Radiant Advisors

Toronto Edition!

HADOOP AND THE BI ENVIRONMENT

TDWI Toronto Chapter

March 1, 2012

An introduction to Hadoop, its components, and position in tomorrow's BI environments

 v1.00.000a

HADOOP AND THE BI ENVIRONMENT

Abstract

Since its release into the open source community, Hadoop has been adopted rapidly to meet the challenges of analyzing big data environments, and case studies are plentiful. As Hadoop comes into the spotlight, more and more companies ask, "How does Hadoop fit into my BI strategy and architecture?" This course will introduce BI practitioners to Hadoop and its history. It examines Hadoop components and architecture, then explores strengths, weaknesses, and integration with current BI architectures.

Hadoop is no silver bullet for the many needs of business intelligence yet it is a new a powerful next step in data management and analysis. We will take a look at the new skills and training that are needed to work with Hadoop and other NoSQL databases in the open source environment. We'll close with some thoughts on what's next after Hadoop as well.

You Will Learn

- Introduction to Hadoop and its history
- Understand the components and architecture for implementing Hadoop
- Explore how best to integrate with your BI strategy

Geared Towards

- BI Directors and Managers
- Architects
- Developers
- Enthusiasts

 v1.00.000a

Radiant Advisors

Principal and Founder, Radiant Advisors

JOHN O'BRIEN

With over 20 years of experience delivering value through data warehousing and BI programs, John O'Brien's unique perspective comes from the combination of his roles as a practitioner, consultant, and vendor in the BI industry. His knowledge in designing, building, and growing enterprise BI systems and teams brings real world insights to each role and phase within a BI program.

Today, through Radiant Advisors John provides research and advisory services that guide companies in meeting the demands of next generation information management, architecture, and emerging technologies.

Instructor 10+ years

As a recognized thought leader in BI, John has been publishing articles and presenting at conferences in North America and Europe for the past 10 years, including The Data Warehousing Institute where he has been invited as one of TDWI's Best Practices judges, Executive Summit presenters and expert panel participants. John has also developed and presented many of his own courses that now comprise the initial Radiant Advisors Learning Catalog.

Experienced

In 2005, John co-founded and became CTO of a data warehouse appliance company that raised $43 million in several rounds of venture capital financing and has many global production customers. As CTO, John's primary role was to focus product development and BI market strategy.

Education

John has aB.S. in Mechanical Engineering from California State University with an emphasis in control systems and instrumentation and an Executive M.B.A. from University of Colorado. He is a Certified Business Intelligence Professional (CBIP) since 2005 with mastery levels in Leadership and Administration, Database Administration and Business Intelligence.

Radiant Advisors

Introduction to Hadoop Components

HADOOP OVERVIEW

History of Hadoop

BIG DATA CHALLENGE

- Big Data requires large amounts of servers and storage
- Scaling costs
  - Must leverage commodity hardware
  - More Cores & Spindles (Cheaper & Faster)
  - Density – computers per rack (Power & Cooling)

Design software to handle hardware failures

Radiant Advisors

History of Hadoop

BIG DATA CHALLENGE

- Bandwidth bottleneck to move data
- Network can't move data to processors
- Massive amount of Nodes with processors and storage to parallelize
- New Paradigm needed

move computations, not data

Radiant Advisors

History of Hadoop

WHAT IS HADOOP?

- The Apache Hadoop software library is a framework that allows for the distributed processing of large data sets across clusters of computers using a simple programming model.
- It is designed to scale up from single servers to thousands of machines, each offering local computation and storage.
- Rather than rely on hardware to deliver high-availability, the library itself is designed to detect and handle failures at the application layer, so delivering a highly-availability service on top of a cluster of computers, each of which may be prone to failures.

Project subprojects:

- Hadoop Common (Core)
- Hadoop Distributed File System (HDFS)
- Hadoop MapReduce

Related Apache projects:

- Avro
- Cassandra
- Chukwa
- Hbase
- Hive
- Mahout
- Pig
- ZooKeeper

Radiant Advisors

History of Hadoop

DEFINITIONS

WIKIPEDIA
The Free Encyclopedia

- Apache Hadoop is a software framework that supports data-intensive distributed applications under a free license.
- It enables applications to work with thousands of nodes and petabytes of data.
- Hadoop was inspired by Google's MapReduce and Google File System (GFS) papers.
- Hadoop is a top-level Apache project being built and used by a global community of contributors, written in the Java programming language.
- Yahoo! has been the largest contributor to the project, and uses Hadoop extensively across its businesses.
- It was originally developed to support distribution for the Nutch search engine project.

Hadoop was created by Doug cutting,
who named it after his son's toy elephant...

Radiant Advisors

History of Hadoop

DEFINITIONS

"Hadoop is an open source framework for writing and running distributed applications that process large amounts of data." - *Chuck Lam, Hadoop in Action*

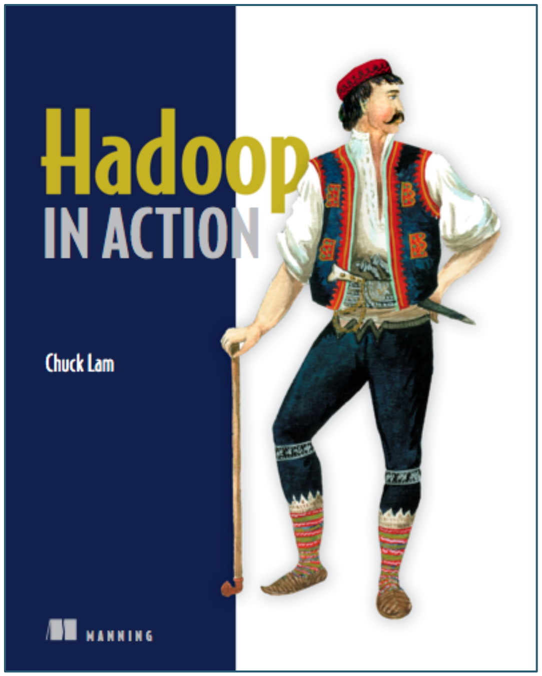

Key distinctions are:

- Accessible – Hadoop runs on large clusters of commodity machines
- Robust – Hadoop is architected with the assumption of frequent hardware malfunctions
- Scalable – Hadoop scales linearly to handle larger data by adding more nodes to the cluster
- Simple – Hadoop allows users to quickly write efficient parallel code

Hadoop cluster is a set of commodity machines networked together in one location.

Radiant Advisors

Extending Hadoop

HADOOP ECOSYSTEM

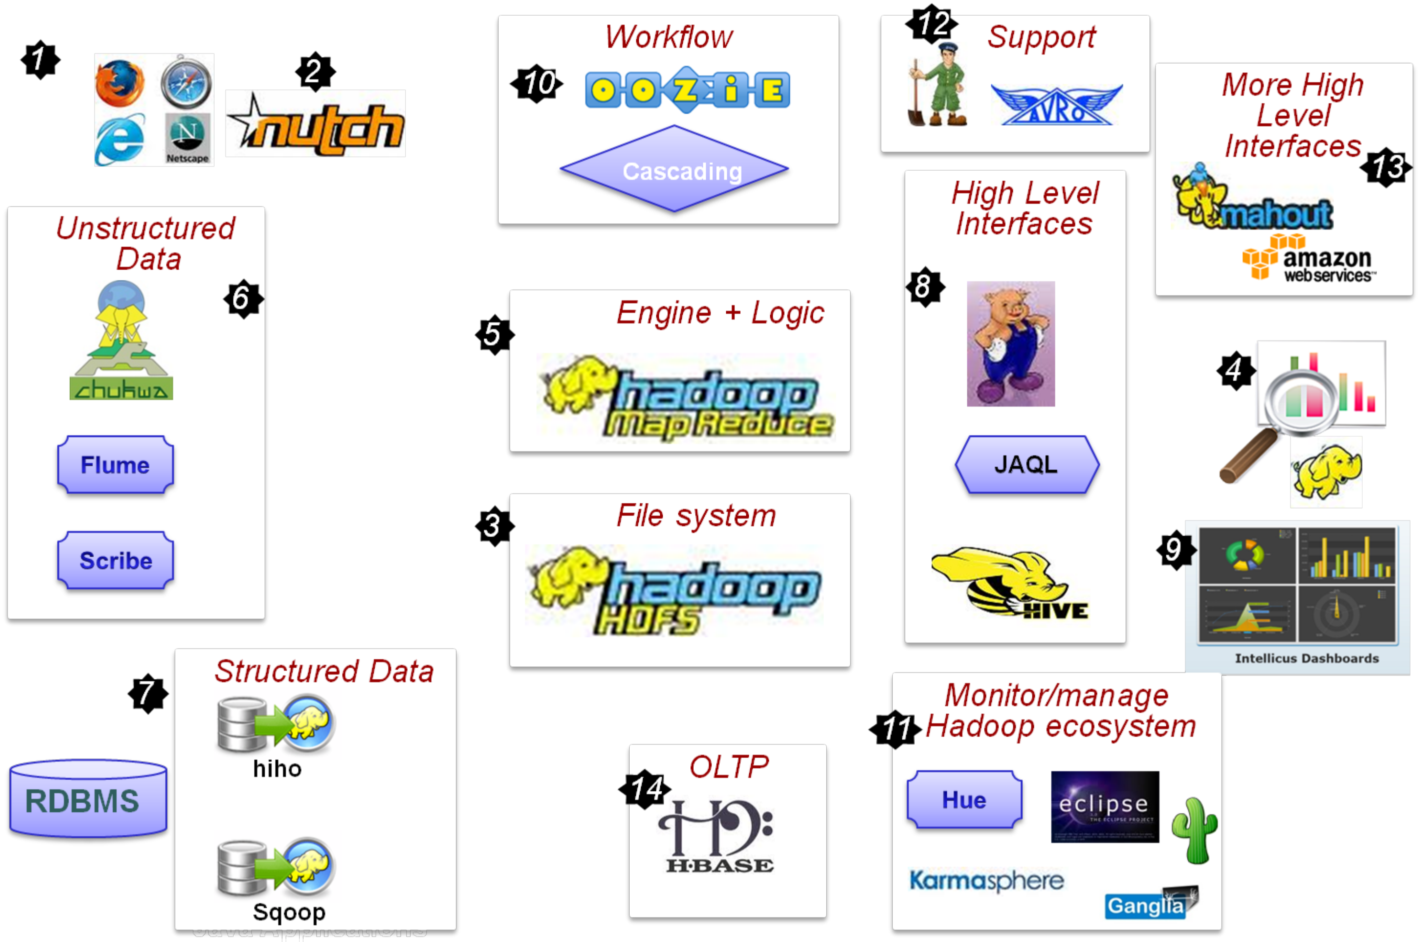

History of Hadoop

DOUG CUTTING

Apache Hadoop Co-founder

1980s

- Full Text Indexes, B-Tree
- Random access
- Seek time is wasted time
- Too slow when updates are frequent

1990s

- Excite
- Big batch processes

2000s

- ***Lucene***
- Open source
- Full-text search library
- Introduction to Apache

2002

- ***Nutch***
- Open Source
- Web search engine
- Sort/Merge Optimizations
- Distributed solution

2004

- Google publishes
- GFS & MapReduce
- Provides automation
- In a few years, 20-40 nodes

2006

- Yahoo! hires Cutting
- ***Nutch*** becomes ***Hadoop***
- Automates maintenance
- 2009 joined Cloudera

2009-10, became Chairman of the Board of the Apache Software Foundation

Radiant Advisors

History of Hadoop

MOVING CODE NOT DATA

Accessing data

with MapReduce Programs

scaling the application

Data stored in

Hadoop Distributed File System

History of Hadoop

HADOOP AT YAHOO!

2004

- Initial versions of Hadoop components prototyped as part of Apache Lucene project

2006

- Apache Hadoop launched by Yahoo! employees
- Yahoo! scales Hadoop to run on 600 server clusters, up from 20 server clusters
- Hadoop deployed to support science at Yahoo!
  - Dramatically speeds innovation, from months to days

2007

- Science results improve Yahoo! Search and Advertising products
- Yahoo scales Hadoop to 1000 server clusters

2008

- Hadoop use evolves beyond science and reaches production across the network
  - Search
  - Advertising optimization
  - Machine Learning
  - Content feeds and optimization

2009

- Over a dozen clusters of up to 4000 servers each
- Released Yahoo! Distribution of Hadoop
- Hadoop adoption:
  - Inventory management
  - Ad delivery optimization
  - Behavioral Targeting
  - Data pipeline

2010

- Yahoo! continues to run the world's largest production Hadoop applications:
  - Yahoo.com Today Module
  - Mail Anti-Spam
  - Advertising Data Systems
  - Content consolidation
  - Data analytics

Radiant Advisors

Architecture Components

MASTER / SLAVE ARCHITECTURE

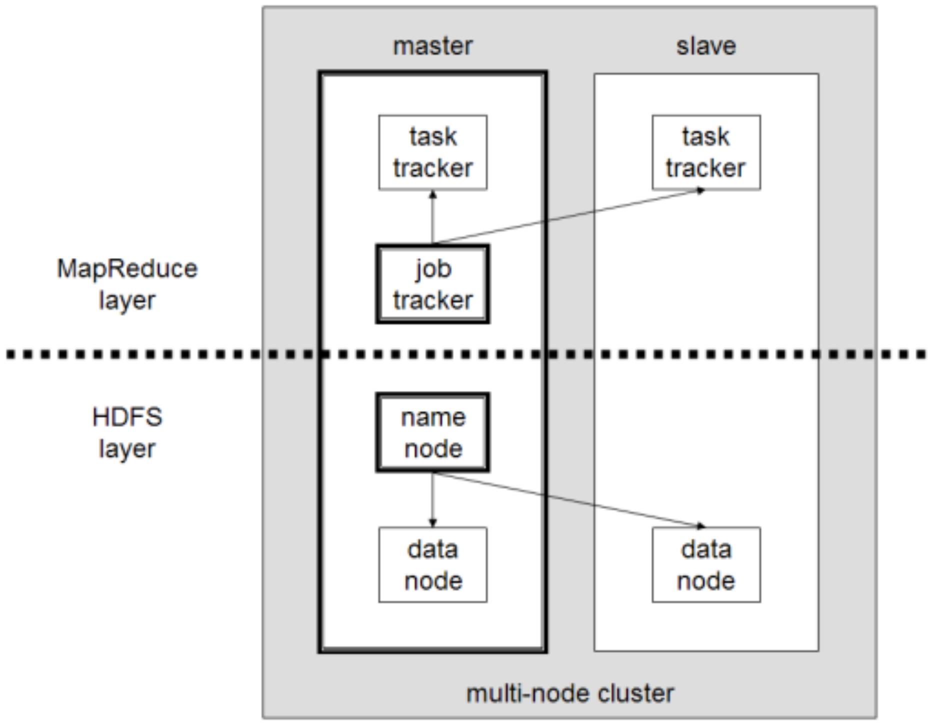

- NameNode
- DataNode
- Secondary NameNode
- JobTracker
- TaskTracker

Master / Slave Architecture

Radiant Advisors

Architecture Components

NAMENODE

Most vital Hadoop daemon:

- Master of HDFS
- Directs slave DataNodes to perform low-level I/O taskes
- Memory and I/O intensive
- Does not perform as DataNode
- Does not perform as TaskTracker
- NameNode tells each client which DataNode each block resides in

Bookkeeper of HDFS:

- Stores file metadata
- How files are broken down into file blocks
- Which nodes store which file blocks
- Overall health of the distributed filesystem

NameNode is the single point of failure of the Hadoop cluster

Radiant Advisors

Architecture Components

DATANODE

DataNode work:

- Each slave machine in cluster will host a DataNode daemon
- Reading and writing HDFS blocks to actual files on local filesystem
- Communicates with other DataNodes to replicate data blocks

Communicate with NameNode:

- NameNode tells client which DataNode has blocks needed
- At initialization, each DataNode informs NameNode of its blocks
- DataNodes continually poll NameNode for instructions to create, move or delete local blocks

DataNodes inform NameNode of its local blocks "mapping"

Radiant Advisors

Architecture Components

SECONDARY NAMENODE (SNN)

Assistant daemon to NameNode:

- Each cluster has one SNN
- Monitors state of cluster HDFS
- No DataNode daemons on SNN
- No TaskTracker daemons on SNN

Communicate with NameNode:

- Takes snapshots of HDFS metadata
- Snapshot intervals defined by cluster configuration
- Helps minimize downtime or loss of data

Secondary NameNode is not a backup or failover for the NameNode

Radiant Advisors

Architecture Components

JOBTRACKER

Master/Slave compute Architecture:

- JobTracker is the master overseeing overall execution of MapReduce job
- Daemon between application and Hadoop
- Automatically relaunches the task if it fails
- Only one JobTracker per Hadoop cluster

JobTracker process:

1. Submit your code to cluster
2. Determines execution plan
3. Which files to process
4. Assigns nodes to different tasks
5. Monitors all tasks as they run

JobTracker is typically run on a server as a master node of the cluster

Radiant Advisors

Architecture Components

TASKTRACKER

Slave compute daemon:

- Executes individual tasks assigned by the JobTracker
- Single TaskTracker per slave node
- Each TaskTracker can spawn multiple JVMs to handle map or reduce tasks in parallel

Communicates with JobTracker:

- JobTracker receives heartbeats
- JobTracker can resubmit task to other nodes in cluster

Each slave nodes has a single TaskTracker and DataNode daemon

Radiant Advisors

Architecture Components

ARCHITECTURE TOPOLOGY

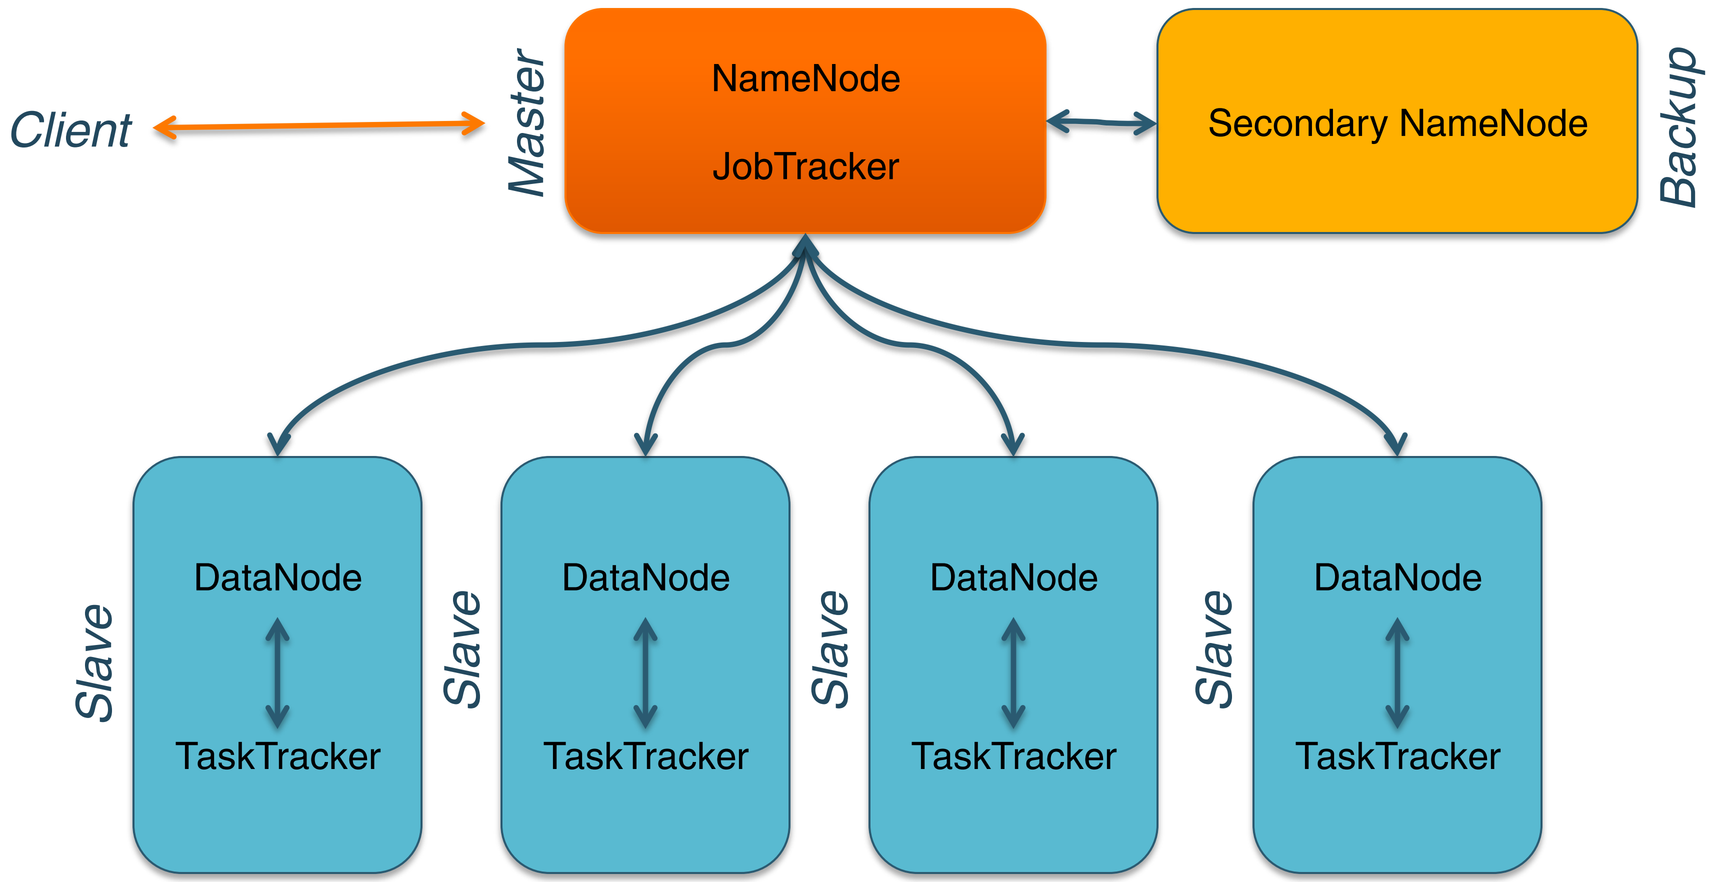

Radiant Advisors

Architecture Components

THREE OPERATING MODES

Standalone Mode

- Default mode for Hadoop
- Assumes single machine
- Chooses conservative
- Assumes minimal config
- No Node communication
- No HDFS
- No Hadoop daemons
- MapReduce development and debugging

Pseudo-distributed

- “Cluster of One”
- All Hadoop daemons running
- Examines memory usage
- HDFS input/out issues
- Daemon interactions
- Further development and debugging mode

Fully-distributed

- Production mode
- Master node of cluster hosts NameNode and JobTracker
- Backup node hosts Secondary NameNode
- Slave boxes running both DataNode and TaskTracker daemons

Radiant Advisors

MapReduce Programming

UNDERSTANDING MAPREDUCE

- MapReduce is a framework for processing huge datasets on certain kinds of distributable problems using a large number of computers (nodes)
- Collectively referred to as a cluster (if all nodes use the same hardware) or as a grid (if the nodes use different hardware)
- Computational processing can occur on data stored either in a filesystem (unstructured) or within a database (structured)

It takes some time to get used to...

Radiant Advisors

MapReduce Programming

HADOOP KEY/VALUE PAIRS

- Fundamental data unit of Hadoop
- Instead of relational tables
- Stores full description in the data with its value
- Flexible to work with less structured data types

MapReduce Programming

MAPREDUCE DATA FLOW

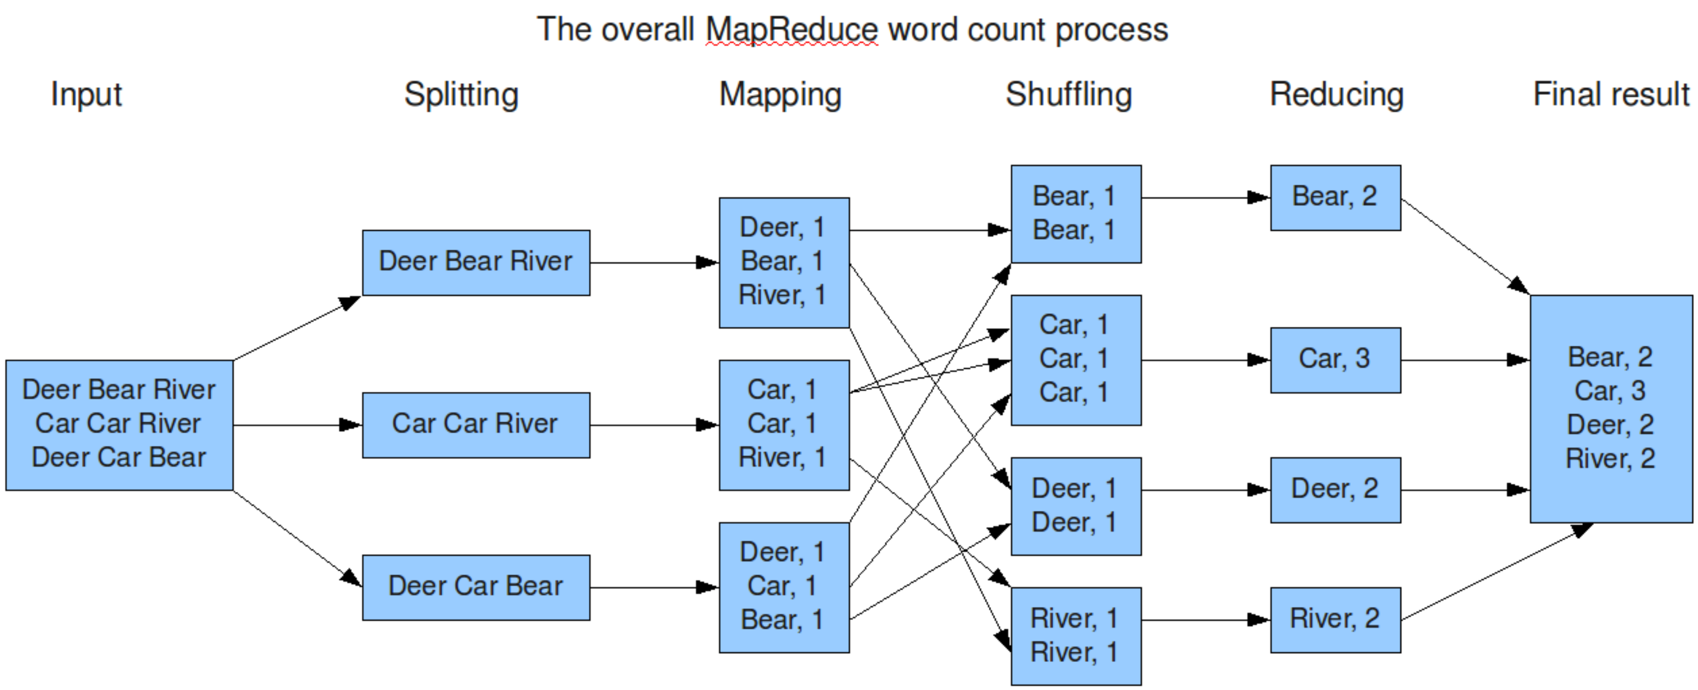

Radiant Advisors

MapReduce Programming

MAPREDUCE DATA FLOW

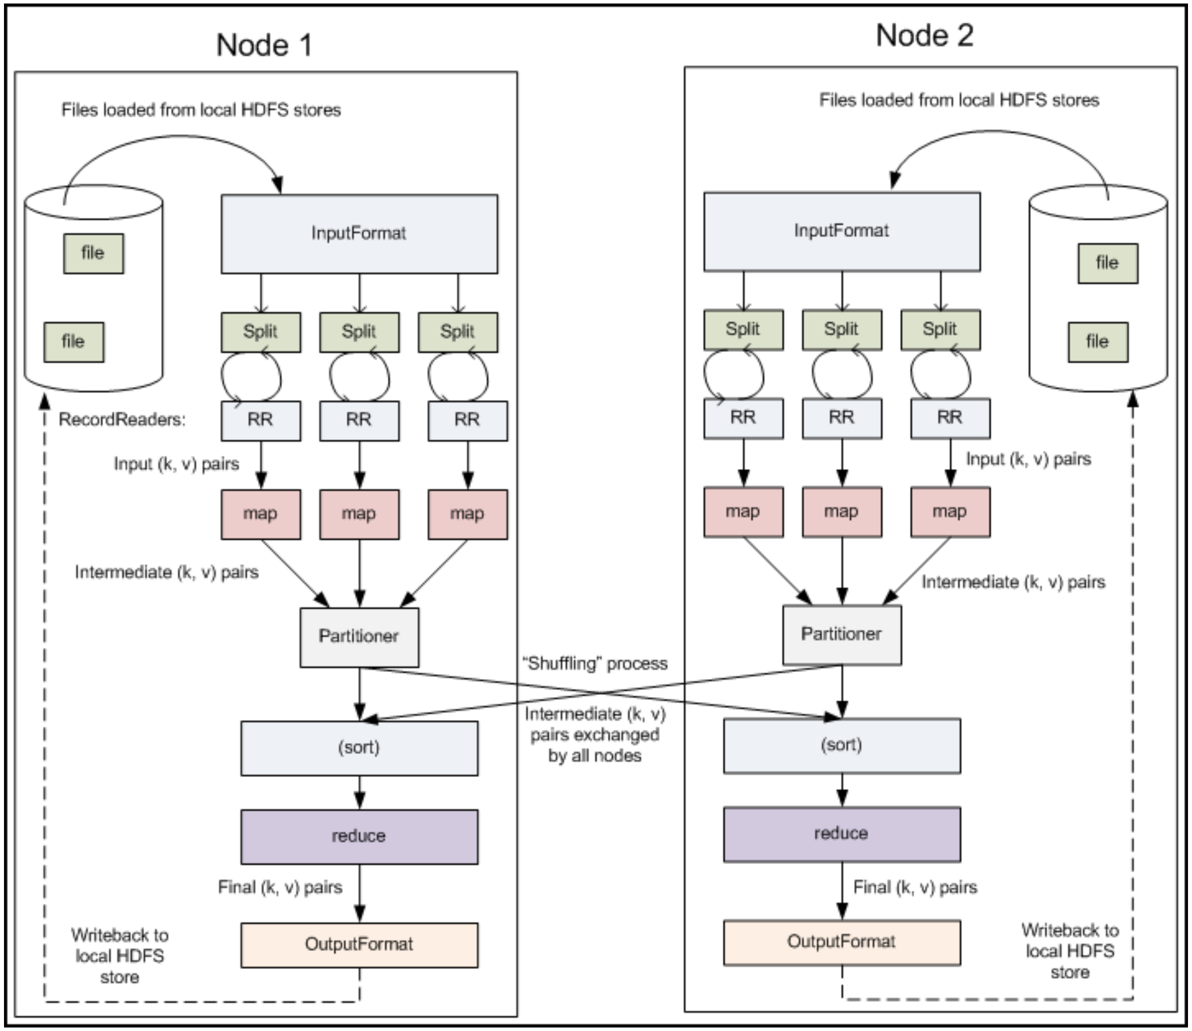

Radiant Advisors

MapReduce Programming

MAP STEP

- The master node takes the input, partitions it up into smaller sub-problems, and distributes those to worker nodes.
- A worker node may do this again in turn, leading to a multi-level tree structure.
- The worker node processes that smaller problem, and passes the answer back to its master node.

MapReduce Programming

REDUCE STEP

- The master node then takes the answers to all the sub-problems and combines them in some way to get the output – the answer to the problem it was originally trying to solve

Radiant Advisors

Extending Hadoop

PIG

Simplifies Hadoop Programming with a high-level data processing language

Two major components:

1. Pig Latin – high level data processing language
2. A Compiler – compiles to an evaluation mechanism, such as Hadoop

- Knowing how to write UDFs is a big part of learning to use Pig
- Pig can operate on data that is relational, nested, semistructured or unstructured
- The Pig interactive shell is known as Grunt

Twitter uses Pig significantly

Yahoo has >40% Pig programs

Extending Hadoop

HIVE

"A SQL like data warehouse infrastructure with HiveQL"

"Hive is a data warehousing package built on top of Hadoop"

Overview:

- Separates the user from the complexity of MapReduce
- Target user is the data analyst more comfortable with SQL for ad-hoc queries, summarization and data analysis
- Hive is a metastore for storing schema information

Features:

- Separates the user from the complexity of MapReduce
- Structured – uses familiar concepts of tables, rows, columns and schema
- Partitions data with use of directory structures
- Connectivity with JDBC, Web GUI or Command Line Interface

Hive began at Facebook processing user and log data

Radiant Advisors

Extending Hadoop

HBASE

A distributed, column-oriented database using HDFS

Overview:

- Targeted for random real-time read/write access of (fairly) structured data
- Designed to support large tables of billions of rows and millions of columns
- HDFS delivers fully distributed and highly available

Hbase is modeled after Google's BigTable

Extending Hadoop

OTHER HADOOP PROJECTS

Avro	data serialization
Cassandra	multi-master database with no single point of failure
ZooKeeper	high performance coordination services of distributed applications
Mahout	scalable machine learning and data mining library
Chukwa	data collection system for managing large distributed systems

Radiant Advisors

History of Hadoop

CLOUDERA & MAPR

Cloudera Distribution Including Apache Hadoop

- Packages CDH ecosystem (Cloudera Distributed Hadoop)
- 100% Apache-licensed components
- Subscription offering
- Simplified installation and update
- Tested and compatible version
- Contributes to Apache Hadoop

MapR Technologies Inc. Distribution

- May 2011- announced the availability of an alternate filesystem for Hadoop
- Replaced the HDFS file system with a full random-access read/write file system
- Advanced features like snaphots and mirrors
- Gets rid of the single point of failure issue of the default HDFS NameNode.

Think of Red Hat for Linux

 v1.00.000a

Radiant Advisors

History of Hadoop

SOME PRODUCTION INSTALLATIONS

- Yahoo! has dozens of 4,000 node clusters and hundreds of petabytes
- Cloudera has 1,000 node customers
- Facebook just moved 30PB cluster July, 2011

Typical Node has 4-8 1TB drives per CPU

(Facebook likes 12 drives)

Radiant Advisors

SUMMARY

"We don't have Hadoop because we have big data.
We have big data because we have Hadoop."

Businesses view Hadoop and NoSQL as enabling technologies that will allow them to discover new opportunities and invent new business models

If you rely only on SQL in an organization, you may be missing opportunities for innovation

Radiant Advisors

Big Data and Hadoop Architectures in BI

BI ARCHITECTURE WITH HADOOP

SCALE OUT ARCHITECTURE

- Predictable linear performance, cost and scalability
- Parallelism through massive number of node servers
- Independent disk access for each nodes improves performance

Big Data Architectures

SCALE OUT ARCHITECTURE

MPP Scale-out:

- Shared Nothing architecture
- Each node has database engine
- Nodes communicate with each other to complete Map and Reduce functions
- Data distribute by column with value, hash or round-robin
- Data access in Master node
- Database includes SQL access

Hadoop Scale-out:

- Separate architecture for Distributed File System for loaders and Name node meta data
- Separate architecture for execution and accessing data
- Each data node communicates with other nodes
- Separate layers for SQL translation (PIG, Hive)

 v1.00.000a

Radiant Advisors

Big Data Architectures

MPP ARCHITECTURE

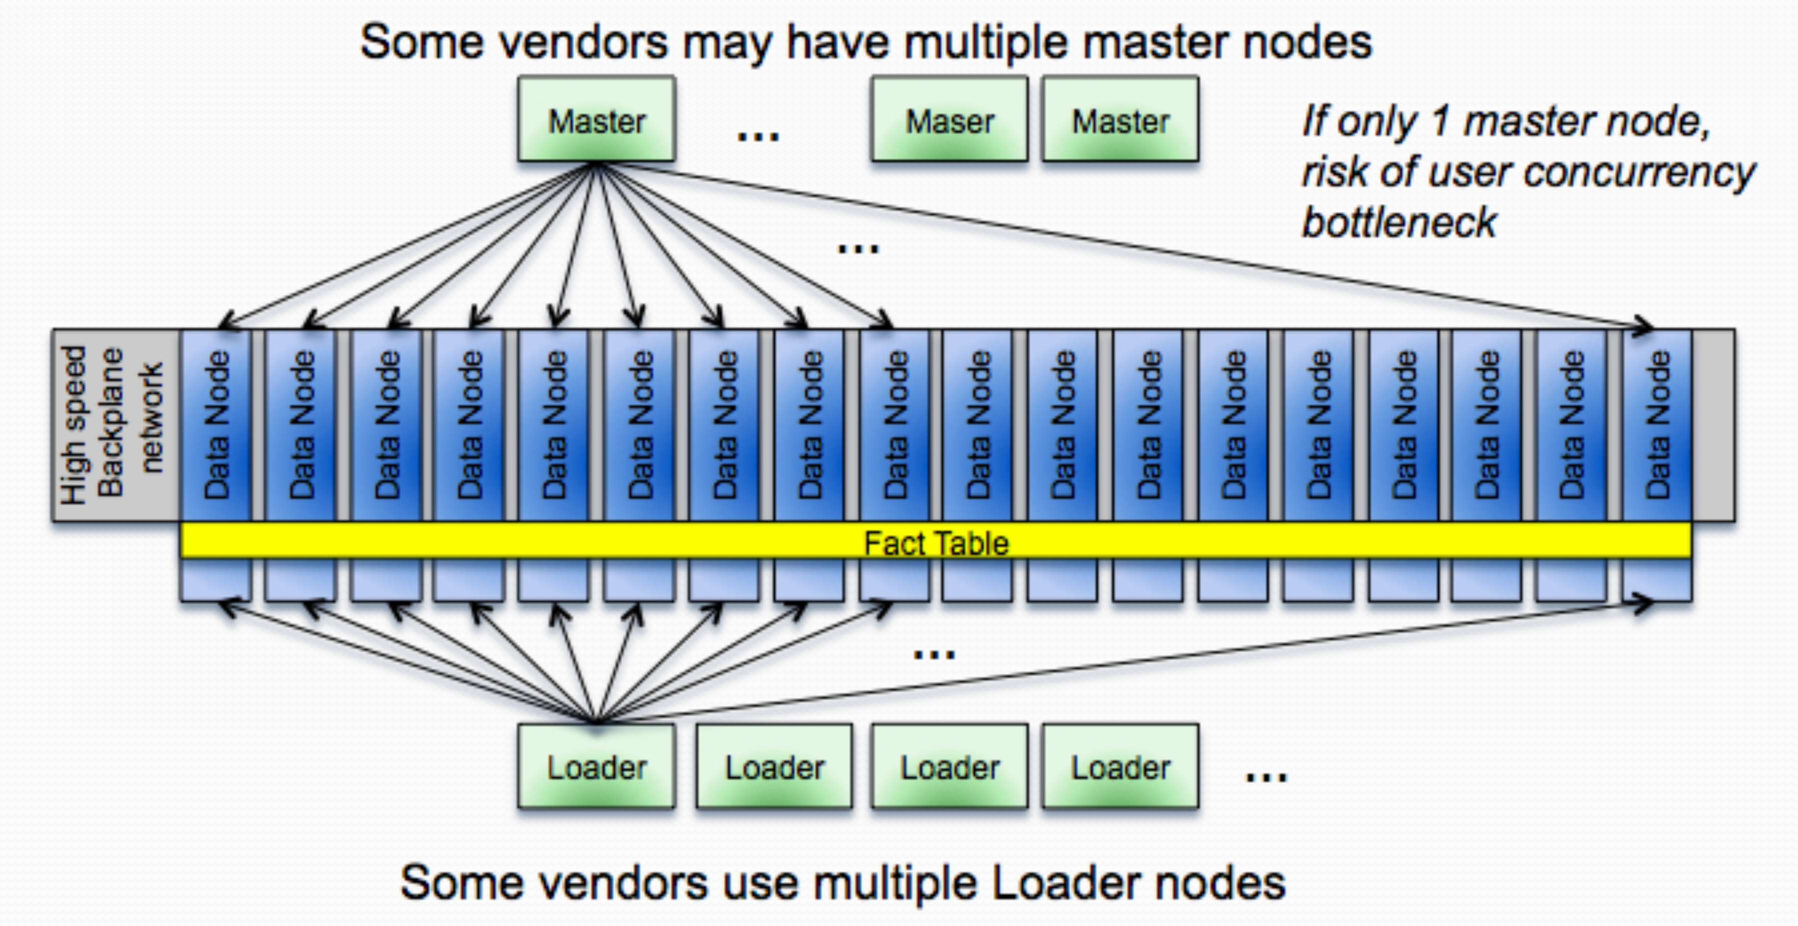

Radiant Advisors

MPP COMMUNICATION

When each node contains a slice of the fact table and a slice of the dimension, cross node joins are required to complete all row joins in the tables

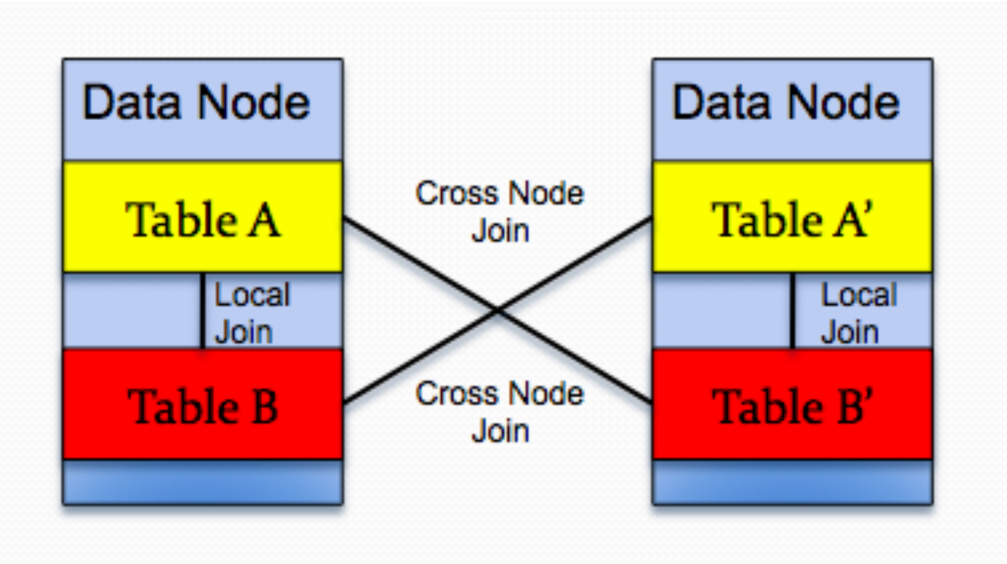

Radiant Advisors

MPP BROADCASTING TECHNIQUE

Specifically for Star schemas, some MPP vendors broadcast (duplicate copy) the dimensions into each node to keep all joins local since Dimensions are much smaller by order of magnitude

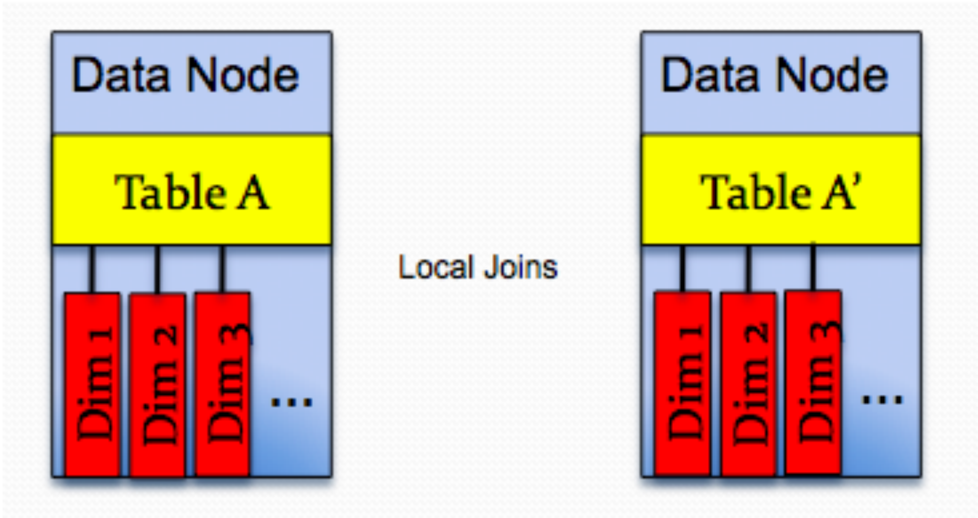

MPP USER DEFINED FUNCTIONS

- Most MPP databases are now MapReduce compatible
- UDFs allow the distributed database engines of MPP to execute User compiled code in the database at the node level for parallelism
  - C/C++, Perl, Python, Ruby, Java
- UDFs useful for complex big data algorithms that are not possible in SQL

Big Data Architectures

APPLIANCES

- Hardware + Software = Tightly integrated
- Plug and Play implementations
- Low administration
- Little to no tuning
- Typically some form of MPP database

- Appliance at what level?
  - System Appliance vs. Component Appliance

Big Data Architectures

COLUMNAR MPP

- Scale-out with RDBMS Structure
- Data written and stored by table column
- High compression rates from grouping data type in block
- Analytics is Read Mainly/Write Once

Most RDBMS offer Hybrid now

Big Data Architectures

COLUMNAR ORIENTATION

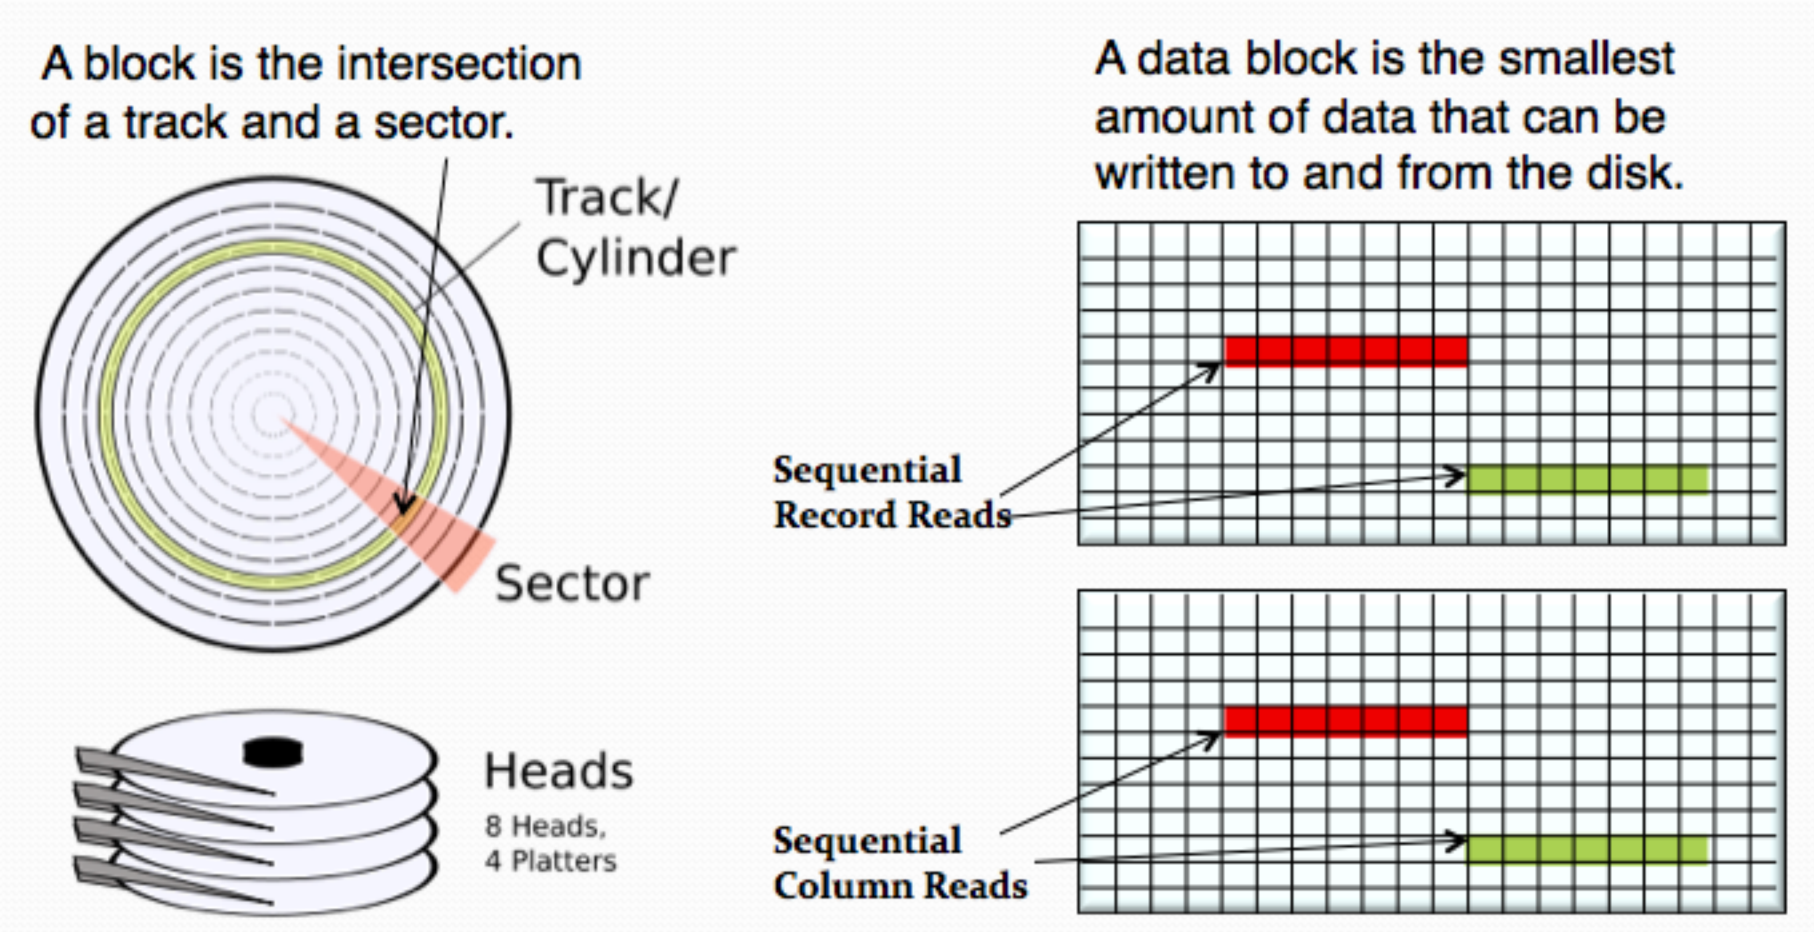

Big Data Architectures

COLUMNAR ORIENTATION

Customer Table

1	John	Kennedy	CO	50000
2	Betty	Smith	CA	55000
3	Sue	Hughes	TX	60000
4	Tom	Jones	NV	65000
5	David	Saunders	CO	70000

Row Serialization: (Written on disk)
1,John,Kennedy,CO,5000,2,Betty,Smith,CA,55000,3,Sue,Hughes,TX,60000,4,Tom,Jones,NV,65000…

Columnar Serialization: (Written on disk)
1,2,3,4,5,John,Betty,Sue,Tom,David,…50000,55000,60000,65000,70000

SELECT State, SUM(Sales) FROM Table1 GROUP BY state ORDER BY 2;

Radiant Advisors

Big Data Architectures

ROW VS COLUMN OPERATION

Reading Rows of Data for Columns

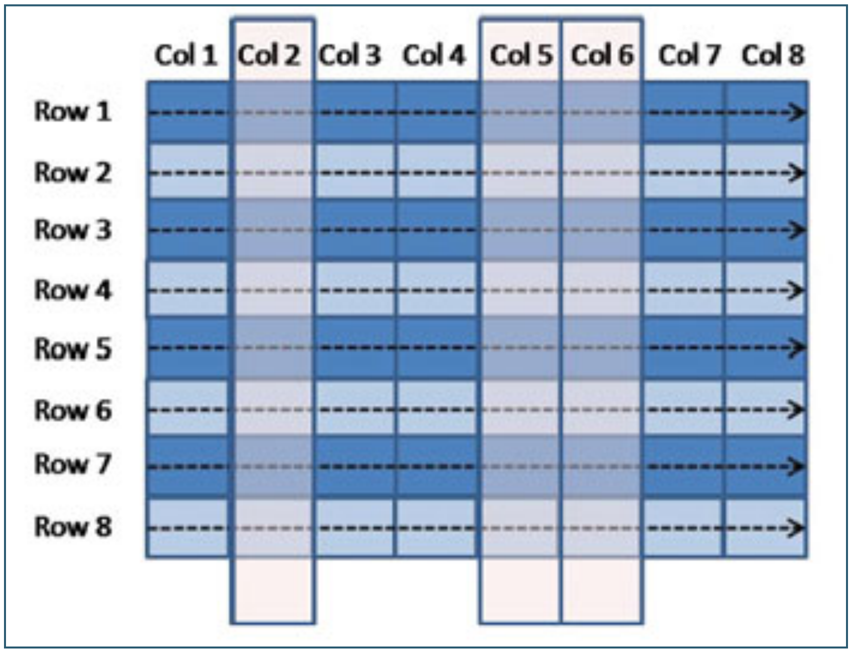

Reading only Columns you need

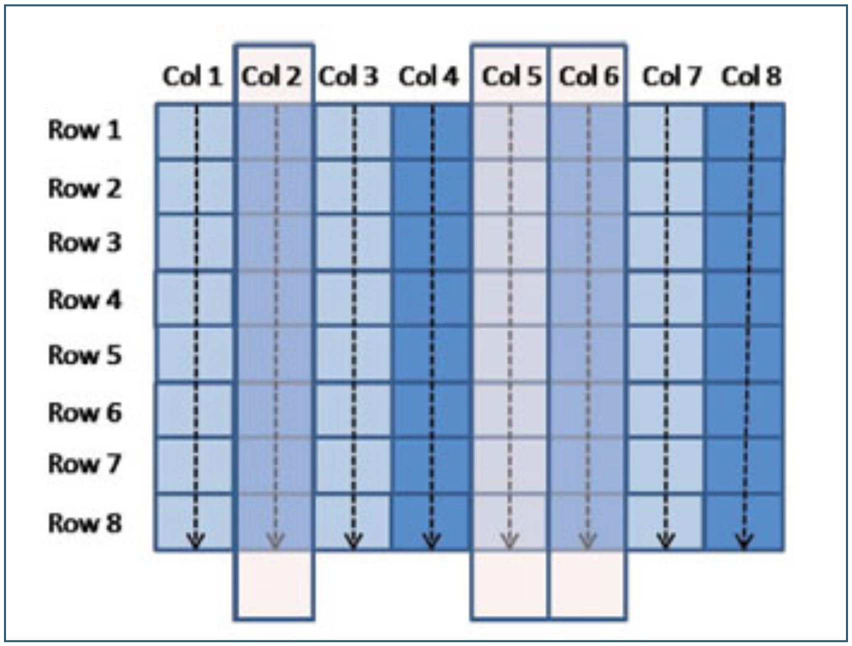

Radiant Advisors

Big Data Architectures

SCHEMA-LESS

Web Logs / App Logs

- Data which is best stored without a fixed schema (table design) since the data keeps changing
- Applications which are constantly changing or use self-describing data formats such as XML

Event Data

- Data which contains standard or fixed data structures which allows it to be stored easily in RDBMS
- However, high volume Event Data presents the challenges of loading, managing and accessing this data

Radiant Advisors

Mixed Technology Architectures

HADOOP INTEGRATION

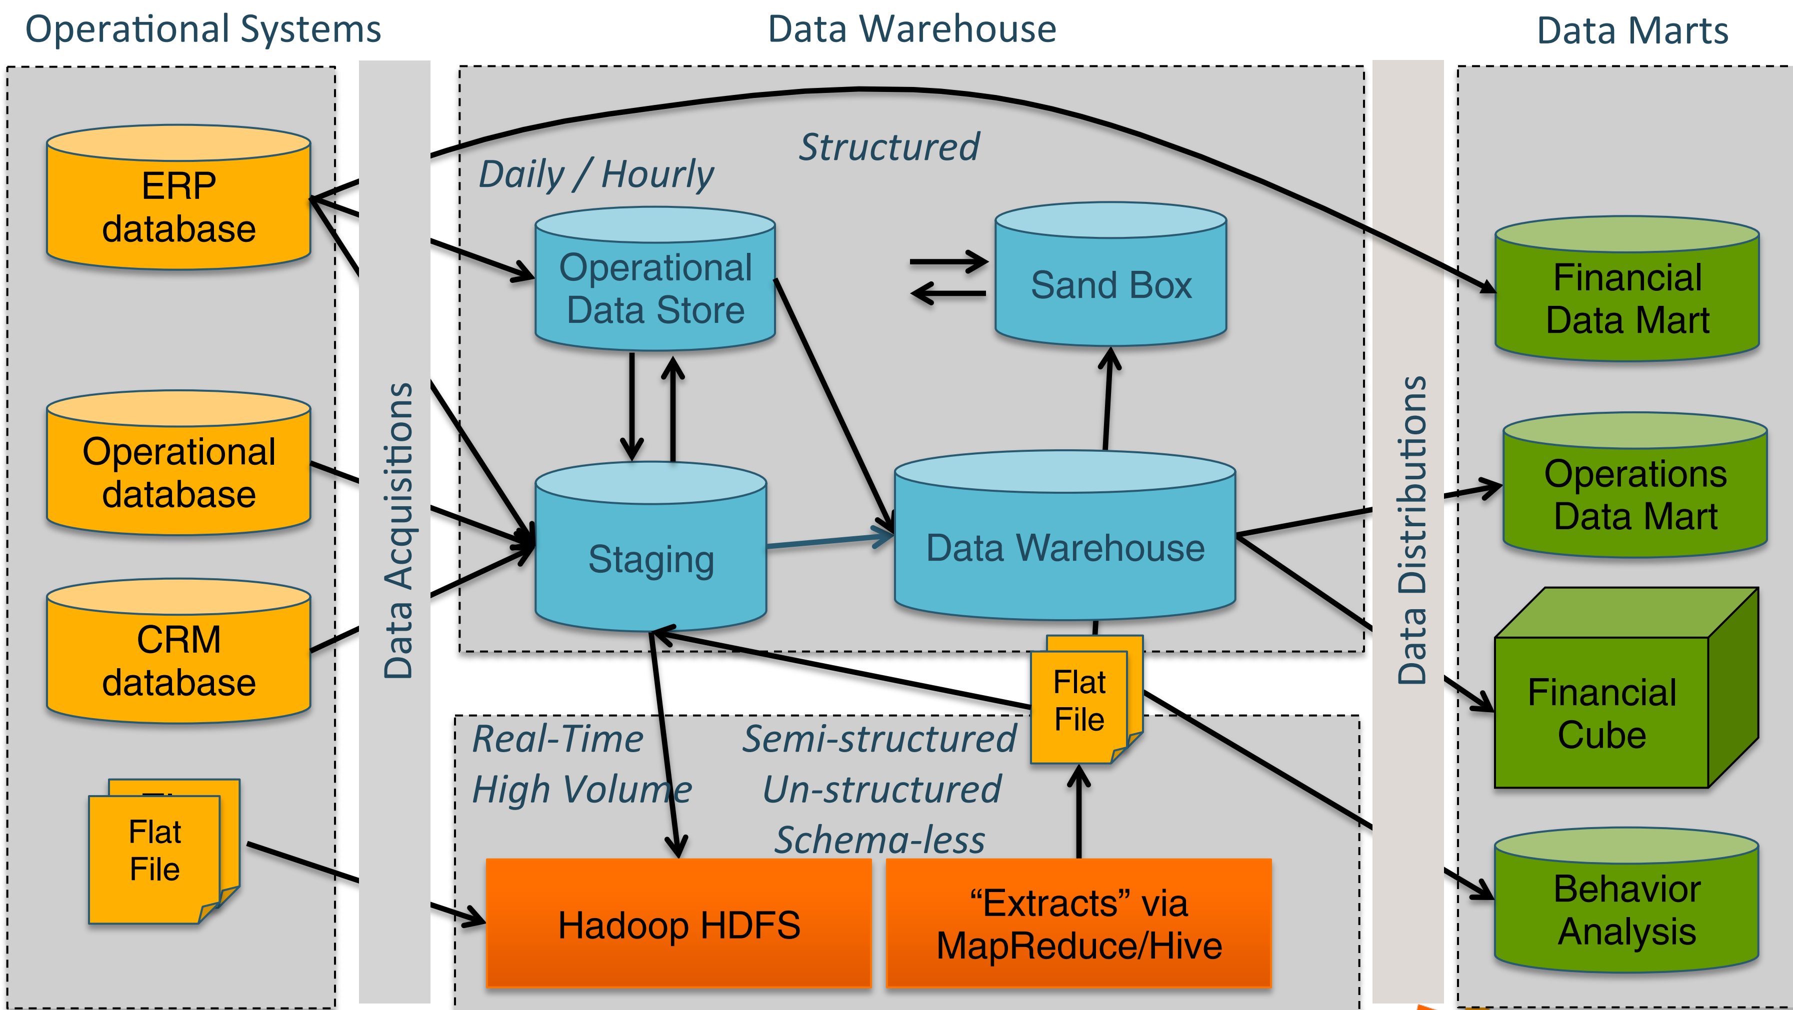

Mixed Technology Architectures

DATA CLASSES FOR BI

Traditional Data Types:

- Business Operations
- Reference Data Subjects
- Event data as Orders, Tickets

Large Enterprise Data Types:

- High Volume Structured Events
- User Driven
- Transactions, Calls, etc

Data Exhaust Data Types:

- Machine generated data
- Instrumented or Sensors
- Semi-structured
- Schema-less
- User behavior on websites
- Non-record based (Unstructured)
- Text based documents

Radiant Advisors

Mixed Technology Architectures

DATA STORE PURPOSE

Traditional Data Types:

- Staging
- Subject oriented data warehouses
- Data marts

Large Enterprise Data Types:

- Data marts
- Data mining
- Operational data marts

Data Exhaust Data Types:

- Staging
- Hadoop Distributed File System
- Data marts
- Sandbox

Mixed Technology Architectures

DATA STORE TECHNOLOGIES

Traditional Data Types:

- RDMBS
- In-memory databases
- OLAP/MDDB

Data Exhaust Data Types:

- Hadoop/MapReduce
- MapR
- HPCC

Large Enterprise Data Types:

- Massively Parallel Processing databases
- MPP based appliances
- Columnar databases

Radiant Advisors

Mixed Technology Architectures

HADOOP AS INDEPENDENT

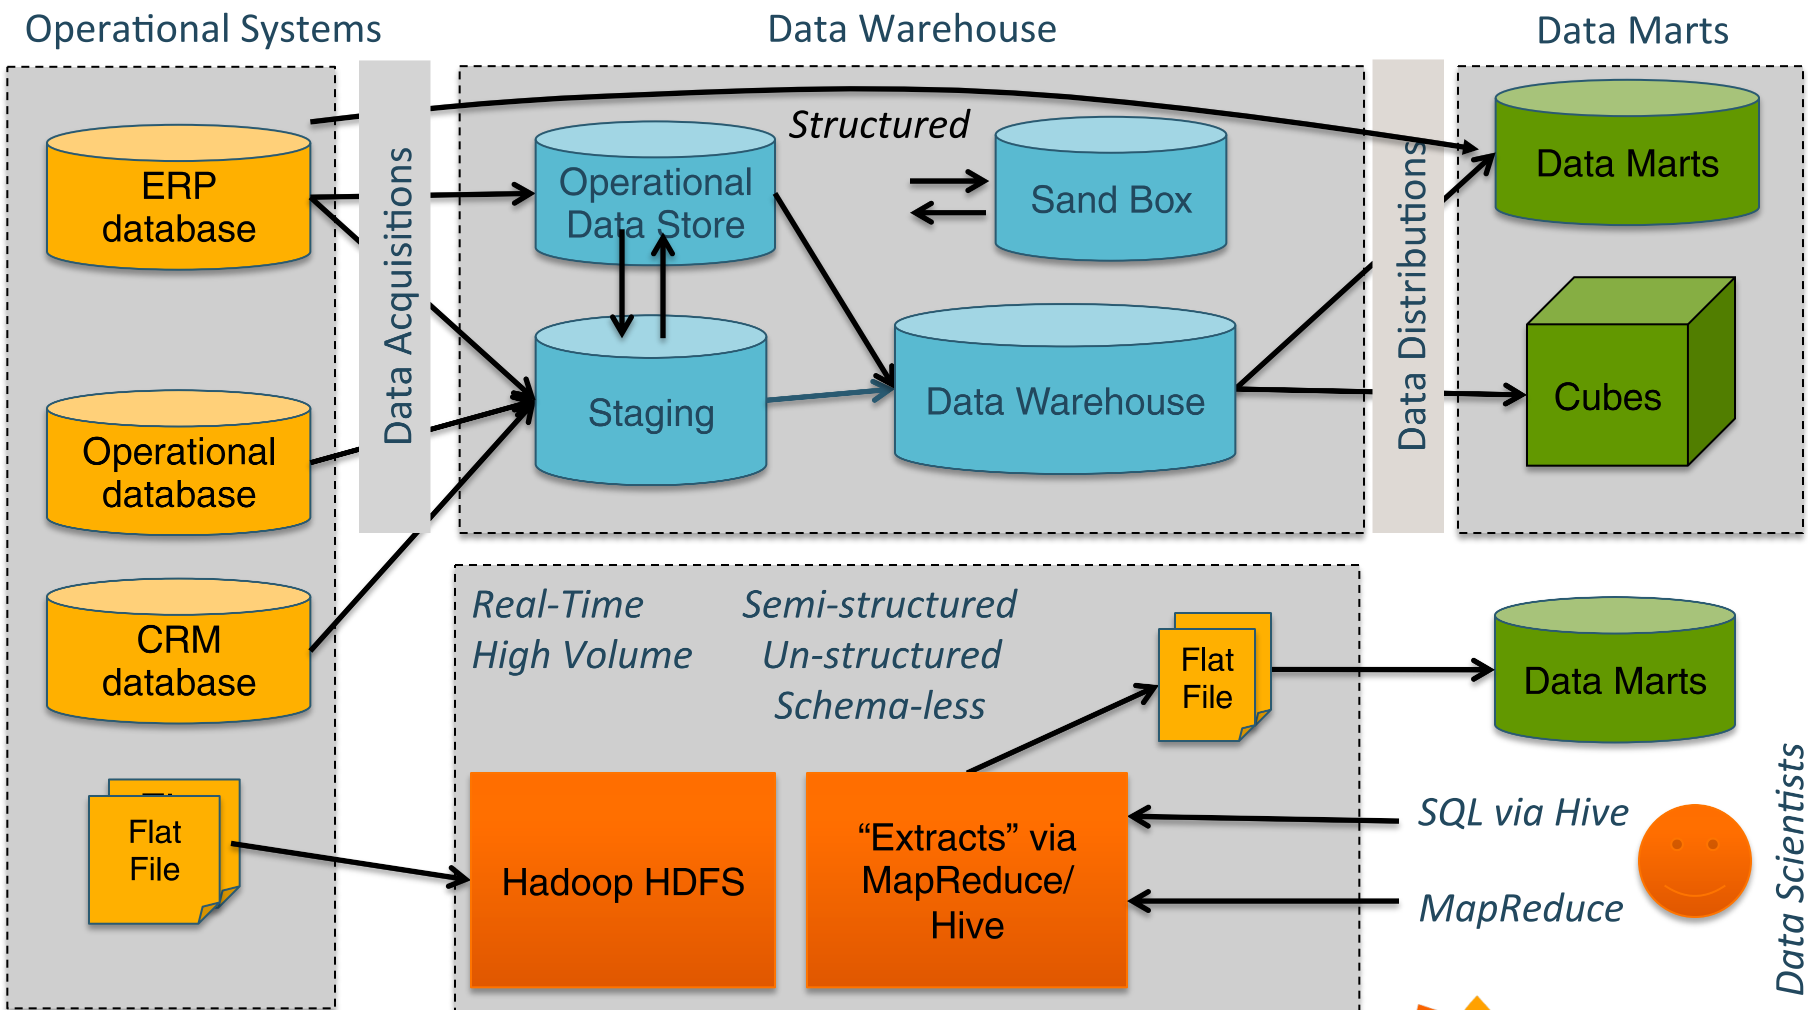

Radiant Advisors

Mixed Technology Architectures

HADOOP AS DATA WAREHOUSE

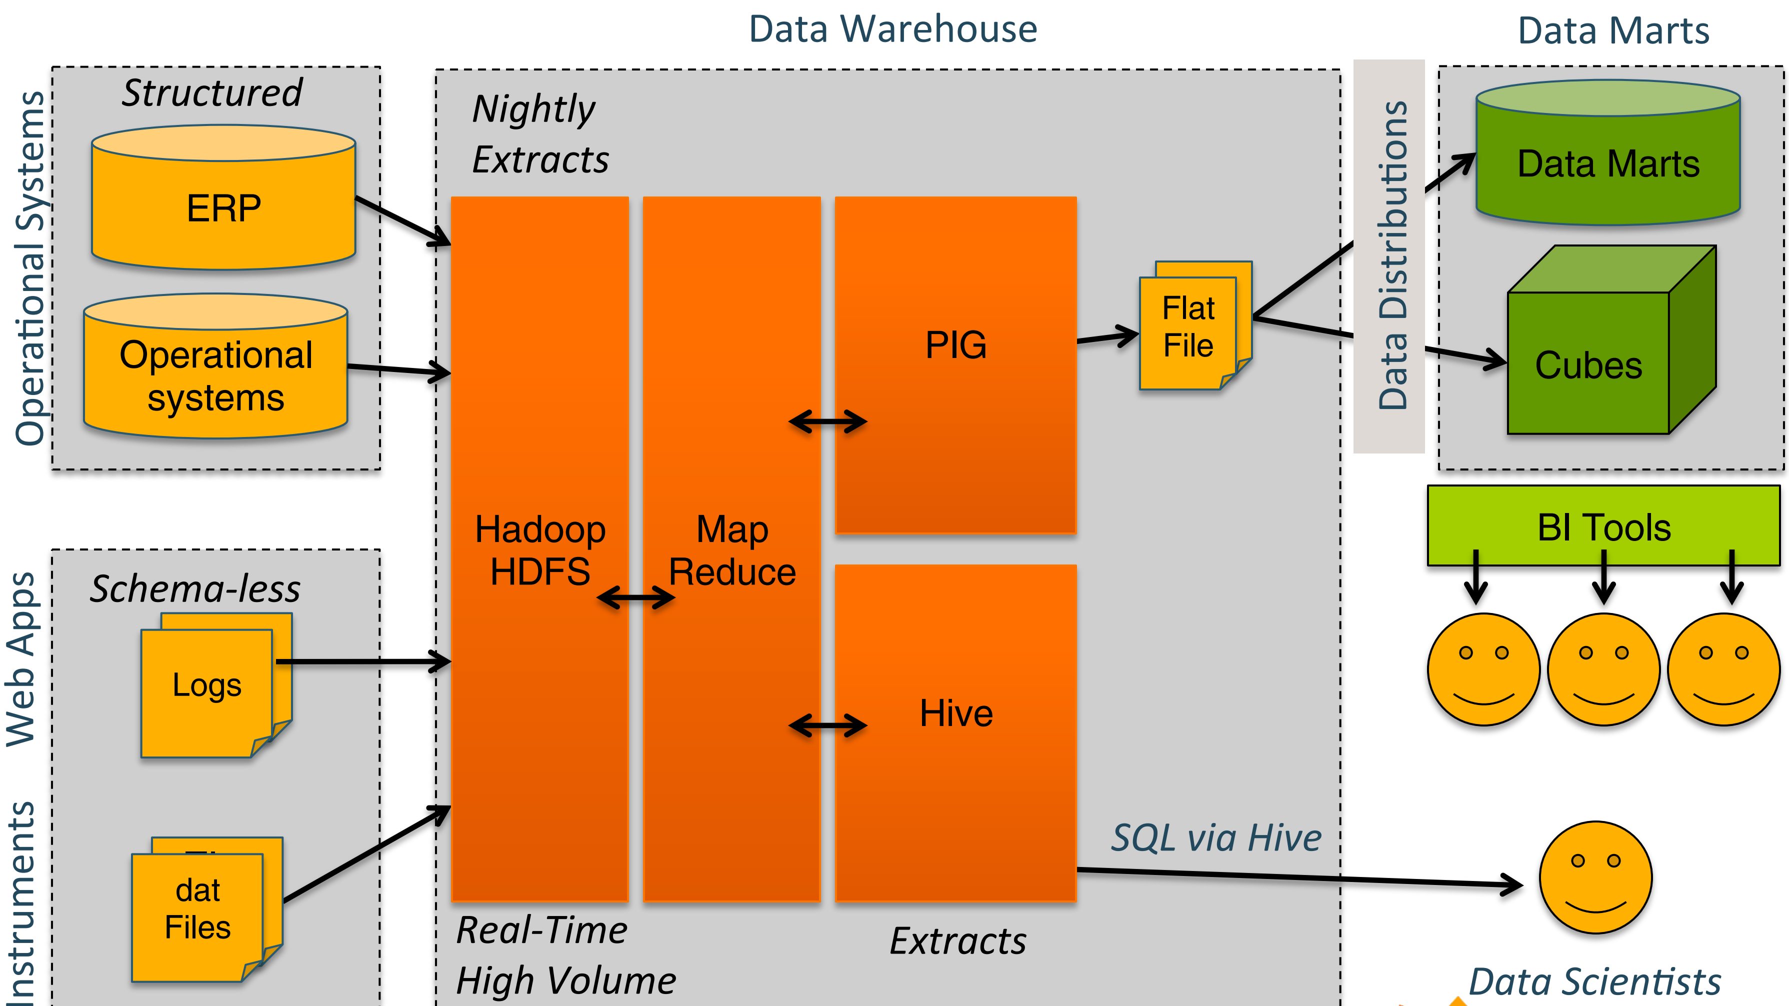

Radiant Advisors

Mixed Technology Architectures

ANALYTIC TECHNOLOGY CATEGORIES

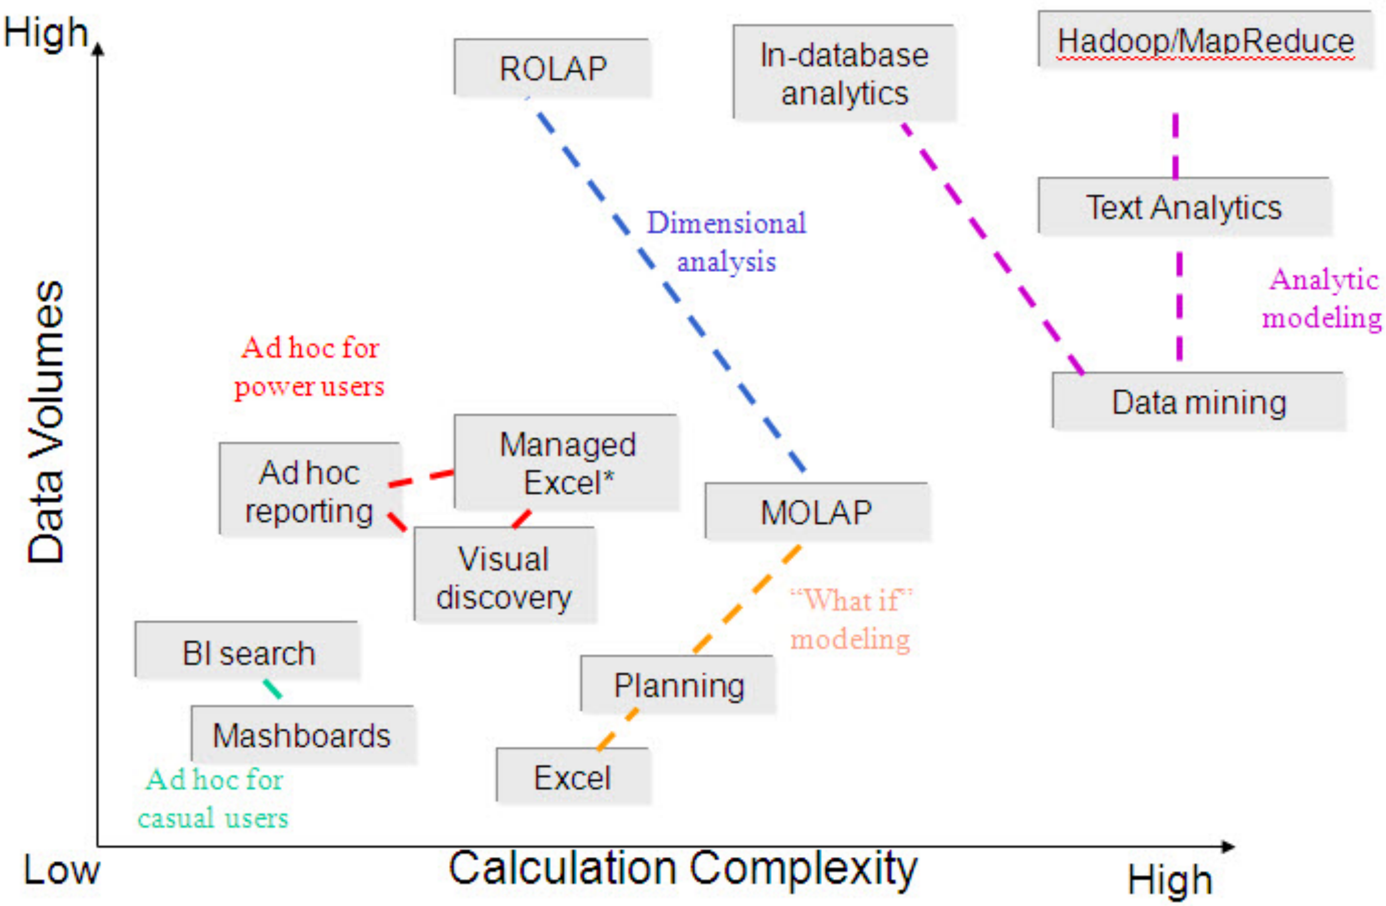

 v1.00.000a

Radiant Advisors

Analytic Architectures

DATA STORE TECHNOLOGIES

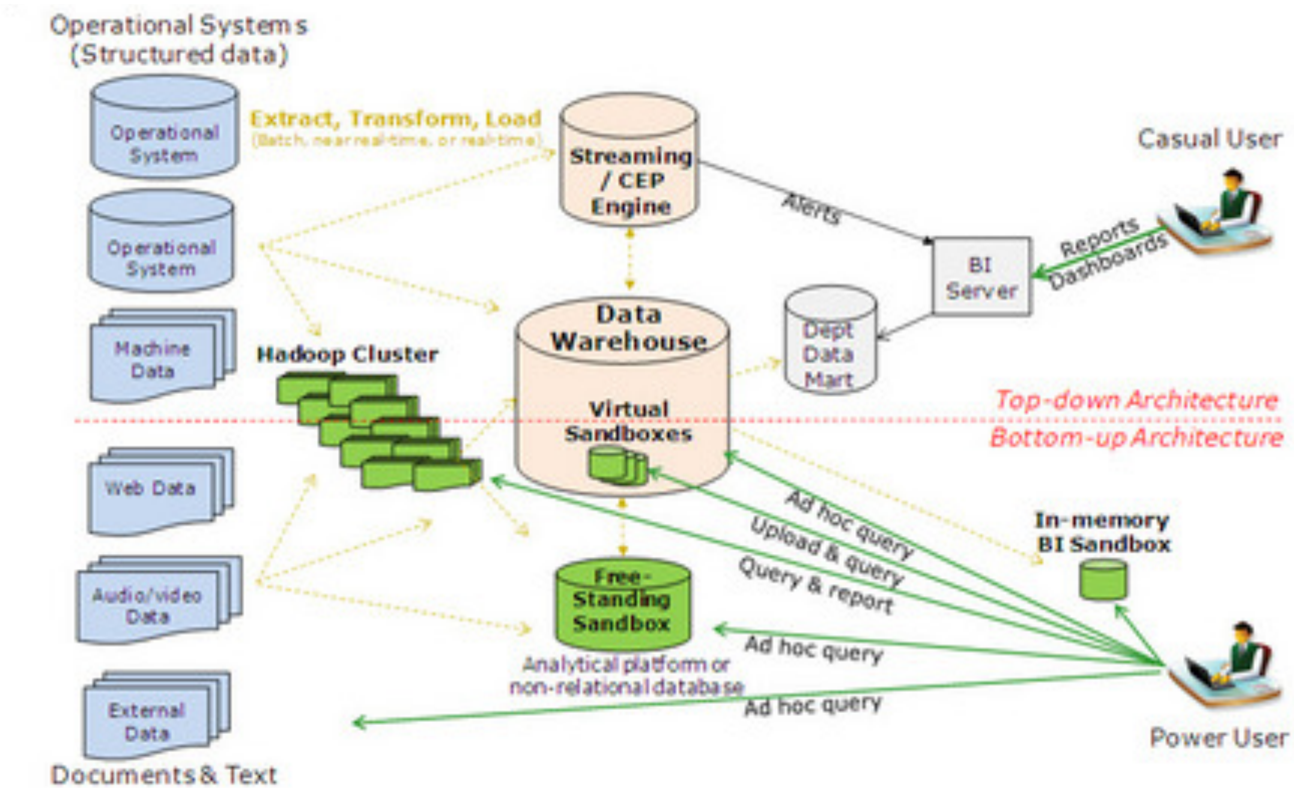

The next-generation BI architecture is more analytical, giving power users greater options to access and mix corporate data with their own data via various types of analytical sandboxes. It also brings unstructured and semi-structured data fully into the mix using Hadoop and nonrelational databases.

Radiant Advisors

Analytic Architectures

TYPES OF ANALYTIC SANDBOXES

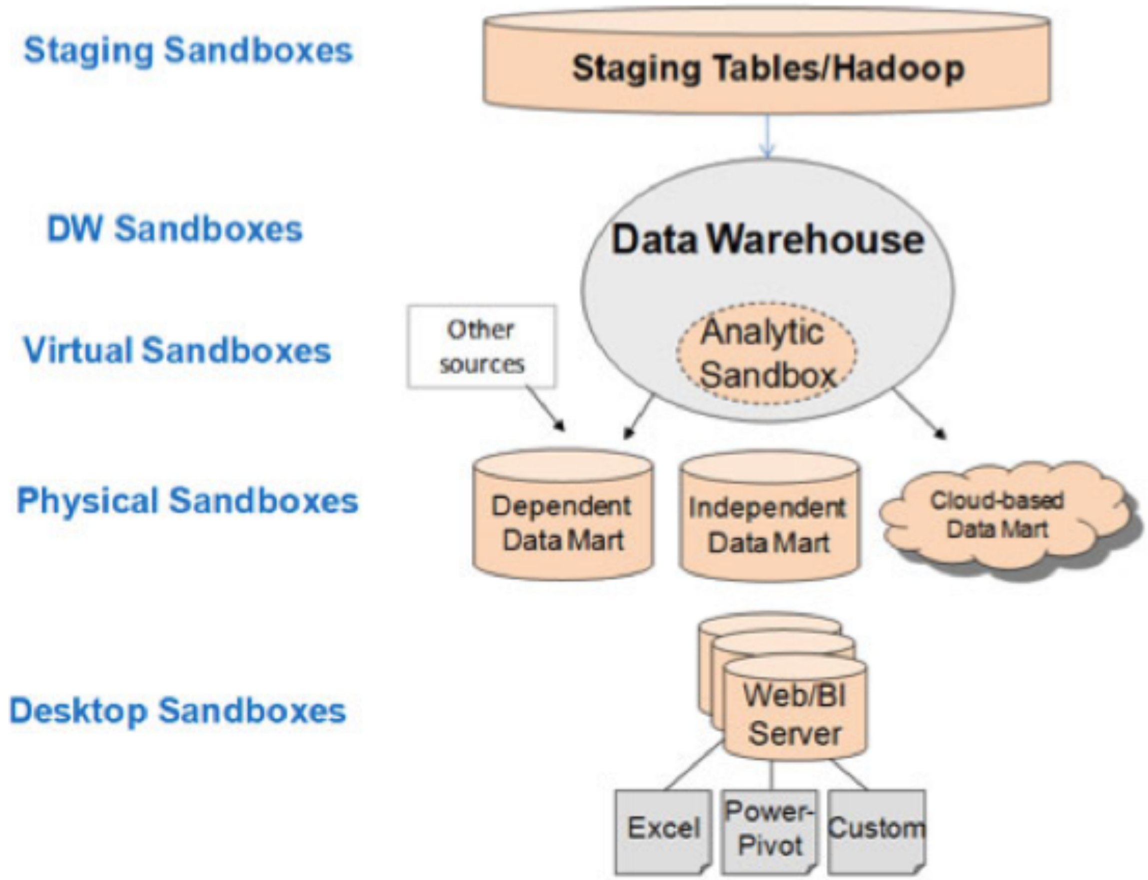

Radiant Advisors

ANALYTIC TECHNOLOGIES

In-Database Operation

- DB must be able to execute models
- PMML or MapReduce
- User Defined Functions
- Should be MPP database for scalability and performance

Non-Database Operation

- External execution to database
- MapReduce like environment (MPP)
- Exports to PMML for compatibility

Analytic Architectures

ANALYTIC TECHNOLOGIES

Technology	Description	Vendor/Product
MPP analytical databases	Row-based databases designed to scale out on a cluster of commodity servers and run complex queries in parallel against large volumes of data.	Teradata Active Data Warehouse, Greenplum (EMC), Microsoft Parallel Data Warehouse, Aster Data (Teradata), Kognitio, Dataupia
Columnar databases	Database management systems that store data in columns, not rows, and support high data compression ratios.	ParAccel, Infobright, Sand Technology, Sybase IQ (SAP), Vertica (Hewlett-Packard), 1010data, Exasol, Calpont
Analytical appliances	Preconfigured hardware-software systems designed for query processing and analytics that require little tuning.	Netezza (IBM), Teradata Appliances, Oracle Exadata, Greenplum Data Computing Appliance (EMC)
Analytical bundles	Predefined hardware and software configurations that are certified to meet specific performance criteria, but the customer must purchase and configure themselves.	IBM SmartAnalytics, Microsoft FastTrack
In-memory databases	Systems that load data into memory to execute complex queries.	SAP HANA, Cognos TM1 (IBM), QlikView, Membase
Distributed file-based systems	Distributed file systems designed for storing, indexing, manipulating and querying large volumes of unstructured and semi-structured data.	Hadoop (Apache, Cloudera, MapR, IBM, HortonWorks), Apache Hive, Apache Pig
Analytical services	Analytical platforms delivered as a hosted or public-cloud-based service.	1010data, Kognitio
Nonrelational	Nonrelational databases optimized for querying unstructured data as well as structured data.	MarkLogic Server, MongoDB, Splunk, Attivio, Endeca, Apache Cassandra, Apache Hbase
CEP/streaming engines	Ingest, filter, calculate, and correlate large volumes of discrete events and apply rules that trigger alerts when conditions are met.	IBM, Tibco, Streambase, Sybase (Aleri), Opalma, Vitria, Informatica

Radiant Advisors

REFERENCES & MORE INFO

Books:

- Hadoop in Action
- Beautiful Data, Edited by Toby Segaran & Jeff Hammerbacher, O'Reilly Media 2009
- Blown to Bits, Your Life, Liberty and Happiness after the Digital Explosion, by Hal Abelson, Ken Leden, Harry Lewis. Addison Wesley published 2008.

Websites:

www.apache.com

www.strata.com

Articles & White Papers:

- Analytic Architectures: Approaches to Supporting Analytics User Workloads. BeyeResearch by Wayne Eckerson
- The CIO's Guide to NoSQL, Dan McCreary and William McKnight, June 2011. Sponsored by Dataversity and 2011 NoSQL Now! Conference.
- "Data, data, everywhere. The Economist Special Report on managing information. Feb 27, 2010
- The End of Theory: The Data Deluge Makes the Scientific Method Obsolete. By Chris Anderson, Wired Magazine Issue 16.07 Jun 23, 2008
- The Web is dead. By Chris Anderson, Wired Magazine Issue 18.09 Sep 2010
- Release 2.0. O'Reilly Radar Issue 2.0.11 Feb 2009
- What is Data Science? By Mike Loukides, O'Reilly Radar Report 2010

Radiant Advisors

For more information

Visit us at: www.radiantadvisors.com

Follow us: Twitter: @RadiantAdvisors

RSS: feed://radiantadvisors.com/feed/

Email us at: info@RadiantAdvisors.com

THANK YOU!

Be Radiant!

 v1.00.000a