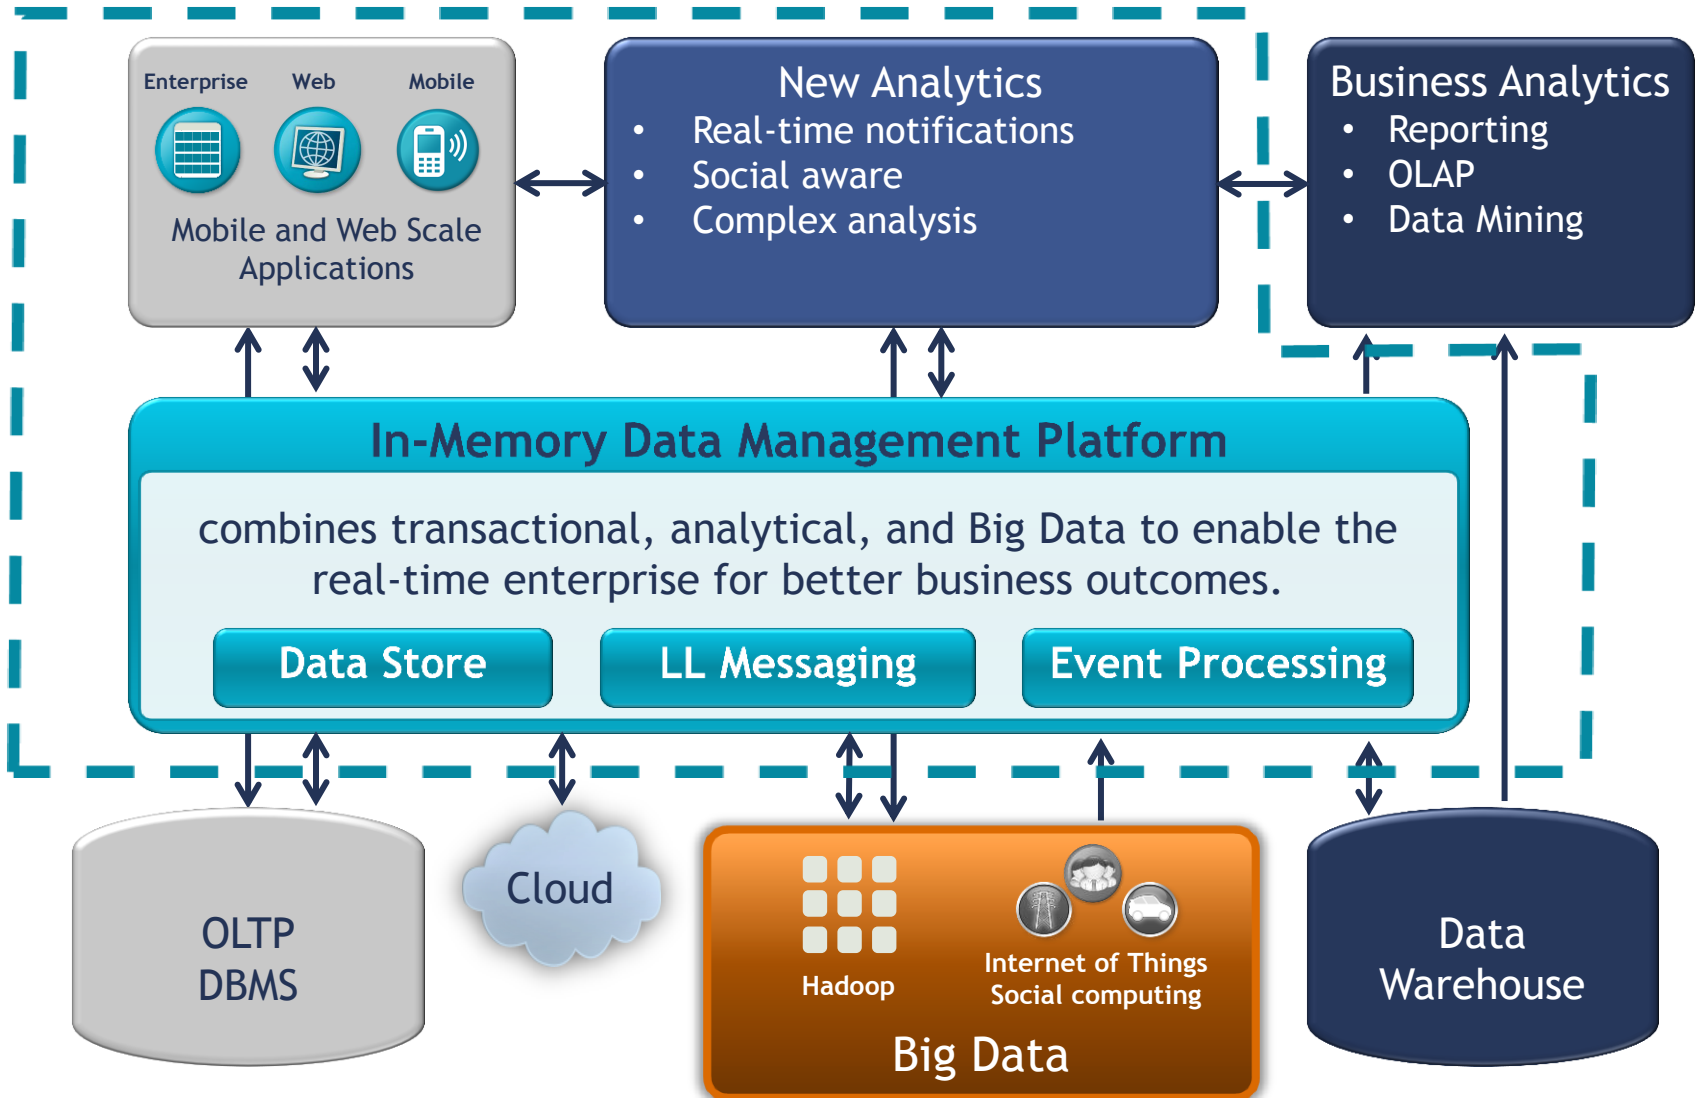
## Global Customer Base





# In-Memory Data Management Platform

**Get There Faster.<sup>TM</sup>**



## Key Advantages

- Leverage enterprise-proven and high-performance in-memory technology (scale-up & -out)
- Process data in real-time (move, manage, analyze)
- Use commodity hardware and software
- Preserve existing business data and data stores
- Implement in modular fashion

**Thank you!**

**Get There Faster.<sup>TM</sup>**

Contact: [juergen.kraemer@softwareag.com](mailto:juergen.kraemer@softwareag.com)