Big Data Platforms, Tools, and Research at IBM

Ed Pednault
CTO, Scalable Analytics
Business Analytics and Mathematical Sciences, IBM Research

Please Note:

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

IBM Big Data Strategy: Move the Analytics Closer to the Data

New analytic applications drive the requirements for a big data platform

- Integrate and manage the full variety, velocity and volume of data
- Apply advanced analytics to information in its native form
- Visualize all available data for ad-hoc analysis
- Development environment for building new analytic applications
- Workload optimization and scheduling
- Security and Governance

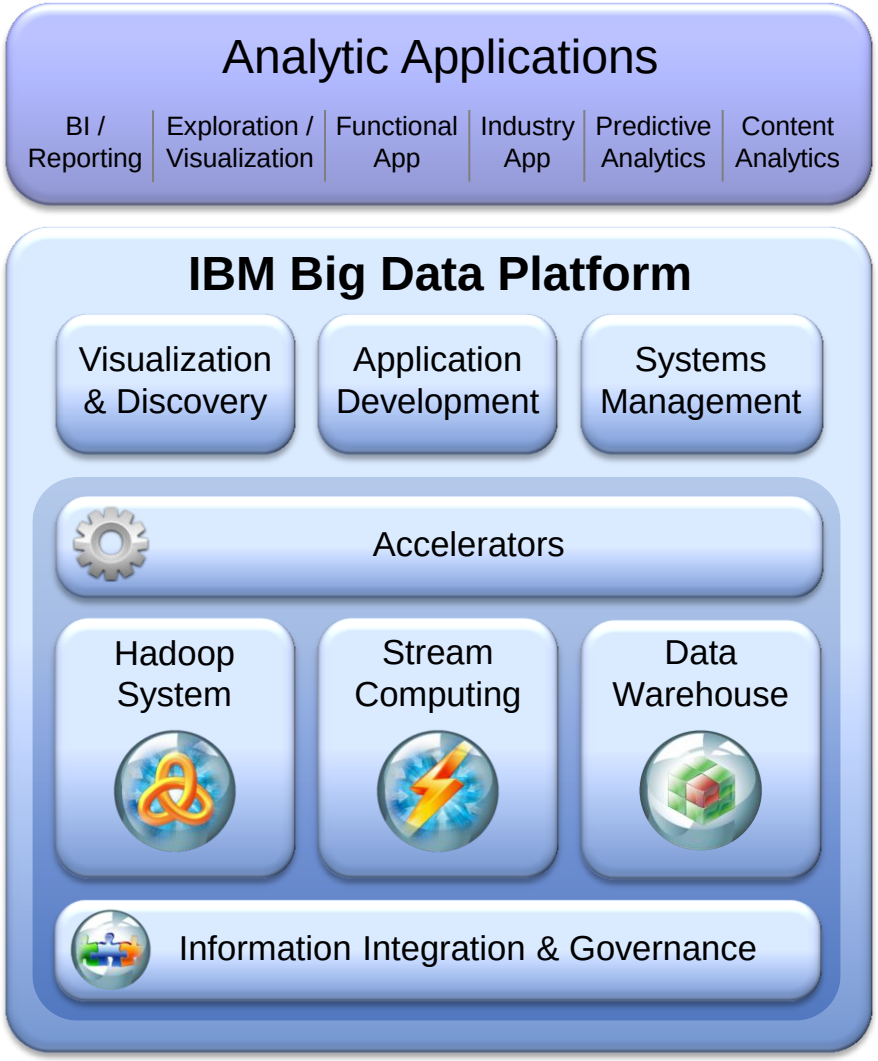

Big Data Platform - Hadoop System

InfoSphere BigInsights

- Augments open source Hadoop with enterprise capabilities
  - Enterprise-class storage
  - Security
  - Performance Optimization
  - Enterprise integration
  - Development tooling
  - Analytic Accelerators
  - Application and industry accelerators
  - Visualization

Enterprise-Class Storage and Security

- IBM GPFS-SNC (Shared-Nothing Cluster) parallel file system can replace HDFS to provide Enterprise-ready storage
  - Better performance
  - Better availability
    - No single point of failure
  - Better management
    - Full POSIX compliance, supports multiple storage technologies
  - Better security
    - Kernel-level file system that can exploits OS-level security
- Security provided by reducing the surface area and securing access to administrative interfaces and key Hadoop services
  - LDAP authentication and reverse-proxy support restricts access to authorized users
    - Clients outside the cluster must use REST HTTP access
  - Defines 4 roles not available in Hadoop for finer grained security:
    - System Admin, Data Admin, Application Admin, and User
    - Installer automatically lets you map these roles to LDAP groups and users
  - GPFS-SNC means the cluster is aware of the underlying OS security services without added complexity

Workload Optimization

Optimized performance for big data analytic workloads

Adaptive MapReduce

- Algorithm to optimize execution time of multiple small and large jobs
- Performance gains of 30% reduce overhead of task startup

Hadoop System Scheduler

- Identifies small and large jobs from prior experience
- Sequences work to reduce overhead

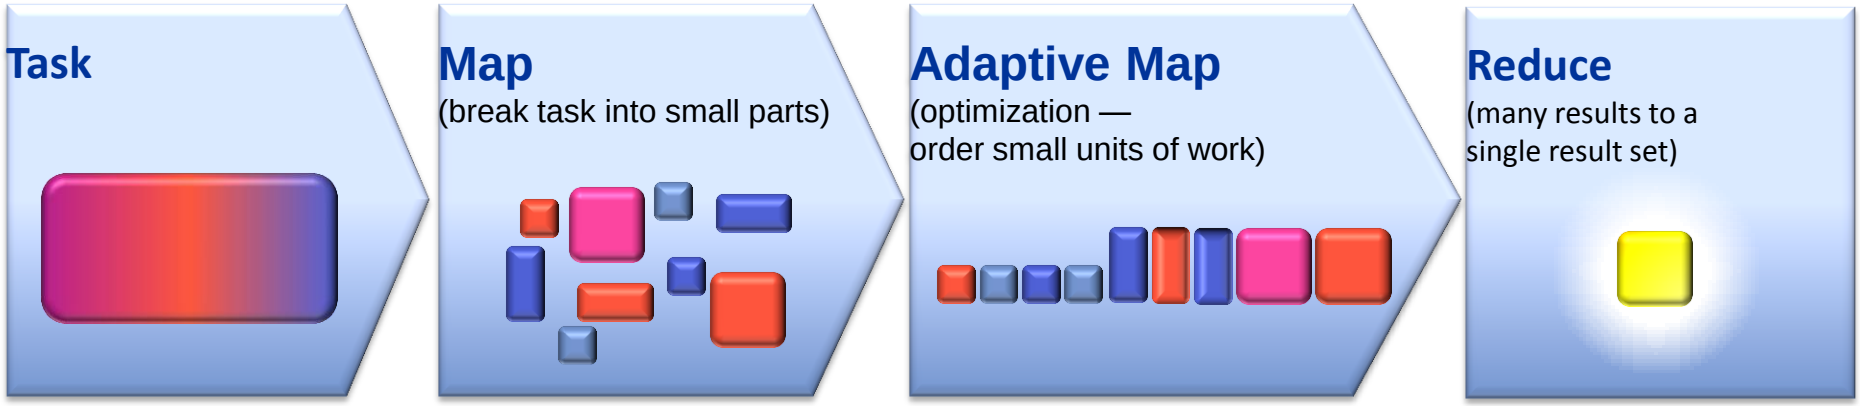

Big Data Platform - Stream Computing

- Built to analyze data in motion
  - Multiple concurrent input streams
  - Massive scalability
- Process and analyze a variety of data
  - Structured, unstructured content, video, audio
  - Advanced analytic operators

InfoSphere Streams exploits massive pipeline parallelism for extreme throughput with low latency

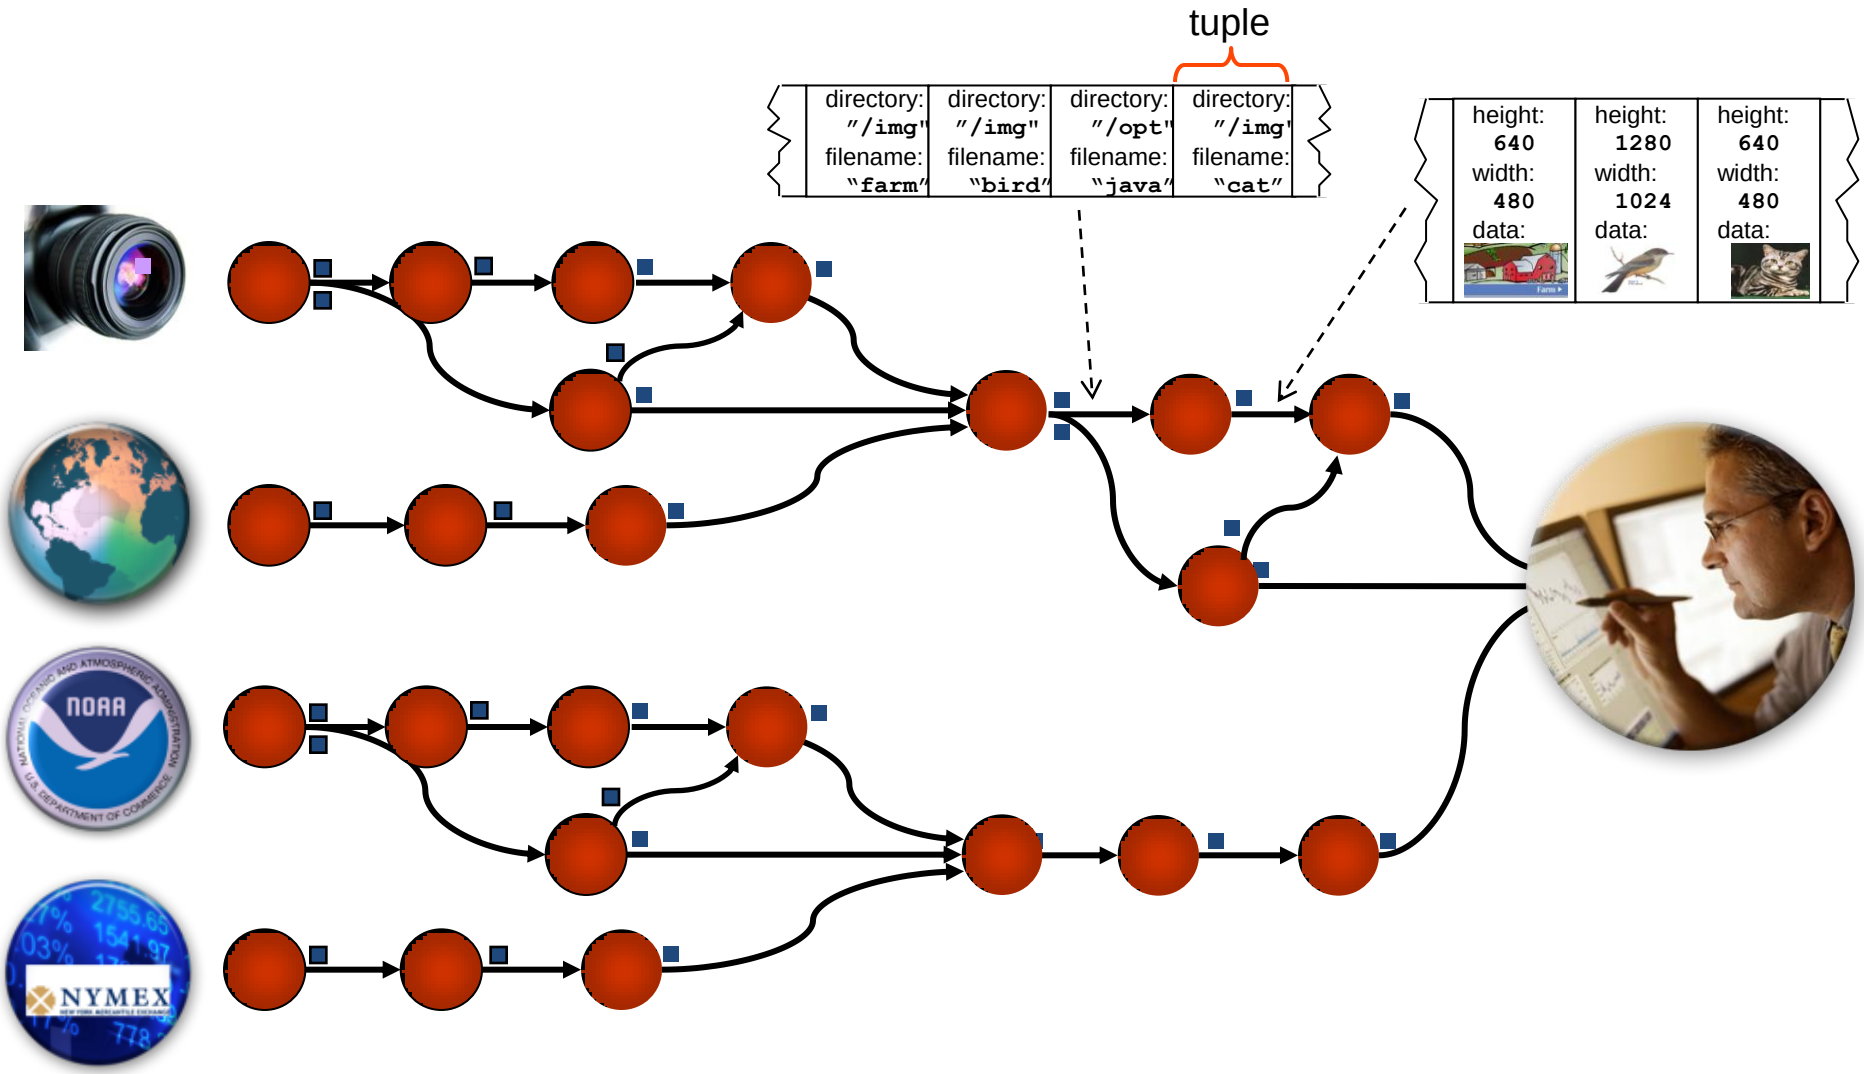

Video analytics

Example – Contour Detection

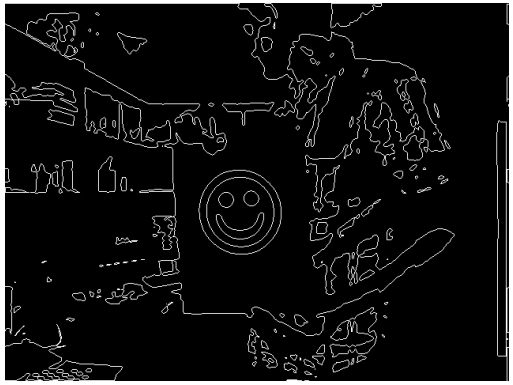

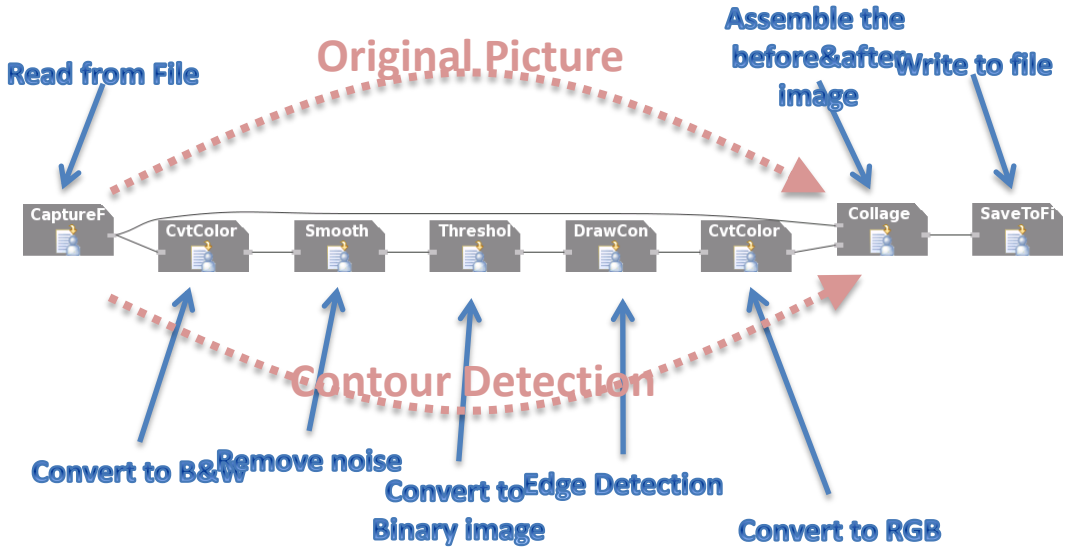


```
[Application]
contours

[Program]

vstream IplImage(channels: Integer,
                 depth: Integer,
                 origin: Integer,
                 width: Integer,
                 height: Integer,
                 data: ByteList)

stream vid(schemaFor(IplImage))
  := CaptureFromFile( ) [file: "$ENV{HOME}/demo3.m4v"; repeat:1] { }
    -> partition["P1"]

stream bw_vid(schemaFor(vid))
  := CvtColor(vid) [ ] { data := ~CV_BGR2GRAY() }
    -> partition["P1"]

stream smooth_bw_vid(schemaFor(vid))
  := Smooth(bw_vid) [iteration: 4] { }
    -> partition["P1"]

stream th_vid(schemaFor(vid))
  := Threshold(smooth_bw_vid) [max:255; threshold:125] { data := ~CV_THRESH_TOZERO( ) }
    -> partition["P1"]

stream cntr(schemaFor(vid))
  := DrawContours(th_vid) [ ] { }
    -> partition["P1"]

stream cntr_rgb(schemaFor(vid))
  := CvtColor(cntr) [channels:3] { data := ~CV_GRAY2BGR( ) }
    -> partition["P1"]

stream src_n_cntr(schemaFor(vid))
 := Collage(vid; cntr_rgb) [ ] { }
    -> partition["P1"]

Nil
  := SaveToFile(src_n_cntr) [filename: "$ENV{PWD}/demo3-out.mp4"; rate:15; fourcc:"fmp4"] { }
   -> partition["P1"]
```

Why InfoSphere Streams?

- **Scalability**
- **Composable with other analytic streams**

- **Use B&W+threshold pictures to compute derivatives of pixels**
- **Used as a first step for other more sophisticated processing**
- **Very low overhead from Streams – pass 200-300 fps per core – once analysis added processing overhead is high but can get 30fps on 8 cores**

Big Data Platform - Data Warehousing

- Workload optimized systems
  - Deep analytics appliance
  - Configurable operational analytics appliance
  - Data warehousing software

- Capabilities
  - Massive parallel processing engine
  - High performance OLAP
  - Mixed operational and analytic workloads

Deep Analytics Appliance - Revolutionizing Analytics

Purpose-built analytics appliance

Speed: 10-100x faster than traditional systems

Simplicity: Minimal administration and tuning

Scalability: Peta-scale user data capacity

Smart: High-performance advanced analytics

Netezza is architected for high performance on Business Intelligence (OLAP) workloads

- Designed to processes data it at maximum disk transfer rates
- Queries compiled into C++ and FPGAs to minimize overhead

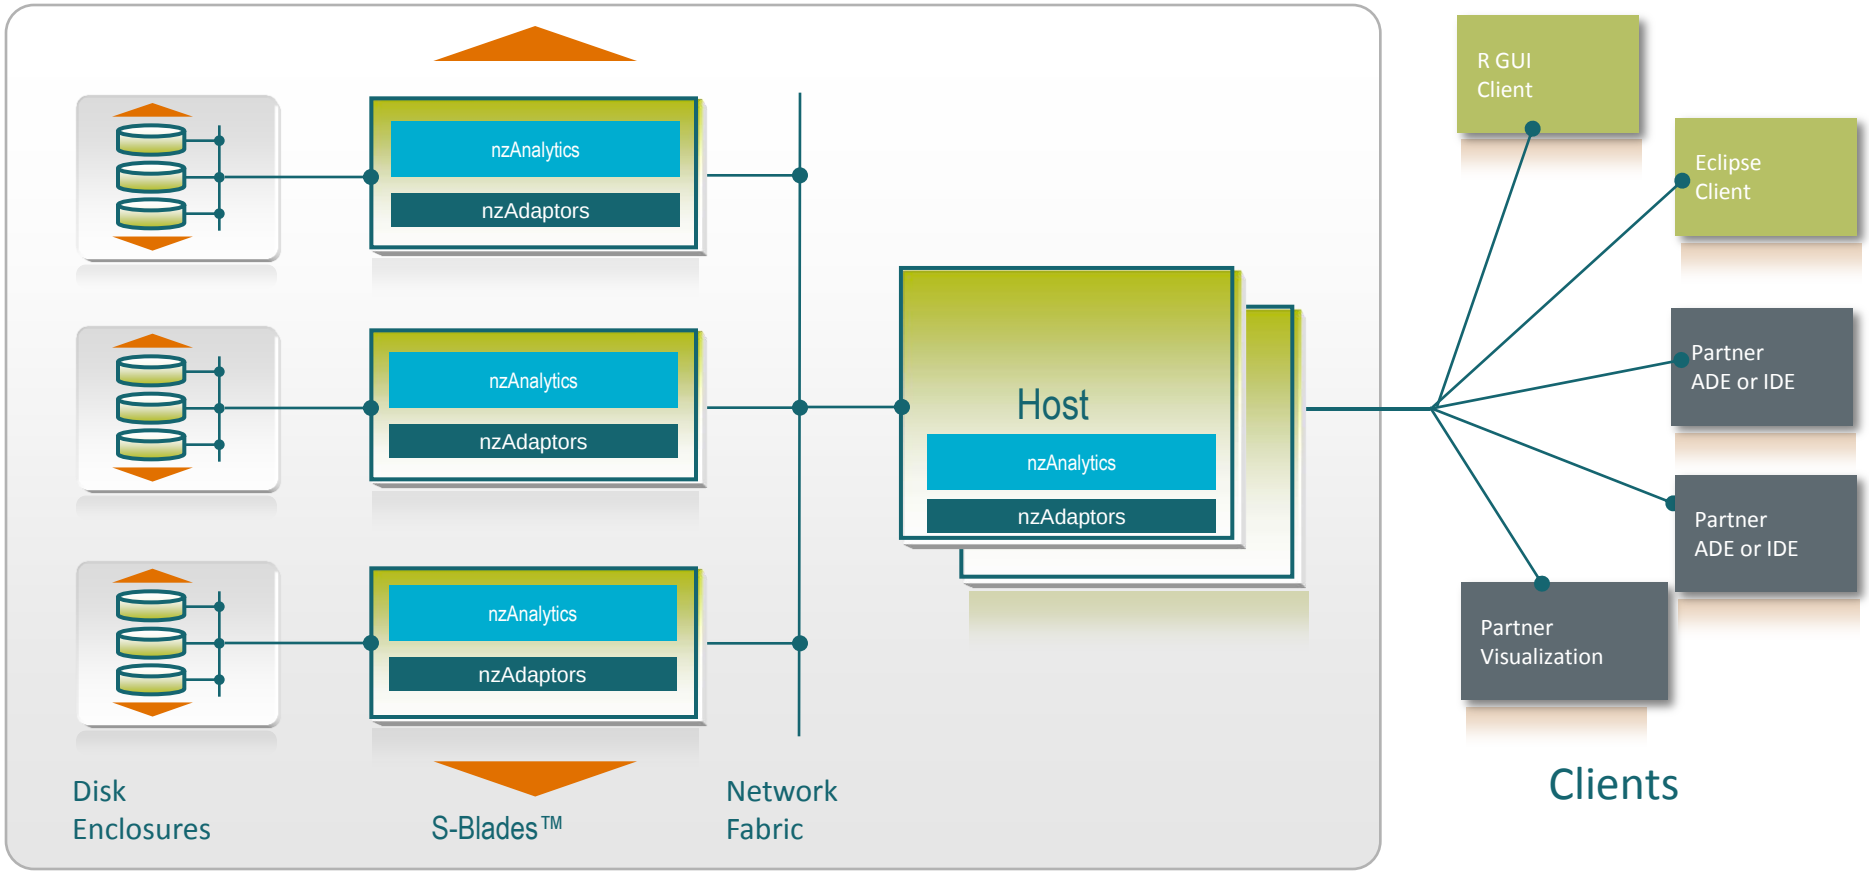

Discovering Patterns in Big Data using In-Database Analytic Model Building

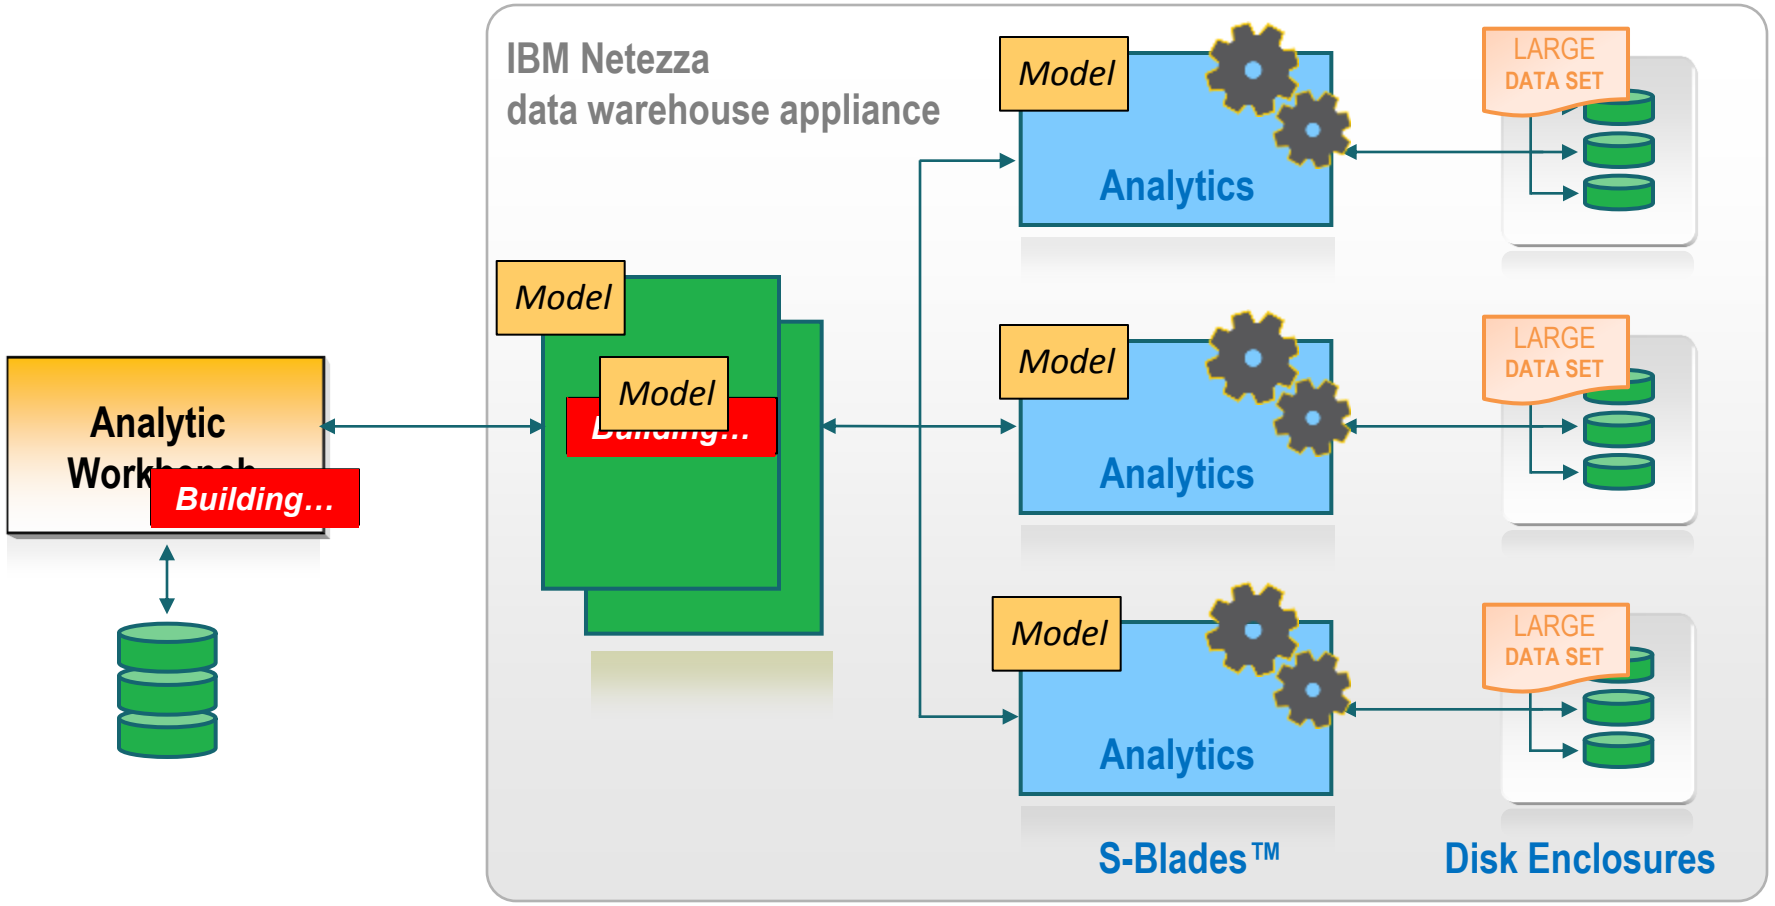

Big Data Platform - Information Integration and Governance

- Integrate any type of data to the big data platform
  - Structured
  - Unstructured
  - Streaming
- Governance and trust for big data
  - Secure sensitive data
  - Lineage and metadata of new big data sources
  - Lifecycle management to control data growth
  - Master data to establish single version of the truth

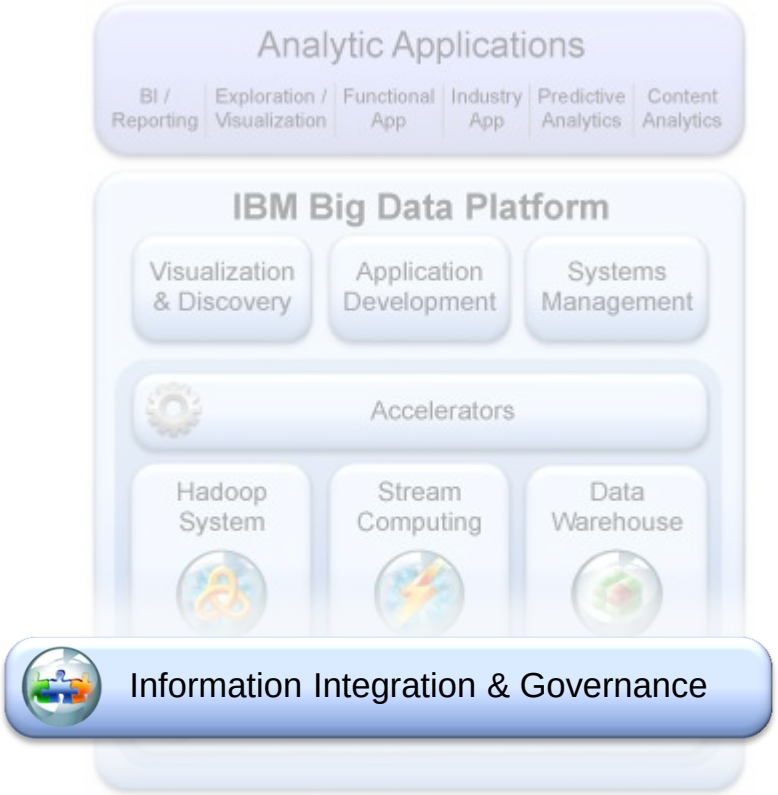

Leverage purpose-built connectors for multiple data sources

Connect any type of data through optimized connectors and information integration capabilities

- Massive volume of structured data movement
  - 2.38 TB / Hour load to data warehouse
  - High-volume load to Hadoop file system
- Ingest unstructured data into Hadoop file system
- Integrate streaming data sources

InfoSphere DataStage for structured data

DataStage

Integrate, transform and deliver data on demand across multiple sources and targets including databases and enterprise applications

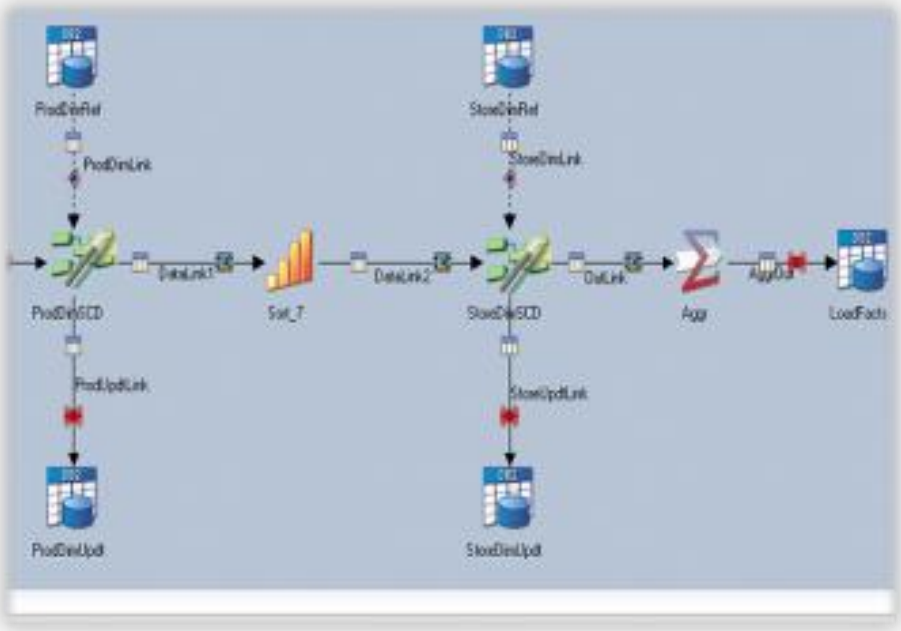

Requirements

- Integrate and transform multiple, complex, and disparate sources of information
- Demand for data is diverse – DW, MDM, Analytics, Applications, and real time

Benefits

- Transform and aggregate any volume of information
- Deliver data in batch or real time through visually designed logic
- Hundreds of built-in transformation functions
- Metadata-driven productivity, enabling collaboration

Hutchinson 3G (3) in UK Up to 50% reduction in time to create ETL jobs.

The Orchestrate engine originally developed by Torrent Systems with funding from NIST provides parallel processing

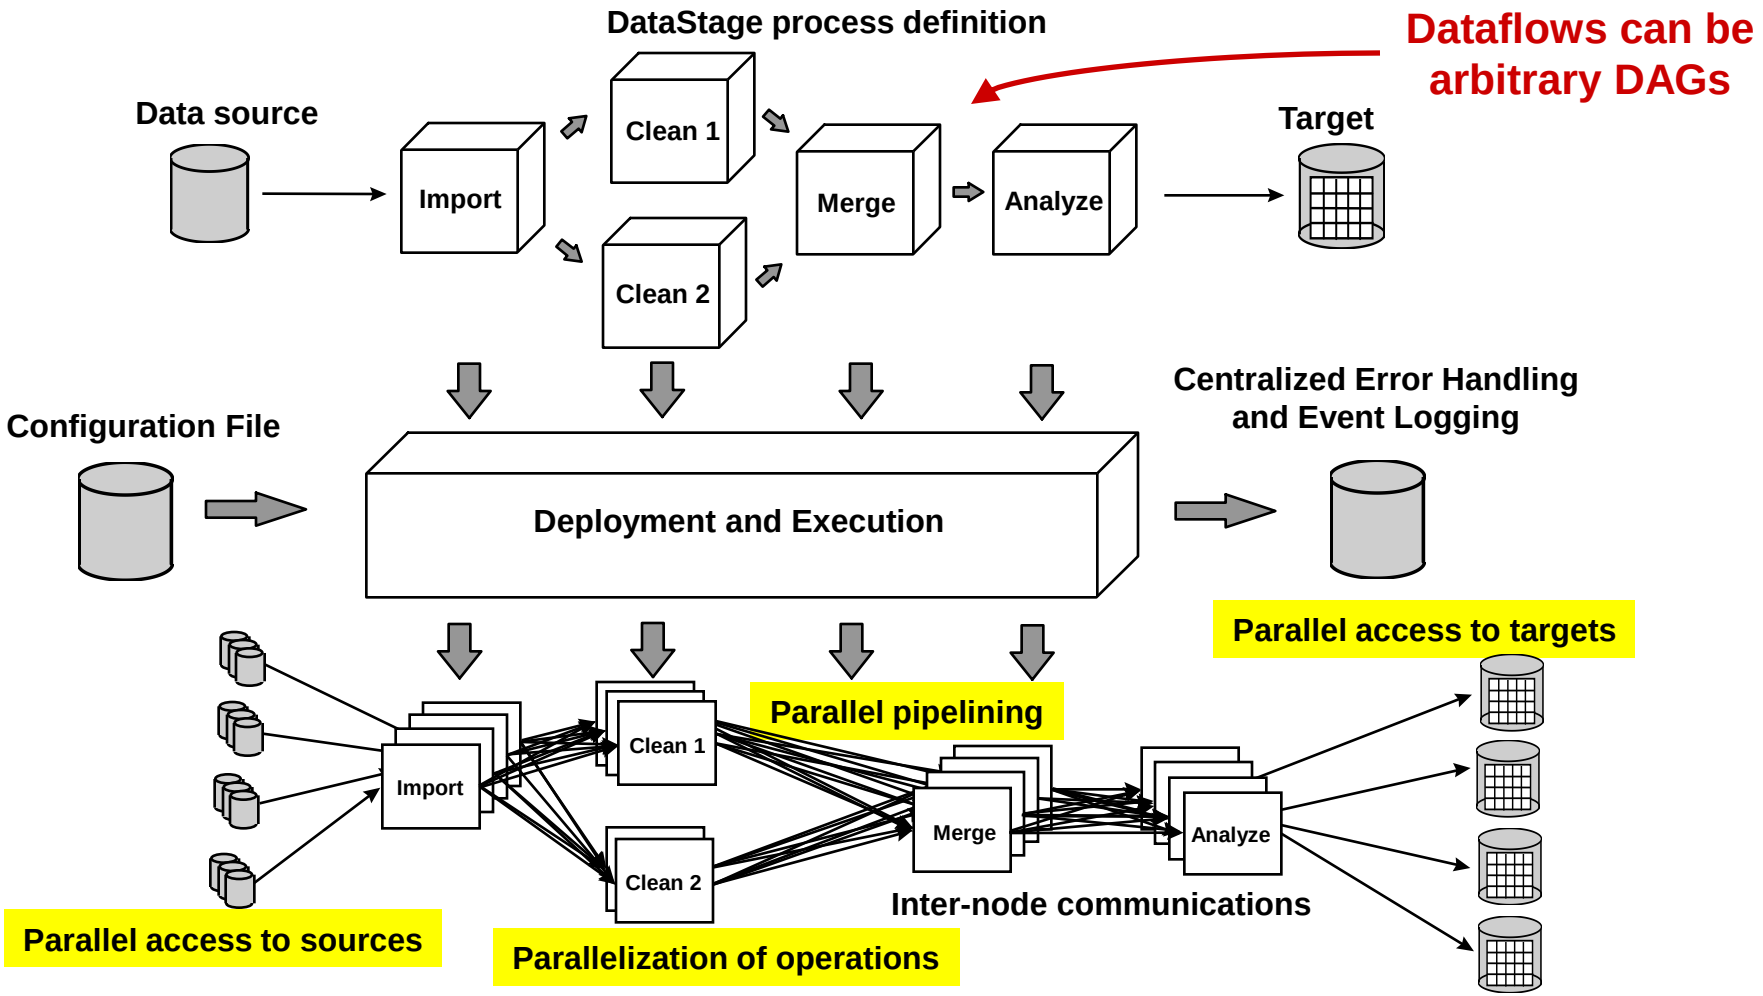

Instances of operators run in OS-level processes interconnected by shared memory/sockets

We connect to EVERYTHING

RDBMS

DB2 (on Z, I, P or X series)
Oracle
Informix (IDS and XPS)
MySQL
Netezza
Progress
RedBrick
SQL Server
Sybase (ASE & IQ)
Teradata
HP NeoView
Universe
UniData
Greenplum
PostresSQL
And more…..

General Access

Sequential File
Complex Flat File
File / Data Sets
Named Pipe
FTP
External Command Call
Parallel/wrapped 3rd party apps

Standards & Real Time

WebSphere MQ
Java Messaging Services (JMS)
Java
Distributed Transactions
XML & XSL-T
Web Services (SOAP)
Enterprise Java Beans (EJB)
EDI
EBXML
FIX
SWIFT
HIPAA

Legacy

ADABAS
VSAM
IMS
IDMS
Datacom/DB

3rd party adapters:
Allbase/SQL
C-ISAM
D-ISAM
DS Mumps
Enscribe
FOCUS
ImageSQL
Infoman
KSAM
M204
MS Analysis
Nomad
NonStopSQL
RMS
S2000
And many more….

Enterprise Applications

JDE/PeopleSoft EnterpriseOne
Oracle eBusiness Suite
PeopleSoft Enterprise
SAS
SAP R/3 & BI
SAP XI
Siebel
Salesforce.com
Hyperion Essbase
And more…

CDC / Replication

DB2 (on Z, I, P, X series)
Oracle
SQL Server
Sybase
Informix
IMS
VSAM
ADABAS
IDMS

Bold / Italics indicates Additional charge item…

InfoSphere Metadata Workbench

- See all the metadata repository content with InfoSphere Metadata Workbench
- It is a key enabler to regulatory compliance and the IBM Data Governance Maturity Model
- It provides one of the most important view to business and technical people: Data Lineage
- Understand the impact of a change with Impact Analysis
- Cross-tool reporting on:
  - Data movement
  - Data lineage
  - Business meaning
  - Impact of changes
  - Dependencies
  - Data lineage for Business Intelligence Reports

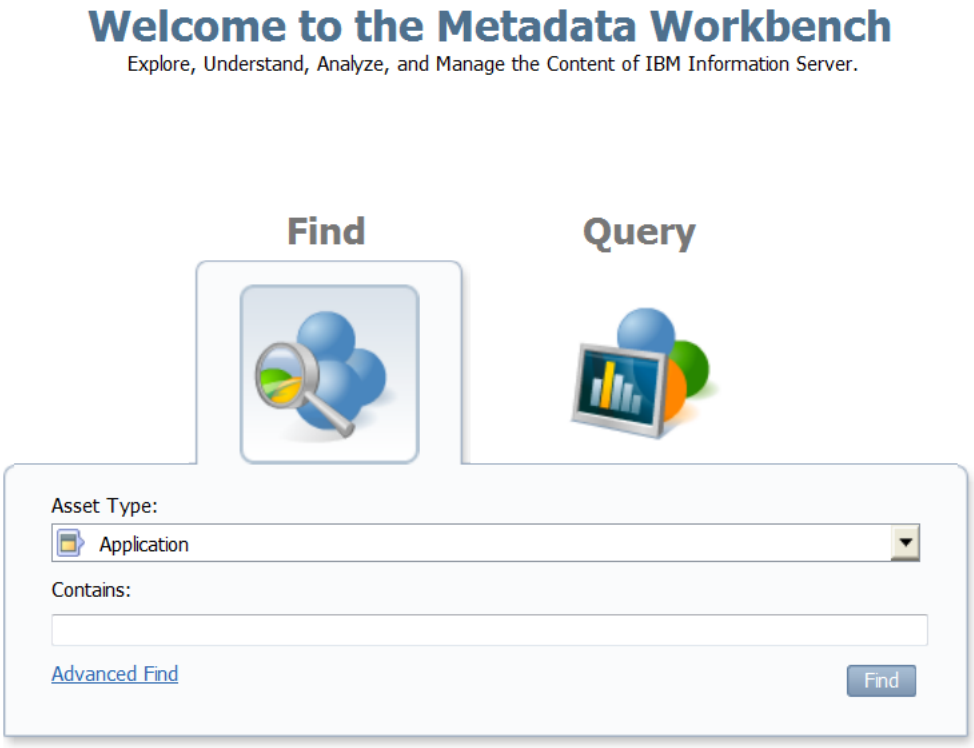

Web-based exploration of Information Assets generated and used by InfoSphere Information Server components

Data Lineage

- Traceability across business and IT domains
- Show relationships between business terms, data model entities, and technical and report fields
- Allows business term relationships to be understood
- Shows stewardship relationships on business terms
- Lineage for DataStage Jobs is always displayed initially at a summarv “Job” level

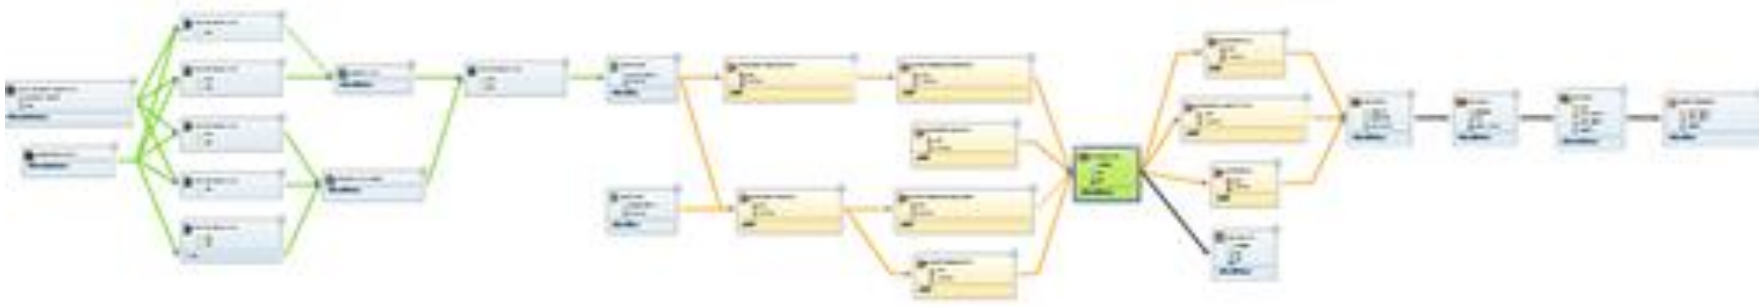

Data Lineage Extender

- Support governance requirements for business provenance
- Extended visibility to enterprise data integration flows outside of InfoSphere Information Server
- Comprehensive understanding of data lineage for trusted information
- Popular business use cases
  - Non-IBM ETL tools and applications
  - Mainframe COmmon Business-Oriented Language (COBOL) programs
  - External scripts, Java programs, or web services
  - Stored procedures
  - Custom transformations

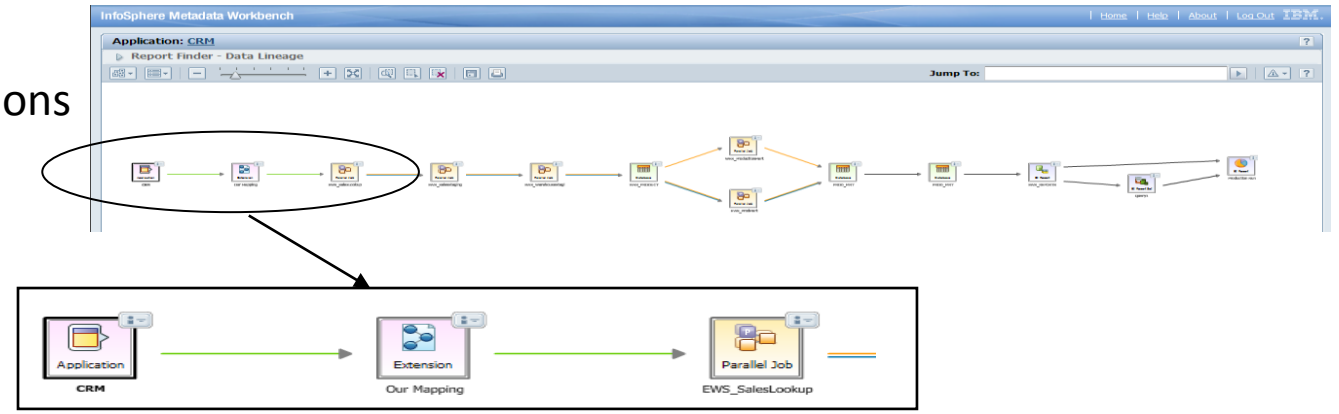

Lineage tracking with BigInsights

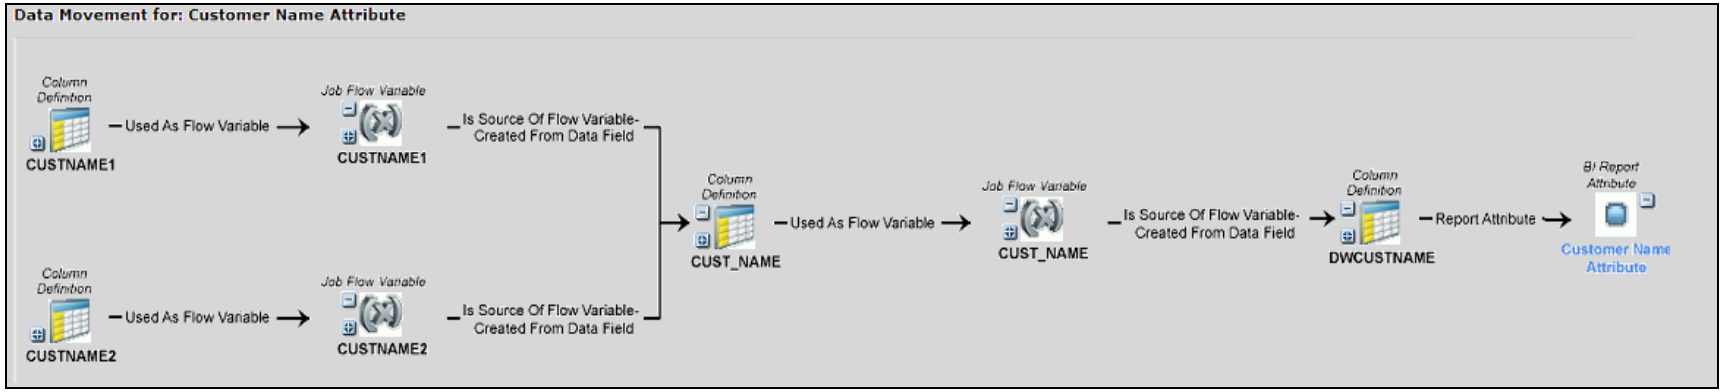

- **Extension Points** easy to define for BigInsights sources and targets
- InfoSphere Metadata Workbench can show lineage/impact of attributes and jobs from end-to-end.
- For this scenario, the current Roadmap includes
  - Better characterization of the metadata of BigInsights data sets and jobs
  - Import of the metadata into Information Server
  - Complete metadata analysis features

Big Data Platform - User Interfaces

- Business Users
  - Visualization of a large volume and wide variety of data
- Developers
  - Similarity in tooling and languages
  - Mature open source tools with enterprise capabilities
  - Integration among environments
- Administrators
  - Consoles to aid in systems management

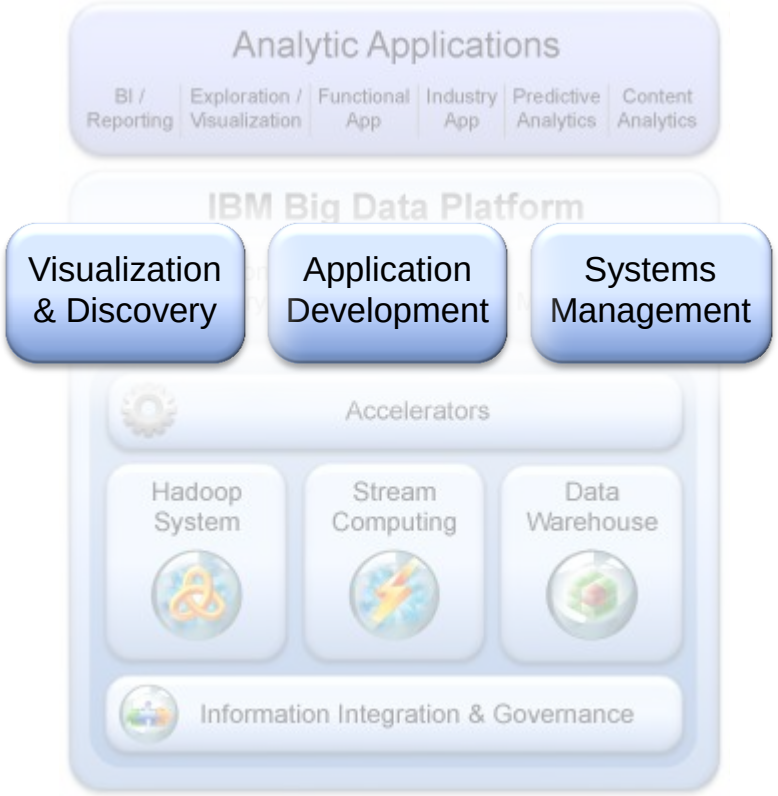

Visualization - Spreadsheet-style user interface

- Ad-hoc analytics for LOB user
- Analyze a variety of data - unstructured and structured
- Browser-based
- Spreadsheet metaphor for exploring/ visualizing data

Gather

Crawl – gather statistically
Adapter–gather dynamically

Extract

Document-level info
Cleanse, normalize

Explore

Analyze, annotate, filter
Visualize results

Iterate

Iterate through any prior step

Object-oriented APIs for implementing data-parallel algorithms

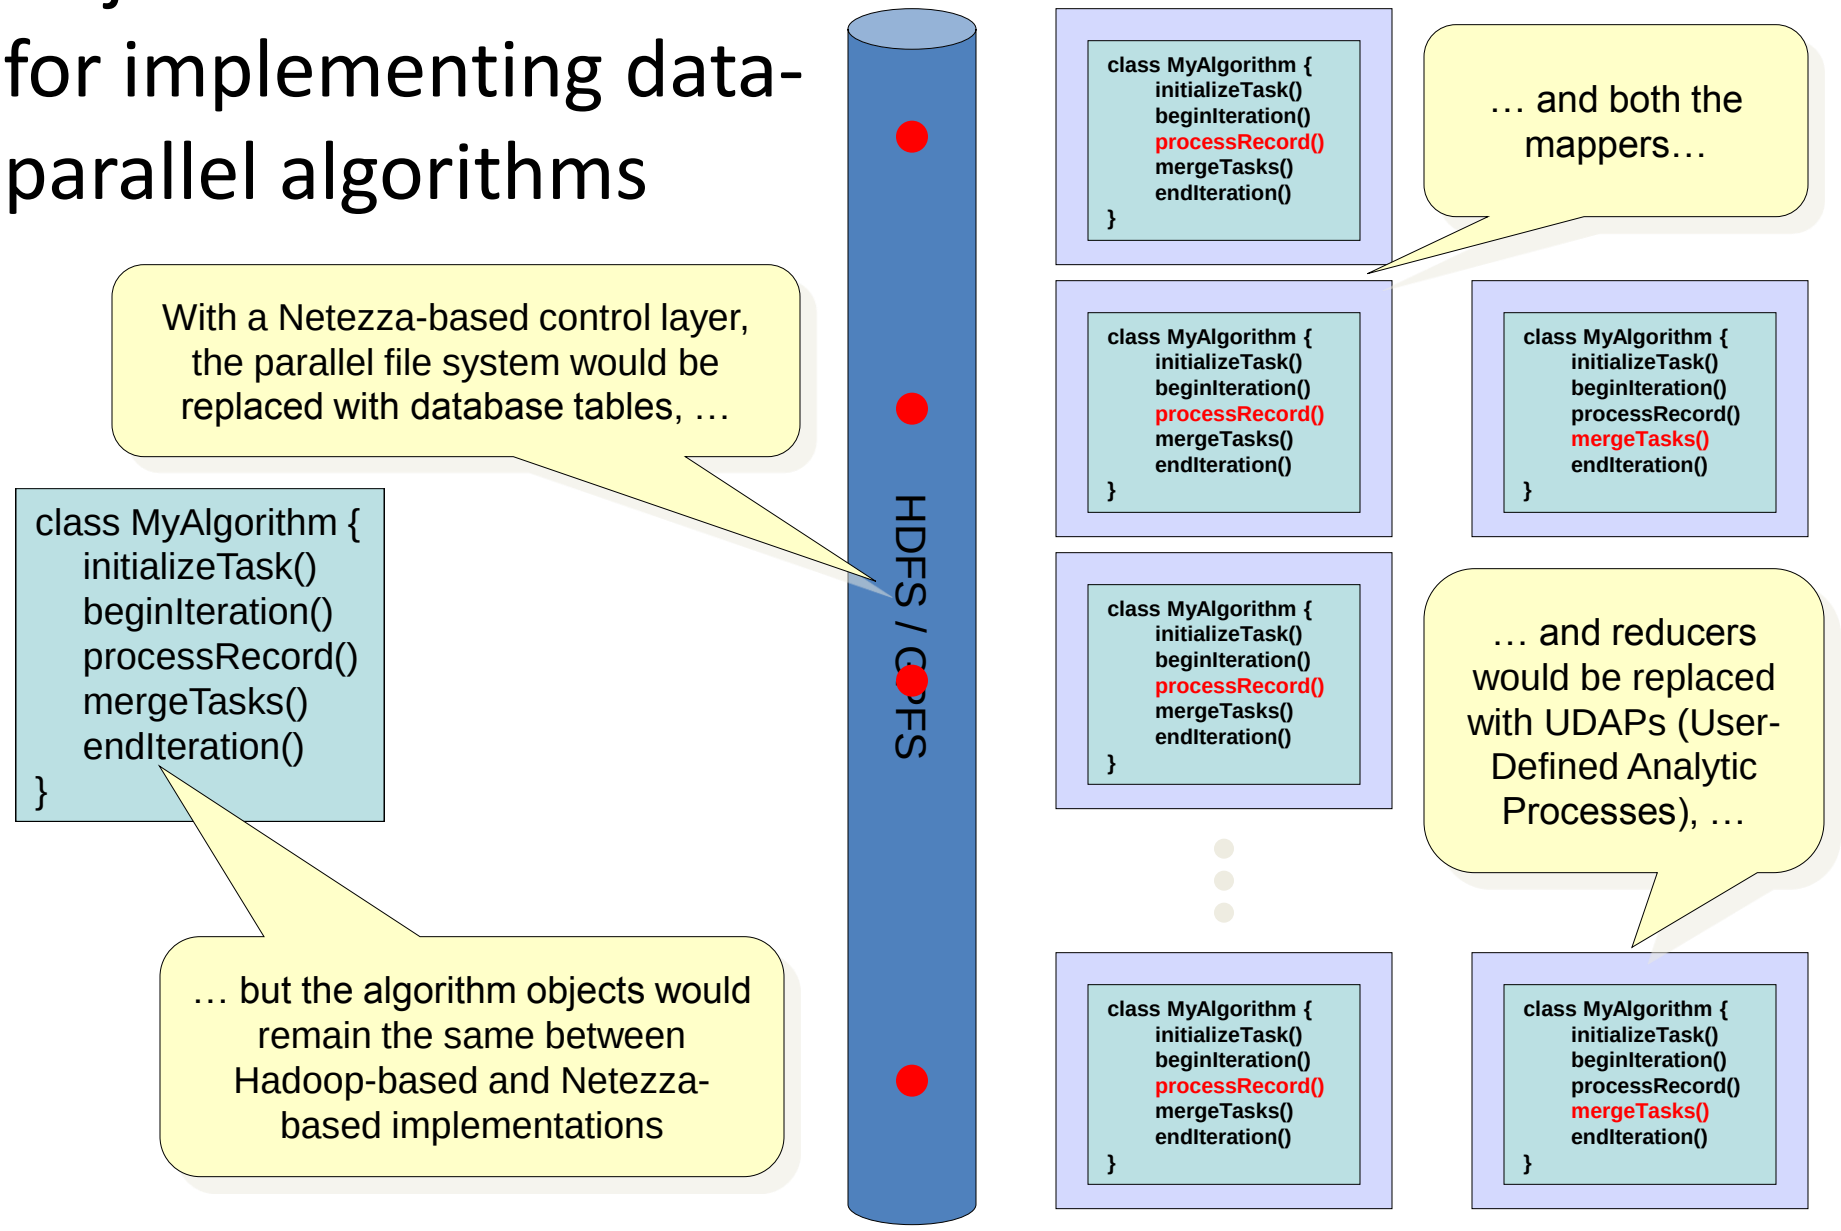

Objects can be connected into workflows with their deployment optimized using semantic properties

- D = 5*(B'*A + A*C)
  - Transpose
    - **BasicOnePassTask**
    - Can execute in Mapper or Reducer
  - MM (matrix multiply)
    - **BasicOnePassMergeTask**
    - Has Map and Reduce components
  - Add (matrix add)
    - **BasicOnePassKeyedTask**
    - Executes in Reducer and can be piggybacked
  - Multiply (scalar multiply)
    - **BasicOnePassTask**
    - Can execute in Mapper or Reducer
- Entire computation can be executed in one map-reduce job due to differentiation of BasicTasks

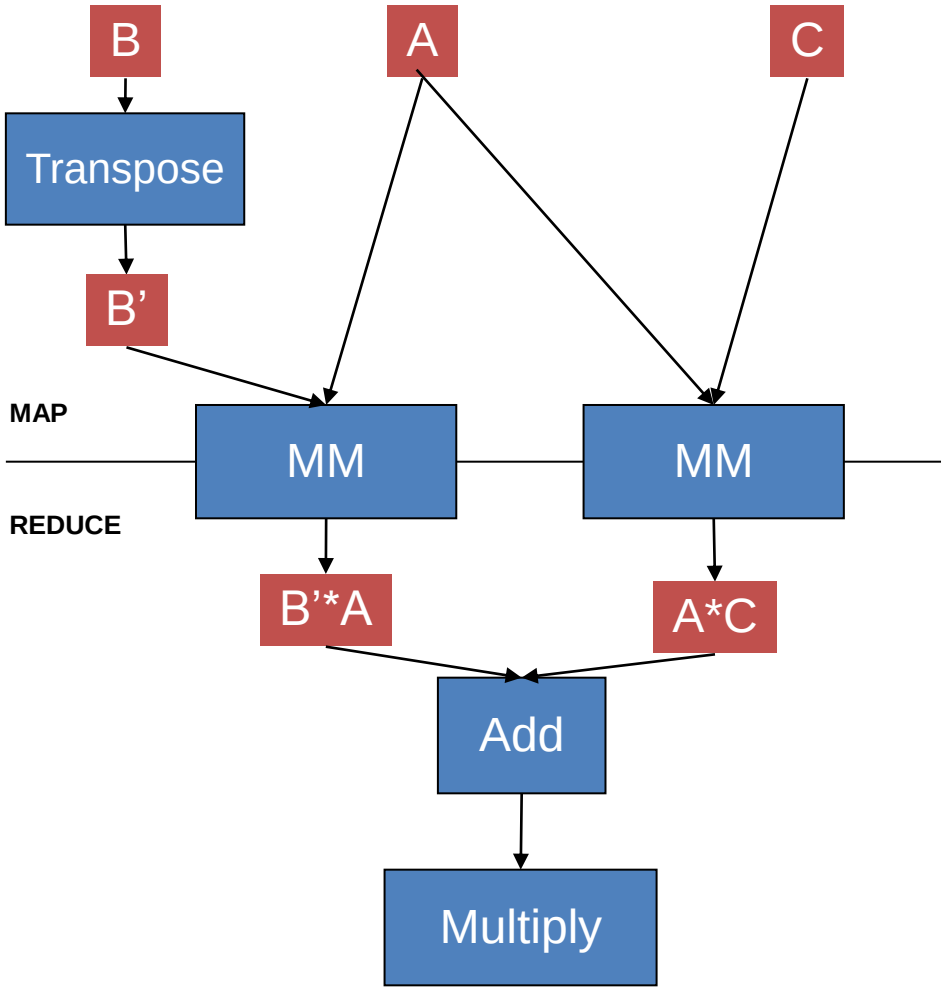

SystemML compiles an R-like language into MapReduce jobs and database jobs

Example Operations

X*Y cell-wise multiplication: $z_{ij} = x_{ij} \cdot y_{ij}$

X/Y cell-wise division: $z_{ij} = x_{ij} / y_{ij}$

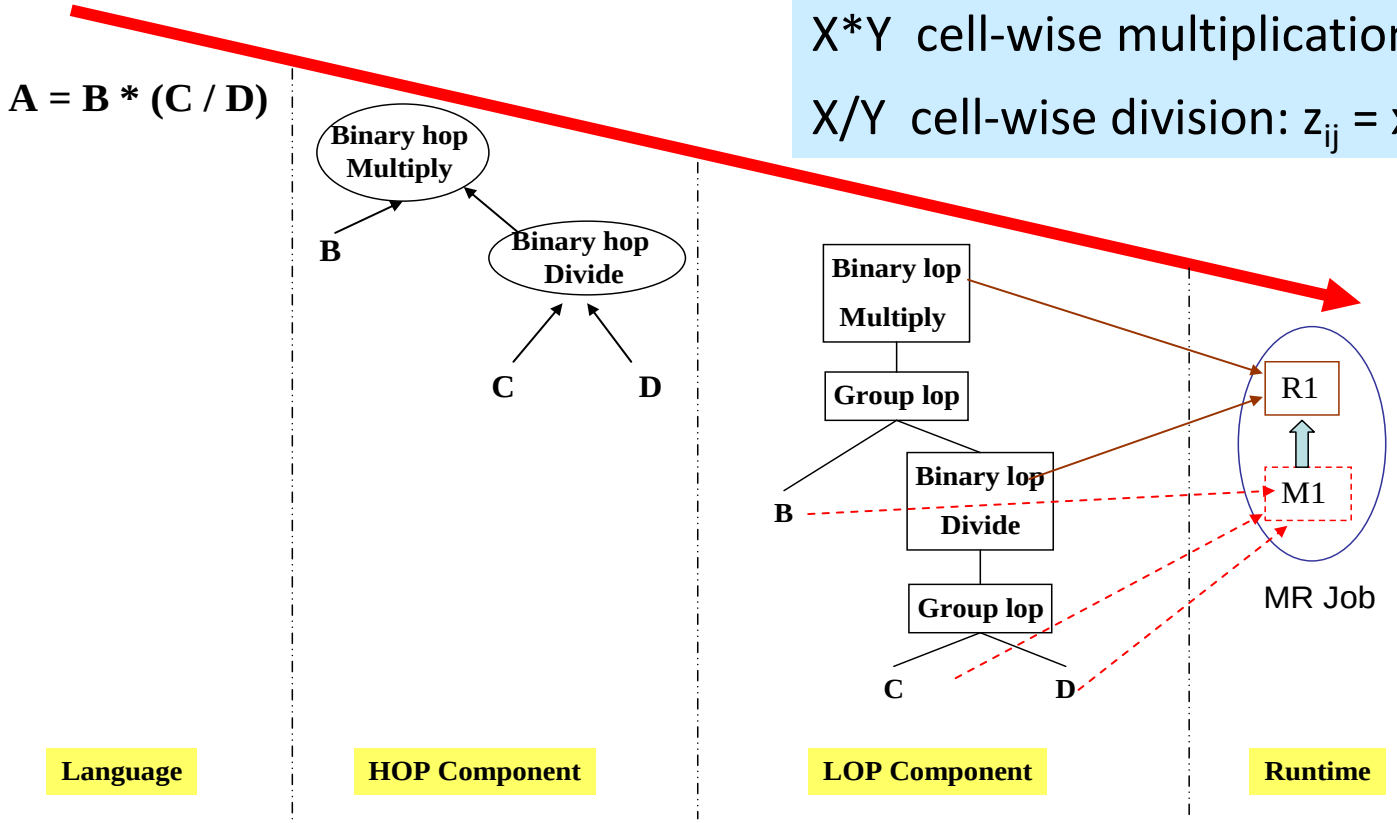

Input DML parsed into statement blocks with typed variables

Each high-level operator operates on matrices, vectors and scalars

Each low-level operator operates on key-value pairs and scalars

Multiple low-level operators combined in a MapReduce job

Approximately thirty data-parallel algorithms have been implemented to date using these and related APIs

- Simple Statistics
  - CrossTab
  - Descriptive Statistics
- Clustering
  - K-Means Clustering
  - Kernel K-Means
  - Fuzzy K-Means
  - Iclust
- Dimensionality Reduction
  - Principal Components Analysis
  - Kernel PCA
  - Non-Negative Matrix Factorization
  - Doubly-sparse NMF
- Graph Algorithms
  - Connected Graph Analysis
  - Page Rank
  - Hubs and Authorities
  - Link Diffusion
  - Social Network Analysis (Leadership)
- Regression Modeling
  - Linear Regression
  - Regularized Linear Models
  - Logistic Regression
  - Transform Regression
  - Conjugate Gradient Solver
  - Conjugate Gradient Lanczos Solver
- Support Vector Machines
  - Support Vector Machines
  - Ensemble SVM
- Trees and Rules
  - Adaptive Decision Trees
  - Random Decision Trees
  - Frequent Item Sets - Apriori
  - Frequent Item Sets - FP-Growth
  - Sequence Mining
- Miscellaneous
  - k-Nearest Neighbors
  - Outlier Detection

MARIO incorporates AI planning technology to enable ease of use

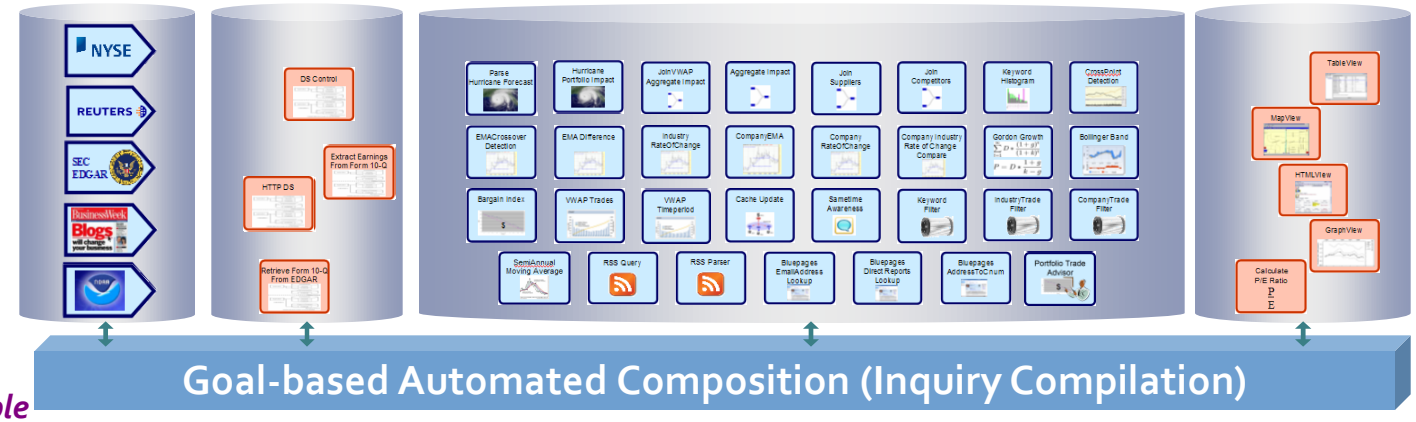

- In response to a *simple processing request*
- MARIO *automatically assembles analytics* into a variety of real-time situational applications

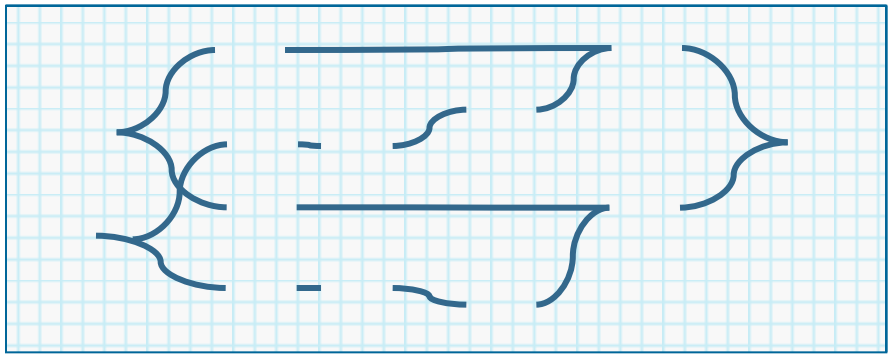

MARIO incorporates AI planning technology to enable ease of use

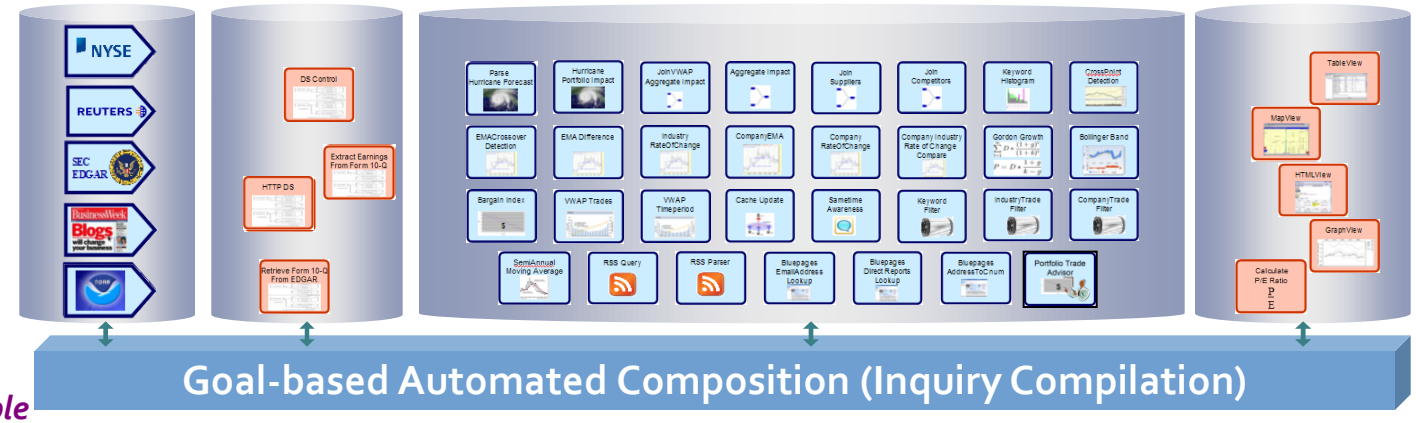

- In response to a *simple processing request*
- MARIO *automatically assembles analytics* into a variety of real-time situational applications
- Deploys application components *across multiple platforms, establishes inter-platform dataflow connections*
- Initiates *continuous processing* of flowing data
- Manages cross-platform operation

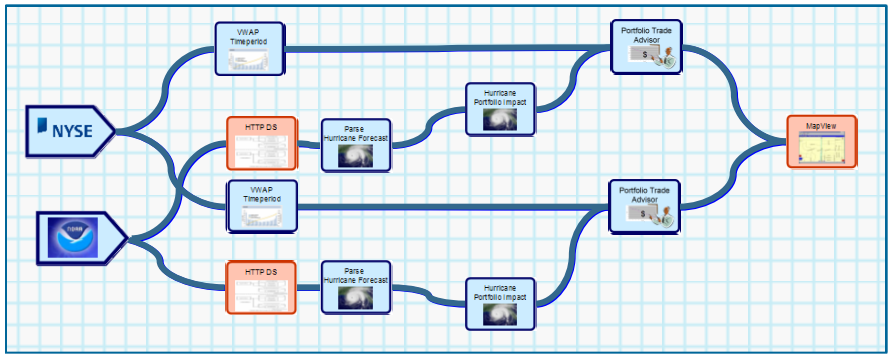

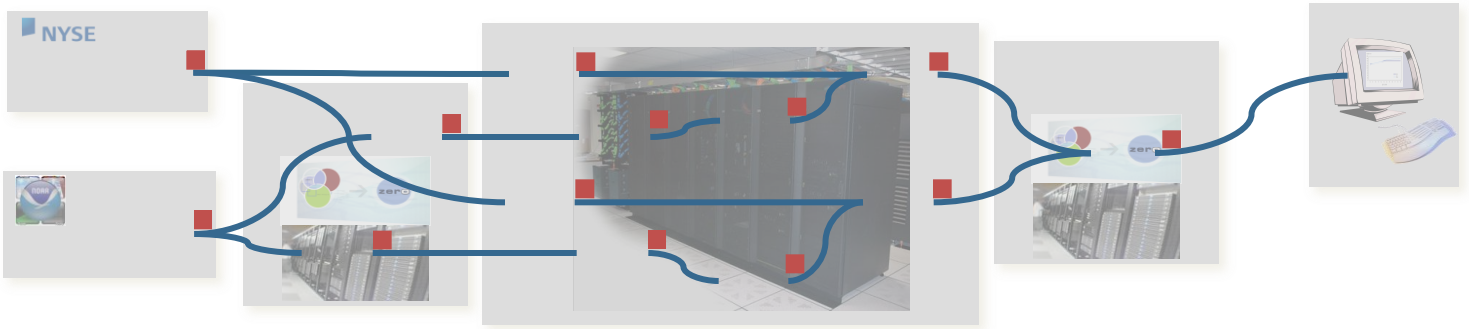

Big Data Platform - Accelerators

- Analytic accelerators
  - Analytics, operators, rule sets
- Industry and Horizontal Application Accelerators
  - Analytics
  - Models
  - Visualization / user interfaces
  - Adapters

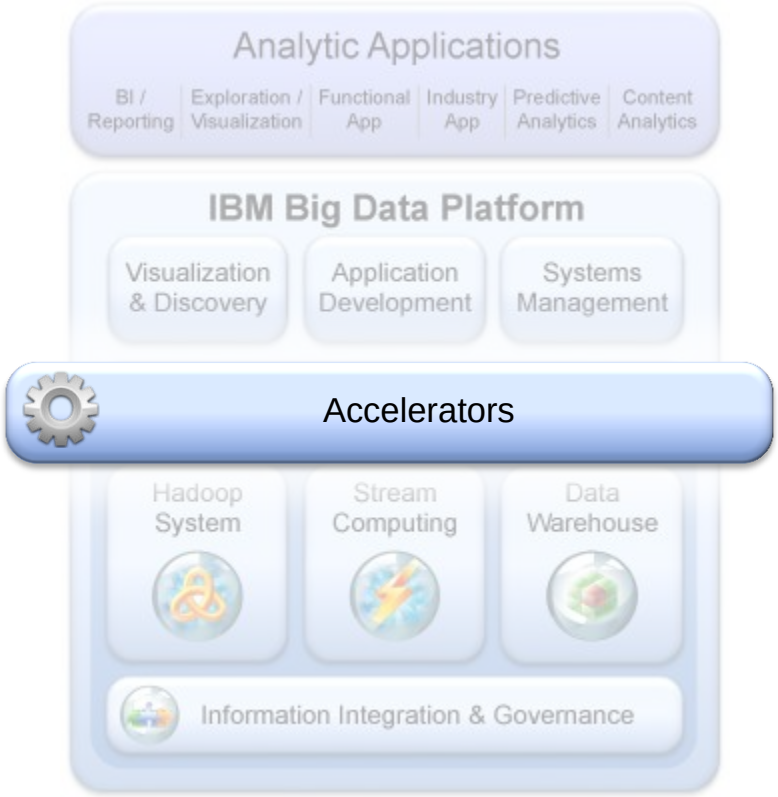

Analytic Accelerators Designed for **Variety**

Text
(listen, verb),
(radio, noun)

Simple & Advanced Text

Acoustic

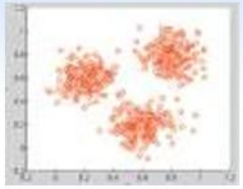

Mining in Microseconds

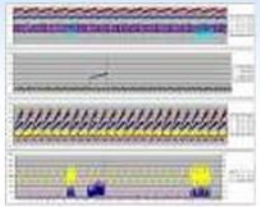

Advanced Mathematical Models

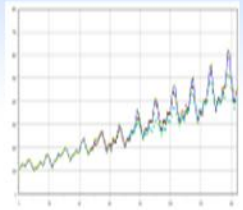

Predictive

$$\sum_{population} R(s_t, a_t)$$

Statistics

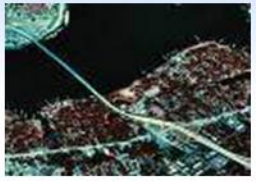

GeoSpatial

Image & Video

Accelerators Improve Time to Value

Telecommunications

CDR streaming analytics

Deep Network Analytics

Retail Customer Intelligence

Customer Behavior and Lifetime Value Analysis

Finance

Streaming options trading

Insurance and banking DW models

Social Media Analytics

Sentiment Analytics, Intent to purchase

Public transportation

Real-time monitoring and routing optimization

Data mining

Streaming statistical analysis

Over 100 sample applications

User Defined Toolkits

Standard Toolkits

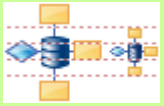

Industry Data Models

Banking, Insurance, Telco, Healthcare, Retail

Big Data Platform - Analytic Applications

Big Data Platform is designed for analytic application development and integration

BI/Reporting – Cognos BI, Attivio

Predictive Analytics – SPSS, G2, SAS

Exploration/Visualization – BigSheets, Datameer

Instrumentation Analytics – Brocade, IBM GBS

Content Analytics – IBM Content Analytics

Functional Applications – Algorithmics**,** Cognos Consumer Insights, Clickfox, i2, IBM GBS

Industry Applications – TerraEchos**,** Cisco, IBM GBS

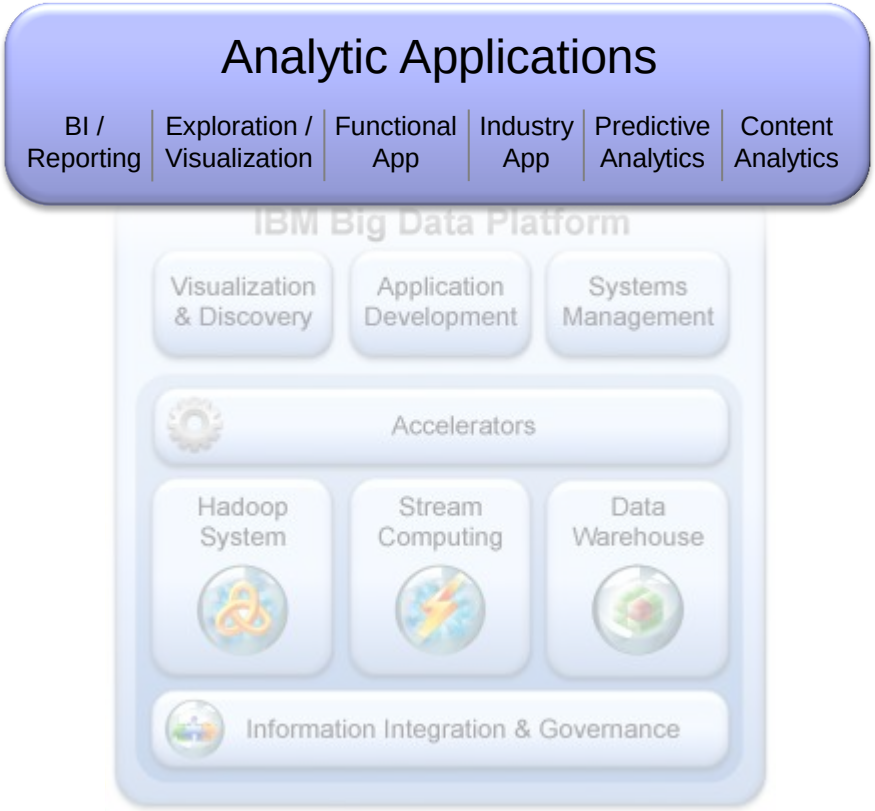


IBM