BIG DATA SYSTEM DEVELOPMENT: AN EMBEDDED CASE STUDY WITH A GLOBAL OUTSOURCING FIRM

Prof. Hong-Mei Chen
IT Management, Shidler College of Business
University of Hawaii at Manoa, USA

Prof. Rick Kazman
IT Management, Shidler College of Business
University of Hawaii at Manoa, USA
Software Engineering Institute, Carnegie Mellon University, USA

Serge Haziyev, Olha Hrytsay
SoftServe Inc.
Austin, TX, USA

OUTLINE

- Research Motivation
- Research Foundations
- Research Method
- Results
- Future Research Directions
- Conclusions

Big Data: Big Promise

- Big hype…
- Big data is the new oil
- Big data is the new gold

????????????

HOW??


BIG
DATA


Big Data


BIG DATA
Text analytics
Semantic Metadata

Challenges

- 5V requirements
- Proliferation of Big Data Technology
- Rapid Big Data Technology Changes
- Complexity
- Paradigm Shifts
- Short history of big data system development in Enterprises

2013 CIO Survey

Big Data Survey http://visual.ly/cios-big-data (Jan. 2013)

55% of big data projects were not completed

Gartner Survey (Dec. 2014): Big Data Investment Grows but Deployments Remain Scarce in 2014

- Hype is wearing thin
- Only 13% of respondents said their IT organizations put big data projects into production this year, but that's 5% higher than last year.
- 24% of those polled voted against the use of big data technologies in their business.

“2013 was the year of experimentation and early deployment; so is 2014”

- 73 percent of respondents have invested or plan to invest in big data in the next 24 months, up from 64 percent in 2013.
- Like 2013, much of the work today revolves around strategy development and the creation of pilots and experimental projects.

- Note: The Gartner survey of 302 Gartner Research Circle members worldwide, which was conducted in June 2014.

Research Objectives

- ✓ To help enterprises navigate through uncharted waters and be better equipped for their big data endeavors.
- ✓ To uncover methodological voids and provide practical guidelines.

Research Questions

1. How does big data system development (processes and methods) differ from "small" (traditional, structured) data system development?
2. How can existing software architecture approaches be extended or modified to address new requirements for big data system design?
3. How can data modeling/design methods in traditional structured database/datawarehouse development be extended and integrated with architecture methods for effective big data system design?

“Small” Data System Development

- **ANSI Standard 3-layer DBMS Architecture**
  - Clear Data-Program Independence (logical and physical data independence)
- **Well-established RAD design process**
  - Iterative design of 7 phases
  - Clear separation of each design phase
  - Mature conceptual design tools: ER, UML, etc.
- **Relational model dominance (95% market)**
  - Relational model easy to understand
  - SQL easy to use, standardized
- **Architecture Choice is relatively simple**
  - N-tier client-server design

Data/program Independence: ANSI 3-Layer DBMS Architecture (1980s)

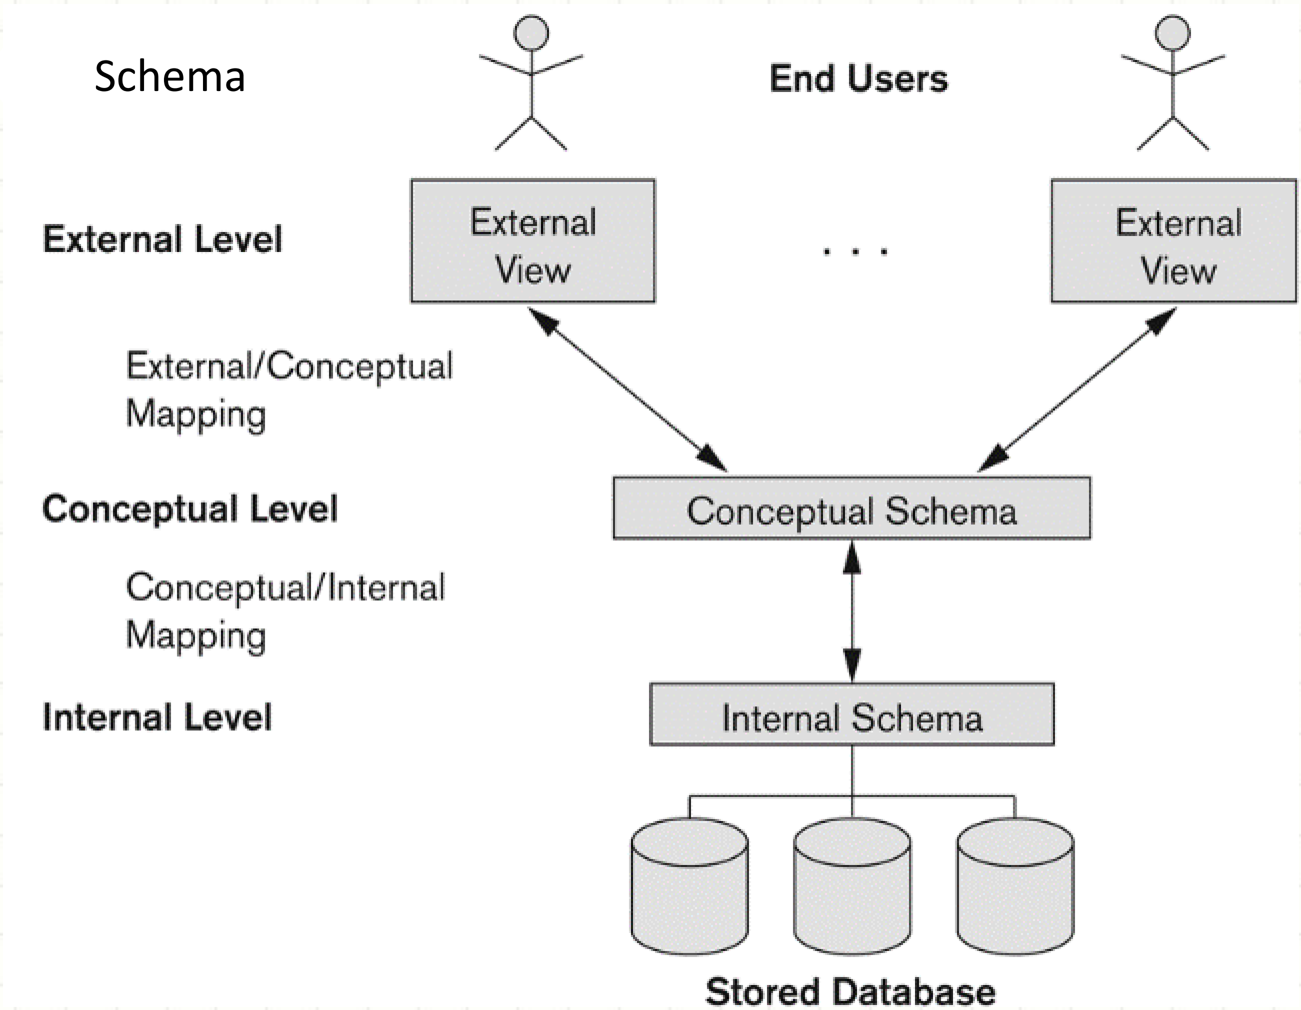

Architecture Design is critical and complex in Big data System Development

I. Volume: Distributed and scalable architecture
II. Variety: Polyglot persistence architecture
III. Velocity: Complex Event processing + → Lambda Architecture
IV. Veracity: Architecture design for understanding the data sources and the cleanliness, validation of each
V. Value: New architecture for hybrid, agile Analytics, big data analytics cloud, integrating the new and the Old (EDW, ETL)
VI. Integration: Integrating separate architectures addressing each of the 5V challenges

Research Questions

1. How does big data system development (processes and methods) differ from "small" (traditional, structured) data system development?
2. How can existing software architecture approaches be extended or modified to address new requirements for big data system design?
3. How can data modeling/design methods in traditional structured database/datawarehouse development be extended and integrated with architecture methods for effective big data system design?

Research Method

- ✓ **Case study research is deemed suitable**:
  - system development, be it big or small data, cannot be separated from its organizational and business contexts.
  - "How" and "Why" research questions.
  - the research is largely exploratory
- ✓ **Multiple cases**: increase methodological rigor
- ✓ **Collaborative Practice Research**
  - SSV, in the outsourcing industry
  - who has successfully deployed 10 big data projects that can be triangulated

  ➔ Embedded Case Study

Reasons for selecting an outsourcer

- Outsourcing is an important and common means to realize a big data strategy
- Big data professional service is the largest segment of big data market and continues to grow.
- Outsourcing mitigates shortages of skills and expertise in the areas where they want to grow.

Big Data Market is Expected to Grow Rapidly

Wikibon

Big Data Market Forecast by Sub-Type, 2011-2017
(in $US billions)

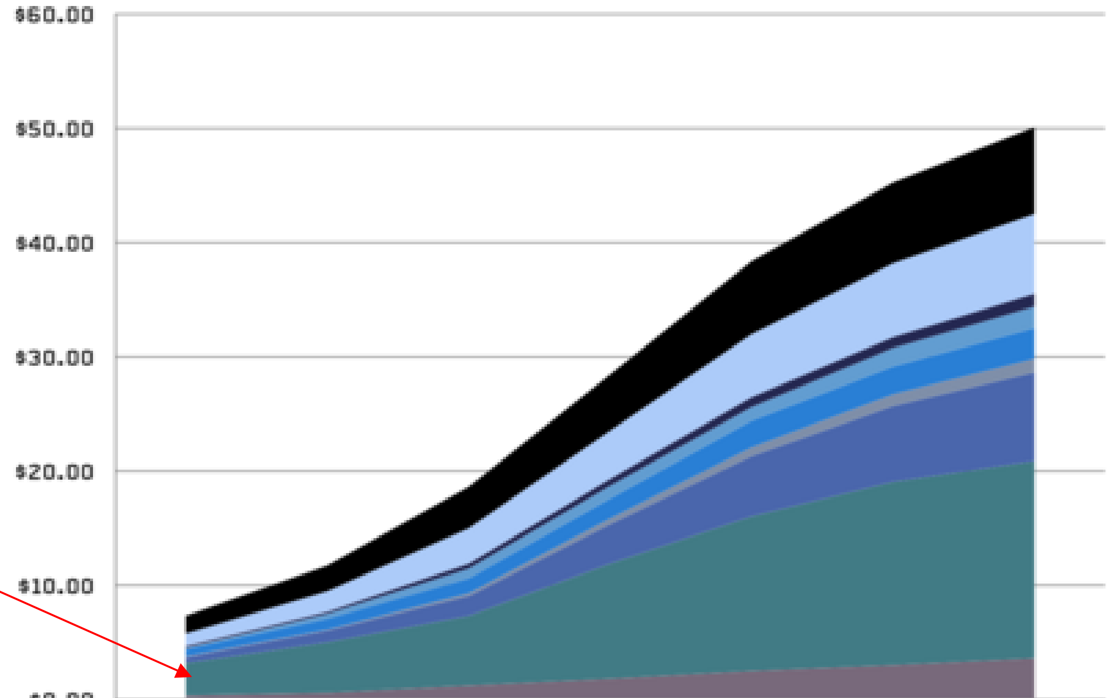

Professional service is the largest growing segment

	2011	2012	2013	2014	2015	2016	2017
Compute	$1.53	$2.29	$3.65	$4.92	$6.40	$7.10	$7.60
Storage	$1.10	$1.75	$3.09	$4.20	$5.50	$6.40	$6.95
Networking	$0.15	$0.23	$0.42	$0.65	$0.85	$1.01	$1.15
Infrastructure Software	$0.14	$0.44	$0.83	$1.08	$1.25	$1.60	$1.90
SQL	$0.62	$0.88	$1.31	$1.75	$2.25	$2.45	$2.70
NoSQL	$0.07	$0.13	$0.29	$0.50	$0.80	$1.00	$1.20
Apps & Analytics	$0.52	$0.99	$1.69	$3.45	$5.29	$6.65	$7.75
Professional Services	$2.80	$4.42	$6.15	$10.10	$13.50	$16.00	$17.20
Cloud	$0.36	$0.62	$1.19	$1.82	$2.52	$3.05	$3.65

Collaborative Practice Research (CPR) Steps in an Iteration

1) Appreciate problem situation
2) Study literature
3) Develop framework
4) Evolve Method
5) Action
6) Evaluate experiences
7) Exit
8) Assess usefulness
9) Elicit research results

Collaborative Practice Research (CPR)

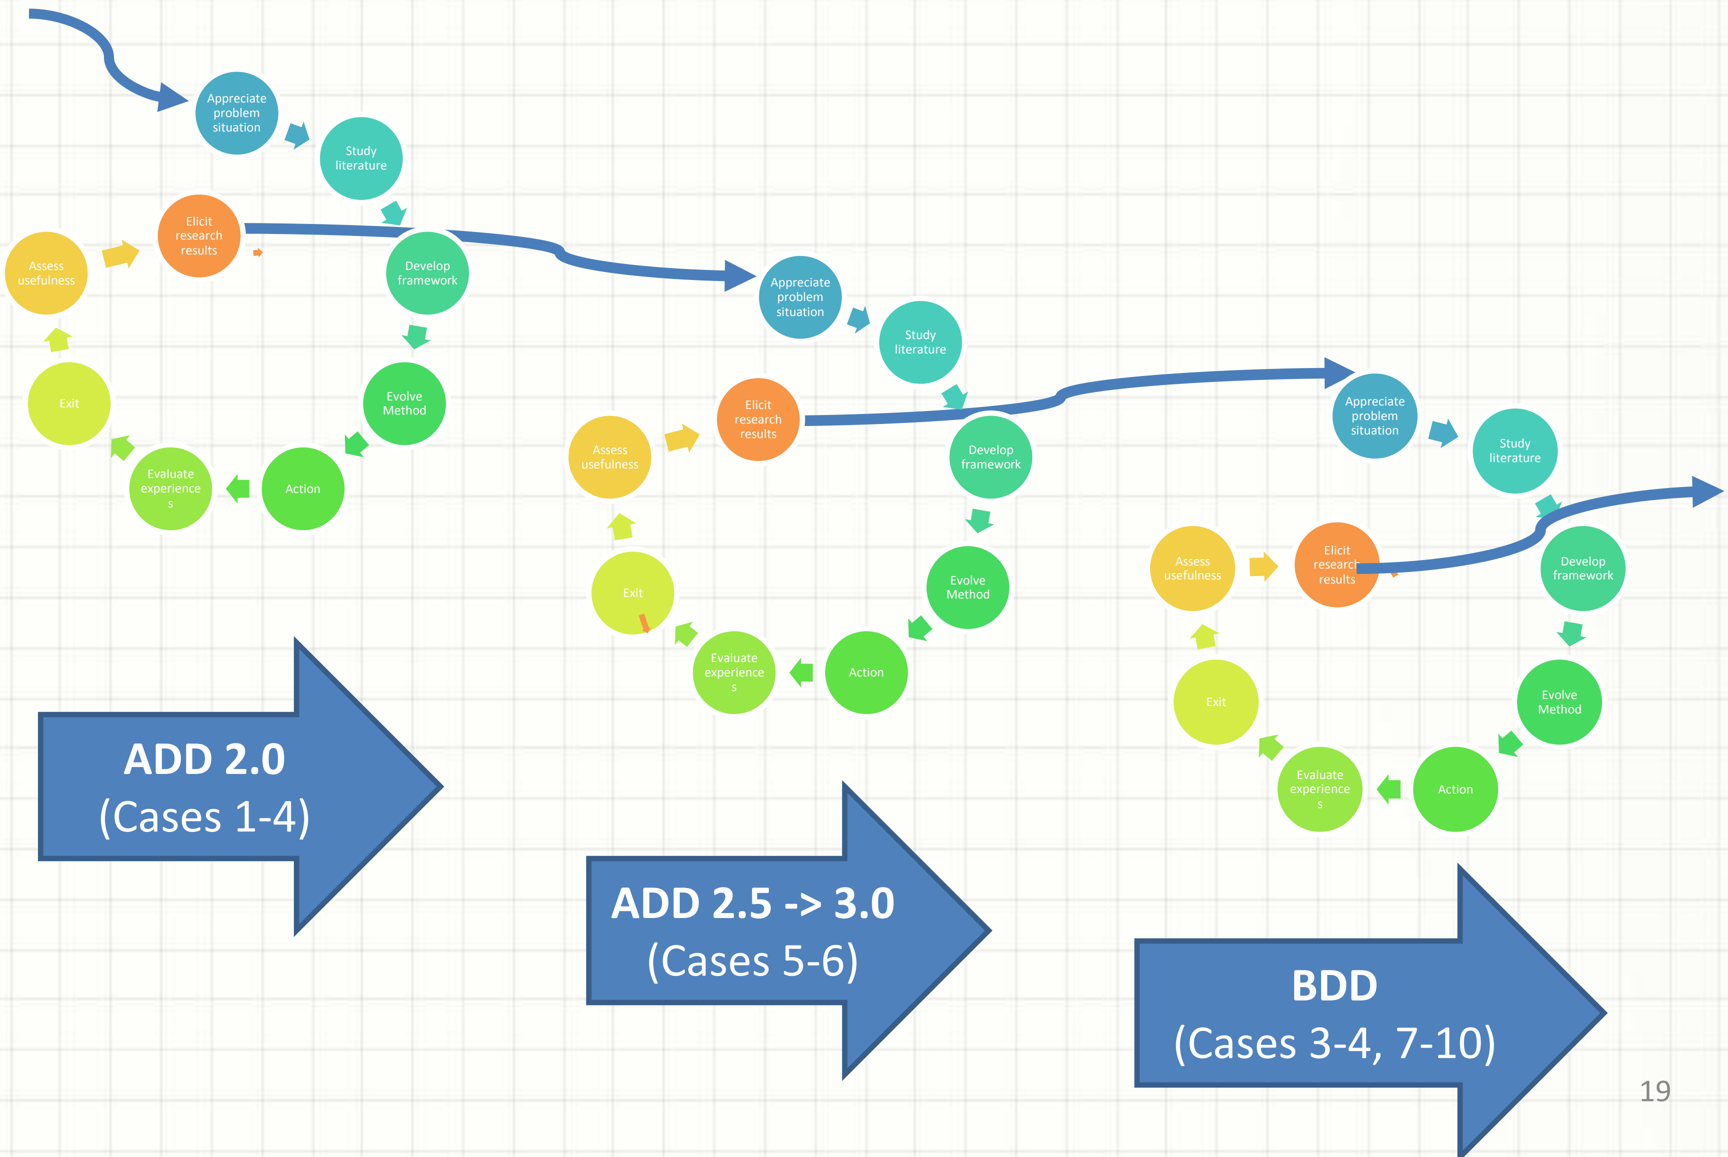

ADD

- ADD (Attribute-Driven Design) is an architecture design method "driven" by quality attribute concerns
  - Version 1.0 released 2000 by SEI.
  - Version 2.0 released November 2006 (on Current SEI site)
  - Version 2.5 published in 2013 by the researcher team
  - Version 3.0 to be published in 2016 by the researcher team.
- The method provides a detailed set of steps for architecture design
  - enables design to be performed in a systematic, repeatable way
  - leading to predictable outcomes.

Embedded Cases 1-3

Case #	Business goals	Start	Big data	Technologies	Challenges
1 **Network Security, Intrusion Prevention** US MNC IT corp. (Employees > 320,000)	• Provide ability for security analysts to improve intrusion detection techniques; • Observe traffic behavior and make infrastructure adjustments: • Adjust company security policies • Improve system performance	Late 2010, 8.5 month	• Machine generated data - 7.5BLN event records per day collected from IPS devices • Near real-time reporting • Reports which "touch" billions of rows should generates < 1 min	•ETL - Talend •Storage/DW – InfoBright EE, HP Vertica •OLAP – Pentaho Mondrian •BI – JasperServer Pro	• High throughput, different device data schemas (versions) • keep system performance at required level when supporting IP/geography analysis: avoid join. • Keep required performance for complex querying over billions rows
2 **Anti-Spam Network Security System** US MNC Networking equipment corp. employees > 74,000	• Validation of the new developed set of anti-spam rules against the large training set of known emails • Detection of the best anti-spam rules in terms of performance and efficacy	2012-2013	• 20K Anti-spam rules • 5M email training set • 100+ Nodes in Hadoop Clusters	• Vanilla Apache Hadoop (HDFS,MapReduce,Oozie,Zookeeper) • Perl/Python • SpamAssassin • Perceptron	• MapReduce was written on Python and Hadoop Streaming was used. The challenge was to optimize jobs performance. • Optimal Hadoop cluster configuration for maximizing performance and minimize map-reduce processing time
3 **Online Coupon Web Analytics Platform** US MNC: World's largest coupon site, 2014 Revenue > US$200M	• In-house Web Analytics Platform for Conversion Funnel Analysis, marketing campaign optimization, user behavior analytics • clickstream analytics, platform feature usage analysis	2012, Ongoing	• 500 million visits a year • 25TB+ HP Vertica Data Warehouse • 50TB+ Hadoop Cluster • Near-Real time analytics (15 minutes is supported for clickstream data)	• Data Lake - (Amazon EMR) /Hive/Hue/MapReduce/Flume/Spark • DW: HP Vertica, MySQL • ETL/Data Integration – custom using python • BI: R, Mahout, Tableau	• Minimize transformation time for semi-structured data • Data quality and consistency • complex data integration • fast growing data volumes, • performance issues with Hadoop Map/Reduce (moving to Spark)

Embedded Cases 4-6

Case #	Business goals	Start	Big data	Technologies	Challenges
4 **Social Marketing Analytical Platform** US MNC Internet marketing (user reviews) '14 Revenue > US$ 48M	• Build in-house Analytics Platform for ROI measurement and performance analysis of every product and feature delivered by the e-commerce platform; • Provide analysis on how end-users are interacting with service content, products, and features	2012, ongoing	•Volume - 45 TB • Sources - JSON • Throughput -> 20K/sec • Latency (1 hour – for static/pre-defined reports /real-time for streaming data)	•Lambda architecture • Amazon AWS, S3 • Apache Kafka, Storm • Hadoop - CDH 5, HDFS(raw data), MapReduce), Cloudera Manager, Oozie, Zookeper • HBase (2 clusters: batch views, streaming data)	• Hadoop upgrade – CDH 4 to CDH 5 • Data integrity and data quality • Very high data throughput caused a challenge with data loss prevention (introduced Apache Kafka as a solution) • System performance for data discovery (introduced Redshift considering Spark) • Constraints - public cloud, multi-tenant
5 **Cloud-based Mobile App Development Platform** US private Internet Co. Funding > US$100M	• Provide visual environment for building custom mobile applications • Charge customers by usage • Analysis of platform feature usage by end-users and platform optimization	2013, 8 month	• Data Volume > 10 TB • Sources: JSON • Data Throughput > 10K/sec • Analytics - self-service, pre-defined reports, ad-hoc • Data Latency – 2 min	• Middleware: RabbitMQ, Amazon SQS, Celery • DB: Amazon Redshift, RDS, S3 • Jaspersoft • Elastic Beanstalk • Integration: Python • Aria Subscription Billing Platform	• schema extensibility • minimize TCO • achieve high data compression without significant performance degradation was quite challenging. • technology selection: performance benchmarks and price comparison of Redshift vs HPVertica vs Amazon RDS).
6 **Telecom E-tailing platform** Russian mobile phone retailer '14 Revenue: 108B rubles	• Build an OMNI-Channel platform to improve sales and operations • analyze all enterprise data from multiple sources for real-time recommendation and sales	End of 2013, (did only discovery)	• Analytics on 90+ TB (30+ TB structured, 60+ TB unstructured and semi-structured data) • Elasticity: through SDE principles	• Hadoop (HDFS, Hive, HBase) • Cassandra • HP Vertica/Teradata • Microstrategy/Tableau	• Data Volume for real-time analytics • Data Variety: data science over data in different formats from multiple data sources • Elasticity: private cloud, Hadoop as a service with auto-scale capabilities

Embedded Cases 7-10

Case #	Business goals	Start	Big data	Technologies	Challenges
7 **Social Relationship Marketing Platform** US private Internet Co. Funding > US$100M	• Build social relationship platform that allows enterprise brands and organizations to manage, monitor, and measure their social media programs • Build an Analytics module to analyze and measure results.	2013 ongoing (redesign 2009 system)	• > one billion social connections across 84 countries • 650 million pieces of social content per day • MySQL (~ 11 Tb) Cassandra (~ 6Tb), ETL (> 8Tb per day)	• Cassandra • MySQL • Elasticsearch • SaaS BI Platform - GoodData • Clover ETL, custom in Java, • PHP, Amazon S3,Amazon SQS • RabbitMQ	• Minimize data processing time (ETL) • Implement incremental ETL, processing and uploading only the latest data.
8 **Web Analytics & Marketing Optimization** US MNC IT consulting co. (Employees > 430,000)	• Optimization of all web, mobile, and social channels • Optimization of recomm-endations for each visitor • High return on online marketing investments	2014, Ongoing (Redesign 2006-2010 system)	• Data Volume > 1 PB • 5-10 GB per customer/day • Data sources – clickstream data, webserver logs	• Vanilla Apache Hadoop (HDFS,MapReduce,Oozie,Zookeeper) •Hadoop/HBase • Aster Data • Oracle •Java/Flex/JavaScript	• Hive performance for analytics queries. Difficult to support real-time scenario for ad-hoc queries. • Data consistency between two layers: raw data in Hadoop and aggregated data in relational DW • Complex data transformation jobs
9 Network Monitoring & Management Platform US OSS vendor Revenue > US$ 22M	•Build tool to monitor network availability, performance, events and configuration. • Integrate data storage and collection processes with one web-based user interface. •IT as a service	2014, Ongoing (Redesign 2006 system)	•collect data in large datacenters (each: gigabytes to terabytes) •real-time data analysis and monitoring (< 1 minute) • types of devices: hundreds	• MySQL • RRDtool • HBase • Elasticsearch	• High memory consumption of HBase when deployed in a single server mode
10 **Healthcare Insurance Operation Intelligence** US health plan provider Employees> 4,500 Revenue> US$10B	• Operation cost optimization for 3.4 million members • Track anomaly cases (e.g. control schedule 1 and 2 drugs, refill status control) • Collaboration tool between 65,000 providers.	2014, Phase 1: 8 months, ongoing	• Velocity: 10K+ events per second • Complex Event Processing - pattern detection, enrichment, projection, aggregation, join • High scalability, High-availability , fault-tolerance	• AWS VPC • Apache Mesos, Apache Marathon, Chronus • Cassandra • Apache Storm • ELK (Elasticsearch, Logstash, Kibana) • Netflix Exhibitor •Chef	• Technology selection constraints by HIPAA compliance: SQS(selected) vs Kafka • Chef Resource optimization: extending/fixing open source frameworks • 90% utilization ratio • Constraints: AWS, HIPAA

RESULTS

- Big Data System Development Framework
- Big Data system Design (BDD) method

BDD Framework

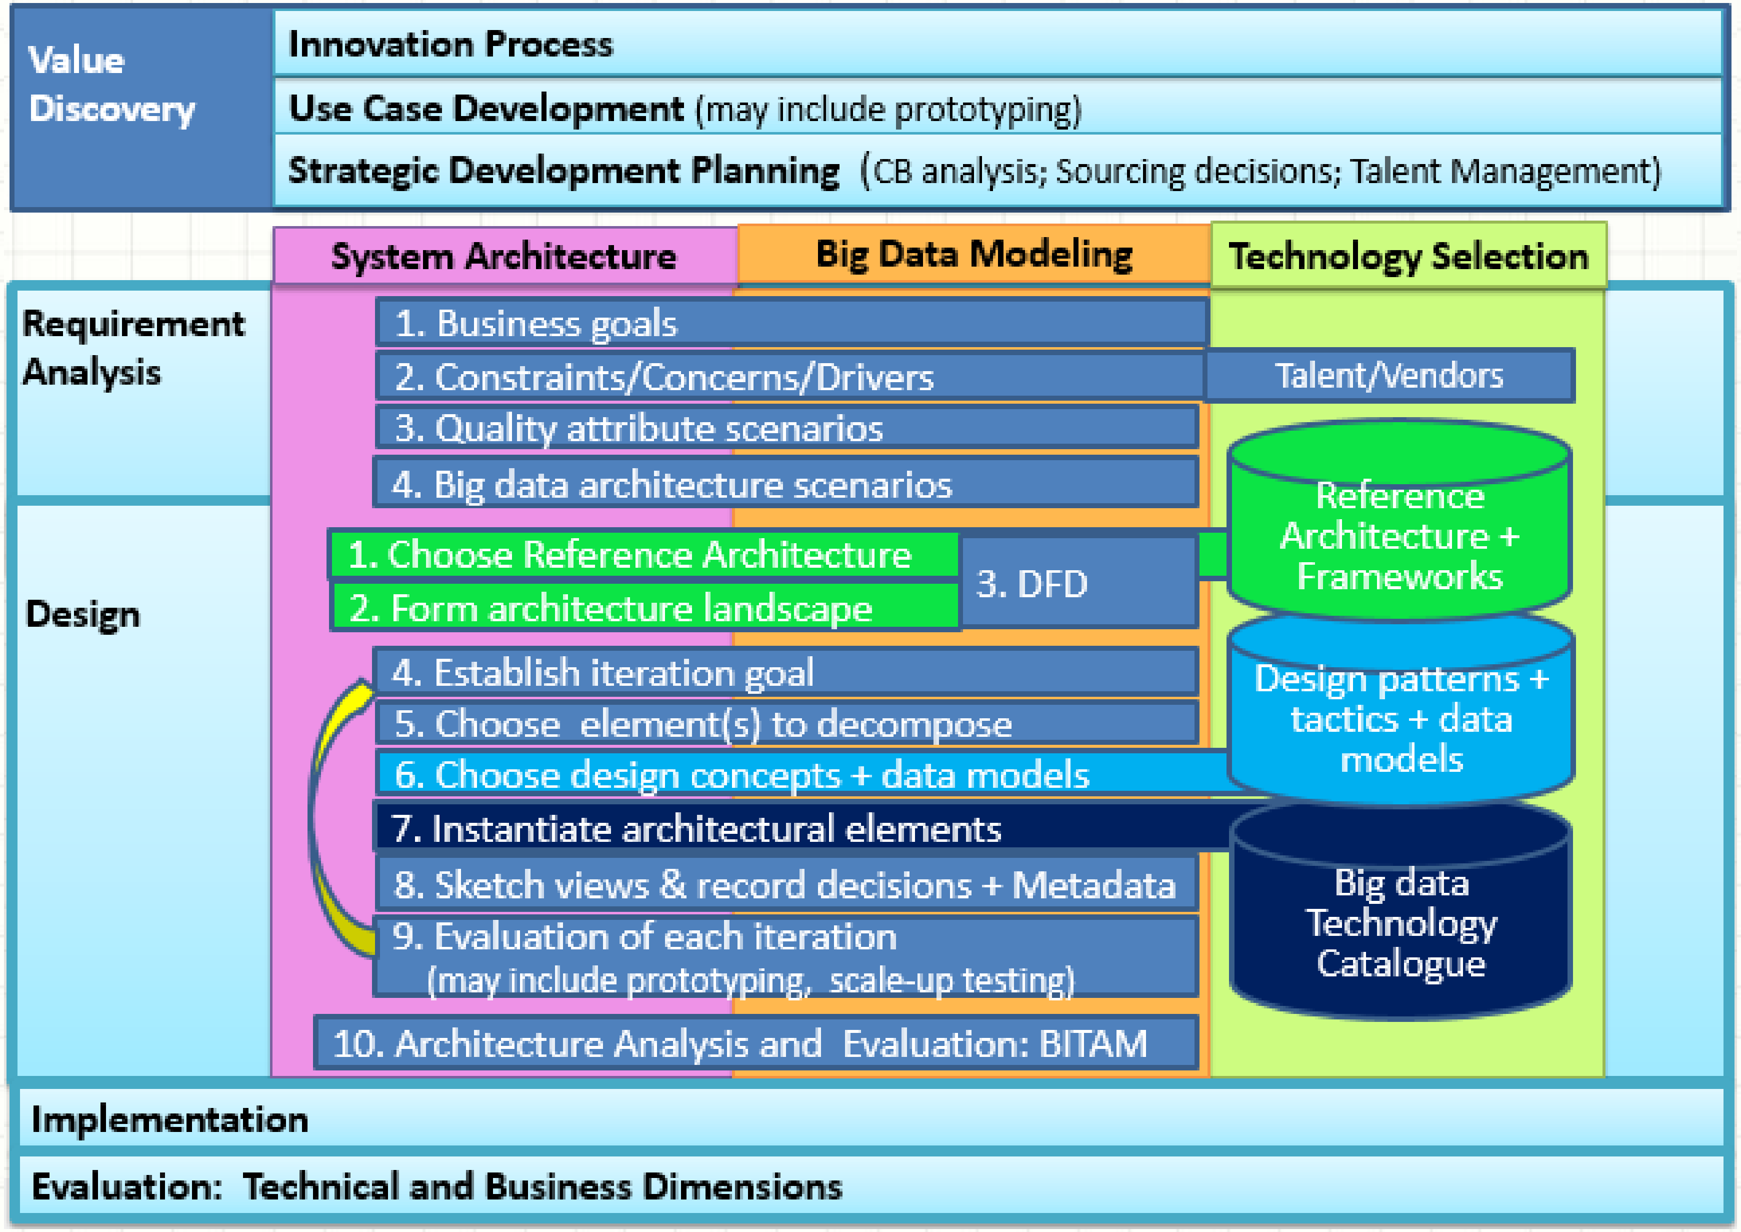

BDD Framework

1. **New Development Process**
   - Data-program independence undone
2. **"Futuring" big data scenario generation for innovation**
   - utilizing Eco-Arch method (Chen & Kazman, 2012).
3. **Architecture design integrated with new big data modeling techniques**:
   - Extended DFD (BD-DFD) , big data architecture template, transformation rules.
4. **Extended architecture design method**
   - ADD 2.0 (by CMU SEI) to ADD 3.0, then to BDD.
5. **Use of design concepts databases** (reference architecture, frameworks, platforms, architectural and deployment patterns, tactics, data models) and a **technology catalogue with quality attributes ratings**.
6. Adding architecture evaluation, BITAM (Business and IT Alignment Model), for risk analysis and ensuring **alignment with business goals** and **innovation desires**.
   - BITAM (Chen et.al. 2005, 2010) extended ATAM.

ECO-ARCH Method (Chen & Kazman, 2012)

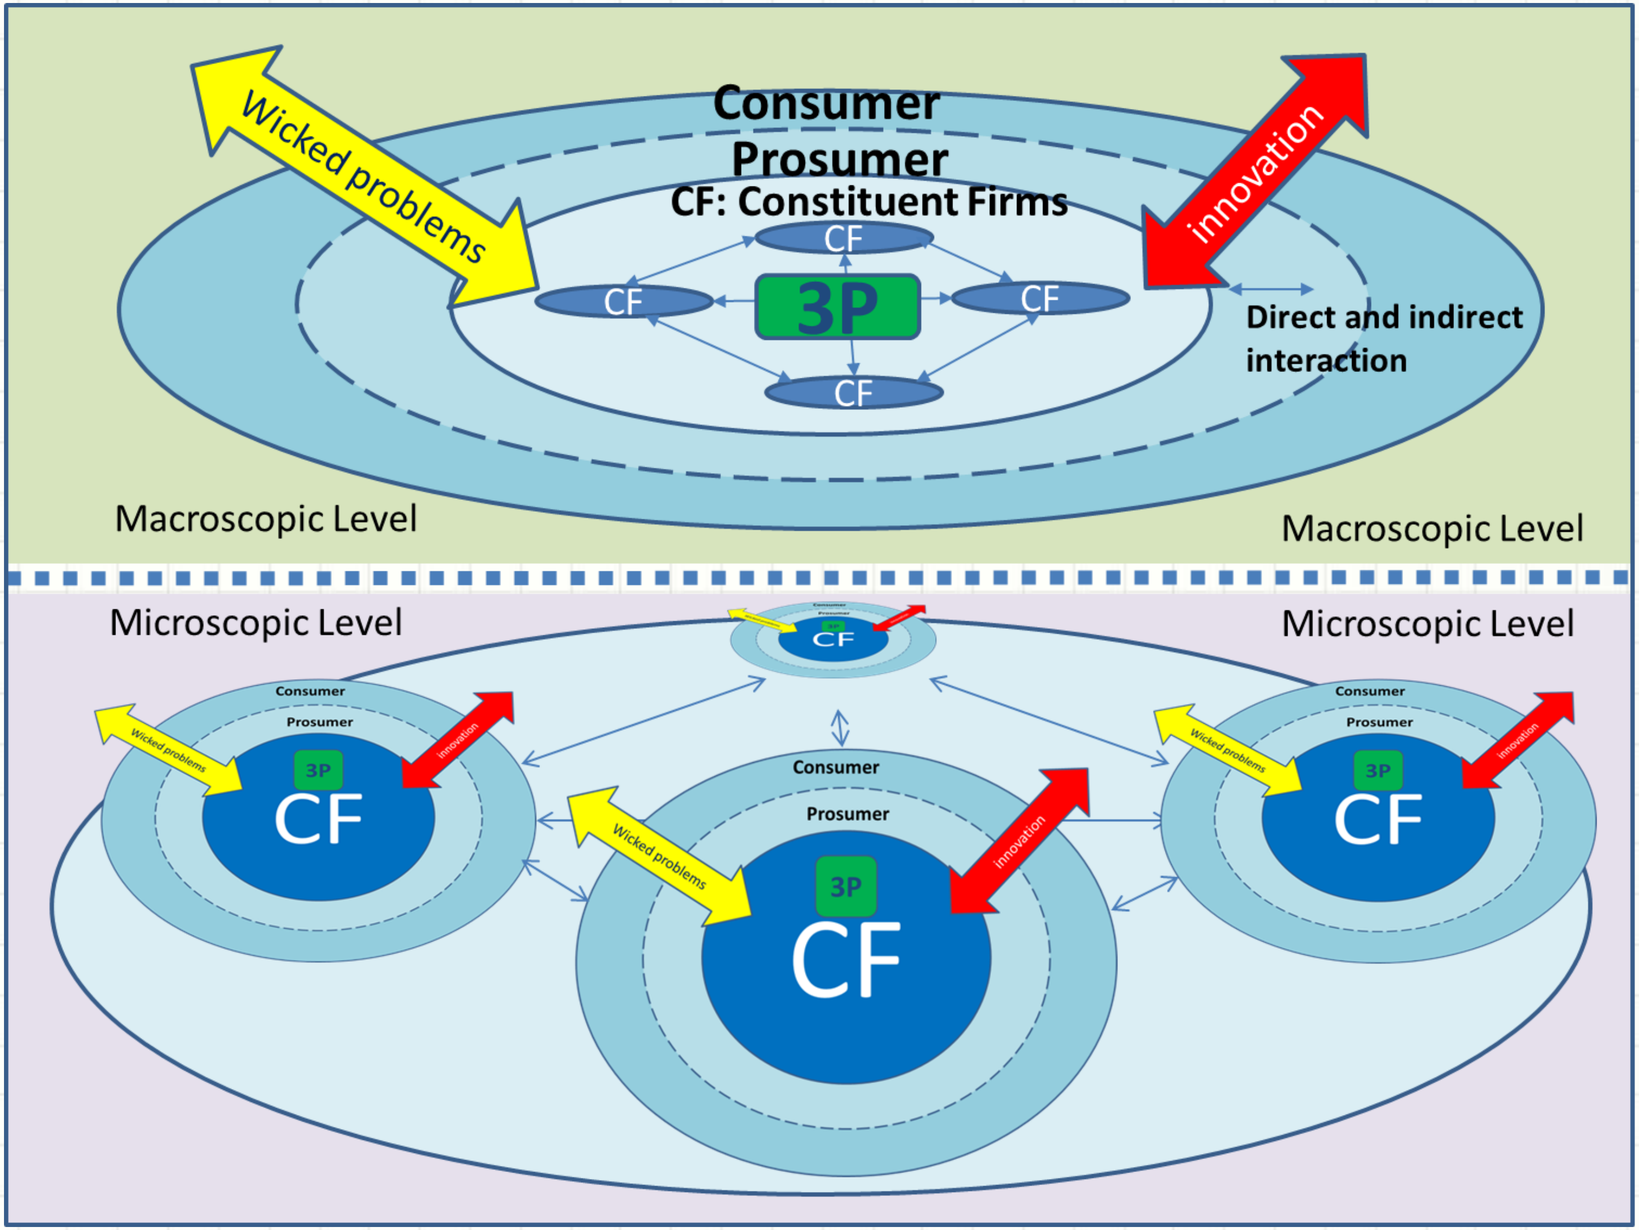

ECO-ARCH Method (Chen & Kazman, 2012)

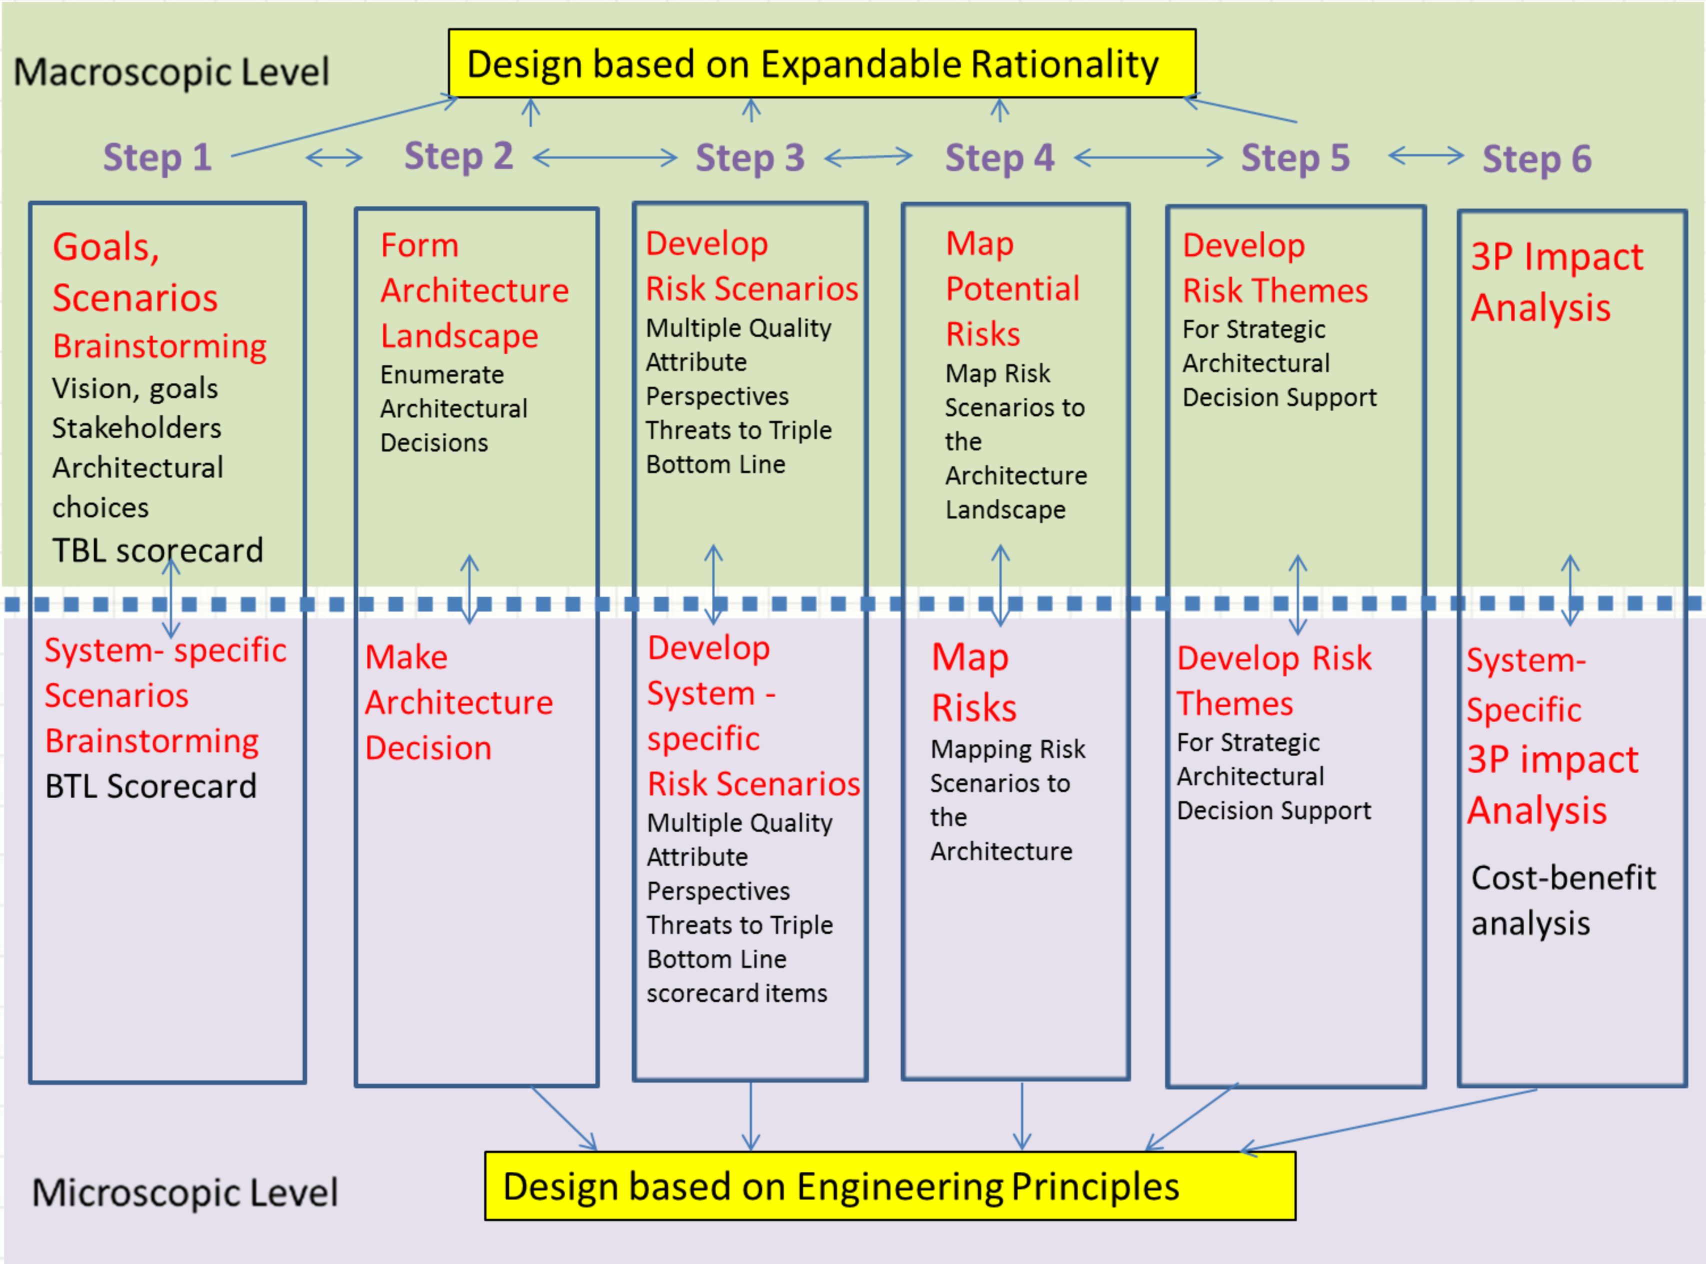

Big Data Architecture Design: Data Element Template

- A Scenario description includes the 6 elements: source, stimuli, environment, artifacts, response, response metrics.

1) **Data sources**: what are the data used in the scenario, where is it (are they) generated? Answer questions below for each source.
2) **Data source quality**: is this data trustworthy? How accurate does it represent the real world element it represents? Such as temperature taken?
3) **Data content format**: structured, semi-structured, unstructured? Specify subtypes.
4) **Data velocity**: what is the speed and frequency the data is generated/ingested?
5) **Data volume and Frequency**: What is the volume and frequency of data?
6) **Data Time To Live (TTL):** How long will the data live during processing?
7) **Data storage** : What is the volume and frequency of the data generated that need to be stored.
8) **Data Life**: how long should the data need to be kept in storage? (Historical storage/time series or legal requirements).
9) **Data Access type:** OLTP (transactional), OLAP (aggregates-based), OLCP (advanced analytics)
10) **Data queries/reports by who**: what questions are asked about the data by who? What reports (real time, minutes, days, monthly?)
11) **Access pattern:** read-heavy, write-heavy, or balanced?
12) **Data read/write frequency:** how often is the data read, written?
13) **Data response requirements:** how fast of the data queries needs to respond?
14) **Data consistency and availability requirements:** ACID or BASE (strong, medium, weak)?

Technology Catalogue: Topology

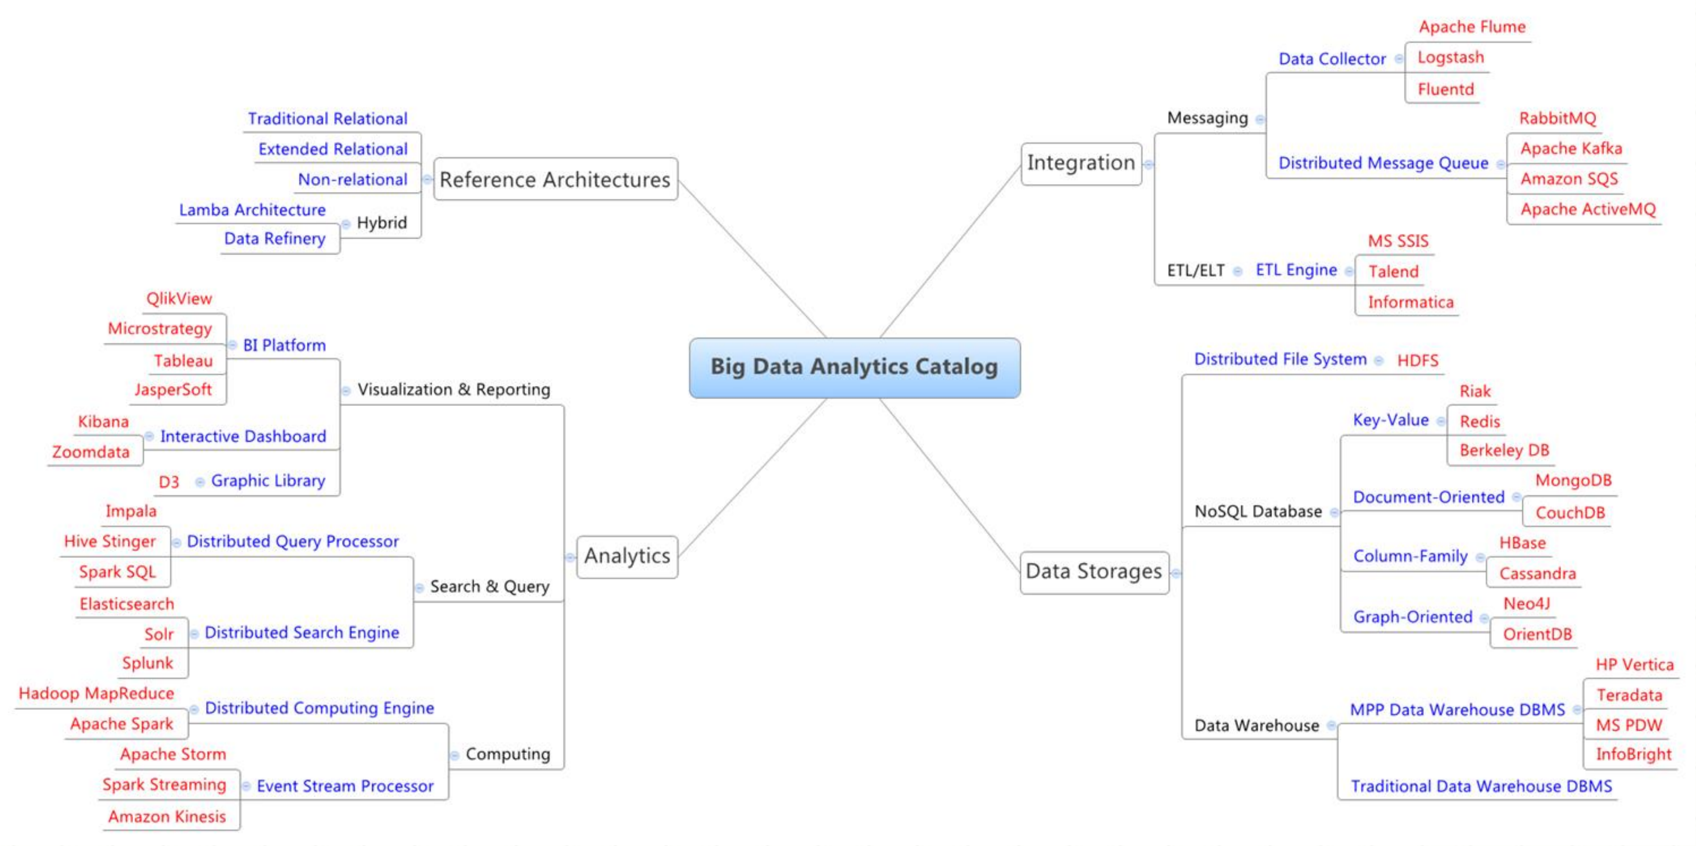

Ratings on Quality Attributes

Cassandra

Technology/Data Storage/NoSQL Database/Column-Family

Description: The Apache Cassandra database is the right choice when you need scalability and high availability without compromising performance. Linear scalability and proven fault-tolerance on commodity hardware or cloud infrastructure make it the perfect platform for mission-critical data. Cassandra's support for replicating across multiple datacenters is best-in-class, providing lower latency for your users and the peace of mind of knowing that you can survive regional outages.

cassandra

Consequences:

★★★ **Performance** – Cassandra is 30%-100% faster (avg) than HBase for both reads and writes due to efficient memory/caching, SSD support, online snapshots, locally-managed storage, effective compaction, etc. It is a winner of most performance benchmarks in its class.

★★★ **Reliability** – one of the most reliable and mature NoSQL databases today

★★ Ad-hoc analysis – Cassandra is not perfect for ad-hoc querying comparing to relational databases

★★★ **Real-time analysis** – fast access to data makes it perfect solution for real-time analysis backend

large community

Impala

Technology/Analytics/Search & Query/Distributed Query Processor

Description: Cloudera's open source massively parallel processing (MPP) SQL query engine for data stored in a computer cluster running Apache Hadoop.

Consequences:

★ **Processing capabilities** – supports the SQL-92 standard, but overall features are limited comparing to HiveQL

★★★ Performance – considered as one of fastest technologies at the moment, significantly faster than Hive

★★ **Compatibility** – supports storages: HDFS, HBase; formats: Parquet, Text, Avro, RCFile, SequenceFile, uses Hive metadata, can work through ODBC/JDBC

★ **Reliability** – designed for short queries, queries must be restarted if a node fails

2013. Besides Cloudera

Spark SQL

Technology/Analytics/Search & Query/Distributed Query Processor

Description: Based on Spark – an in-memory distributed computing engine (alternative to Hadoop MapReduce), Spark SQL allows running SQL and HiveQL queries over large datasets. Spark SQL is an ancestor of Shark.

Spark SQL

Consequences:

★★ **Processing capabilities** – based on SQL-like query language supporting most of HiveQL features including UDFs and SerDes

★★★ Performance – considered as one of fastest technologies at the moment, significantly faster than Hive

★ **Compatibility** – for now supports only HDFS, formats: Text and Parquet, can work through ODBC/JDBC

★★★ **Reliability** – supports long-running queries and mid-query faults recovery

★ **Maturity** – currently in alpha stage, although leverages Shark framework that counts its history since 2011

CouchDB

Technology/Data Storage/NoSQL Database/Document-Oriented

Description: CouchDB is a database that embraces the web by storing data with JSON documents; allowing accessing data via HTTP; indexing, combining, and transforming your documents with JavaScript. CouchDB works well with modern web and mobile apps, supports incremental replication and master-master setups with automatic conflict detection.

Apache CouchDB

Consequences:

★★ **Performance** – fast for direct ID lookups and map-reduce jobs, but that's it. Users reported performance issues.

★ **Reliability** – serious problems with reliability and availability were reported by users despite functionality like replication and automatic conflict resolution. Not yet suitable for highly-available or heavy-loaded solutions.

★ Ad-hoc analysis – CouchDB is generally queried by direct ID lookups and is not designed for ad-hoc despite secondary indexes and full-text search support

★★★ **Real-time analysis** – fast ID lookups and fast aggregation calculation using map-reduce

★★★ Ease of use – HTTP-based API makes it very easy to use and integrate with web applications. Administration of DB could be made using HTTP API too.

MongoDB

Technology/Data Storage/NoSQL Database/Document-Oriented

Description: MongoDB (from "humongous") is an open-source document database, and the leading NoSQL database. Written in C++, MongoDB features: JSON-style documents with dynamic schemas offer simplicity and power, indexes on any attribute, mirror across LANs and WANs for scale, scale horizontally without compromising functionality, flexible aggregation and data processing, etc.

mongoDB

Consequences:

★★ **Performance** – not as fast as simplest key-value storages, but features like auto-sharding, full index support, map-reduce makes it fast enough. Written in C++ rather than Java. With adding nodes throughput is not as efficient as with Cassandra or HBase.

★★ **Reliability** – durability is a known problem (being fixed though), issues with repairing databases, requires replication setup to implement reliability

★★ Ad-hoc analysis – full index support (index on any attribute) and map-reduce

★★★ **Real-time analysis** – one of the most common use-cases, supports schema design, indexing and sharding for real time analytics workloads

★★ Ease of use – requires documents model and JSON understanding

Apache Hive

Technology/Analytics/Search & Query/Distributed Query Processor

Description: The Apache Hive data warehouse software facilitates querying and managing large datasets residing in distributed storage.

HIVE

Consequences:

★★★ **Processing capabilities** – are based on HiveQL, a subset of SQL-92 which offers also extensions such as non-scalar data types, XML/JSON functions, UDFs, custom SerDes and other features

★★ Performance – even with Stinger initiative Hive is still slow comparing to other alternatives such as Impala or Spark SQL

★★ **Compatibility** – supports storages: HDFS, HBase, S3; formats: Text, Avro, RCFile, ORC, SequenceFile, can work through ODBC/JDBC

★★★ **Reliability** – supports long-running queries and mid-query faults recovery

★★★ **Maturity** – introduced in 2008, Apache Hive has been the de-facto SQL solution in Hadoop

BITAM
(Business-IT Alignment Model)

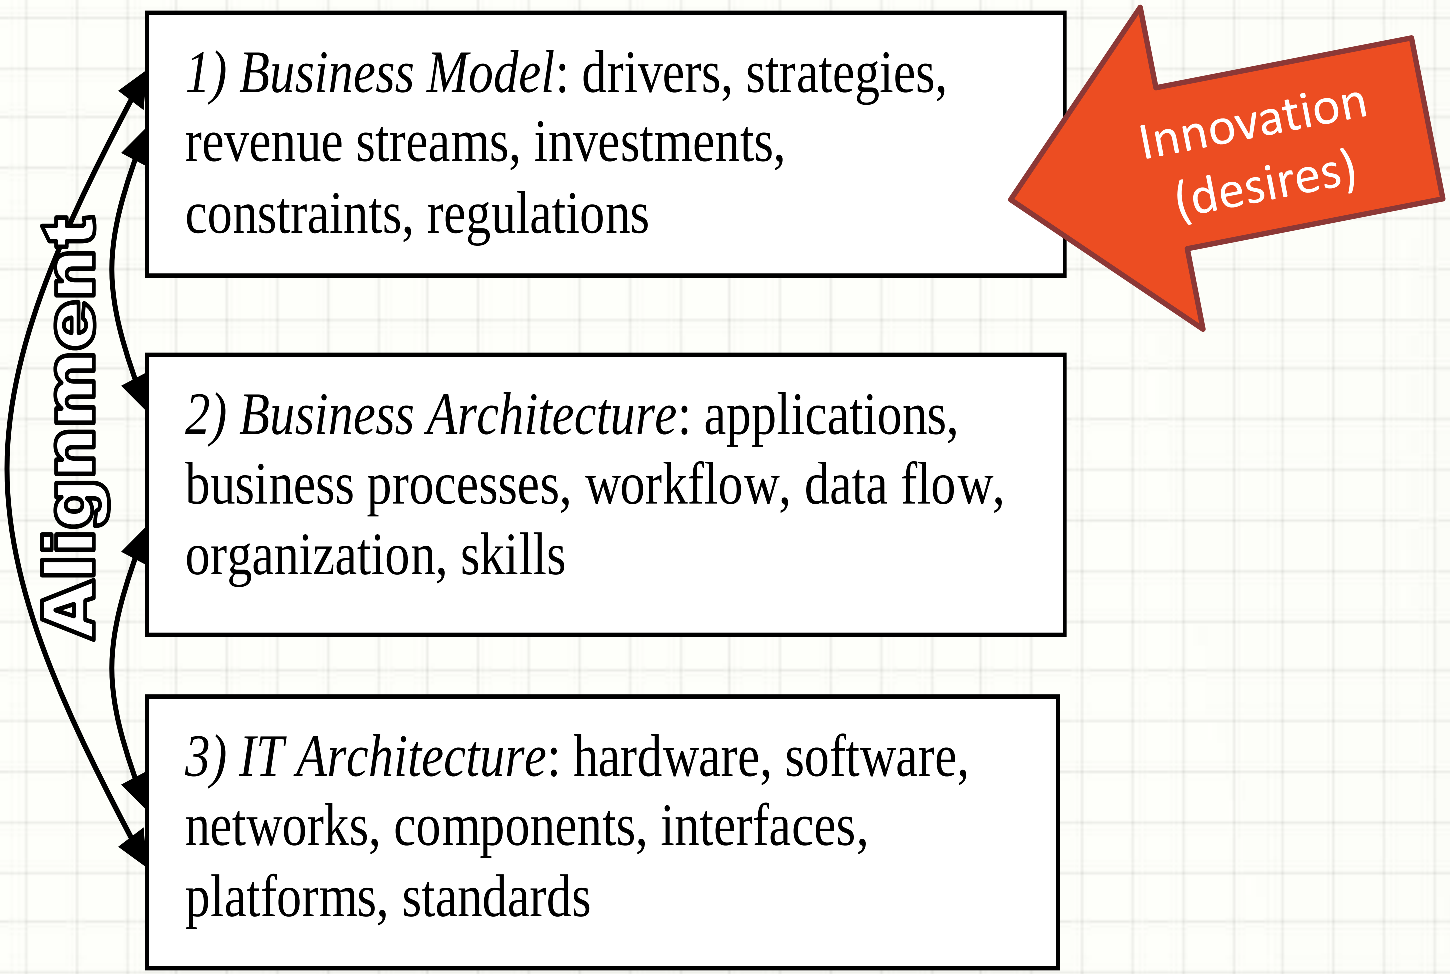

(Chen, Kazman, & Garg, 2005)

Work-in-Progress/Future Research

1. Prototyping vs. Architecture Analysis
2. Eco-Arch extension: More case studies
3. Decision support system (DSS) for knowledge-based big data technology selection
4. Automation of big data technology cataloguing
5. New big data design patterns for hybrid environment
6. Conceptual design for NOSQL data modeling
7. Metadata management for big data
8. Neo-Metropolis Model: BDaaS, etc.

Conclusions (1)

1. CPR approach balance rigor and relevance.

2. BDD framework describes a new process of big data system development, which is dramatically different from “small” data system development, reflecting the paradigm shifts required for big data system development.

3. Paradigm shifts and complexity in big data management underscore the importance of an architecture-centric design approach.

Conclusions (2)

4. BDD method is the first attempt to extend both architecture design methods and data modeling techniques for big data system design and integrate them in one method for design efficiency and effectiveness.
5. BDD method focuses on “futuring” for innovation.
6. BDD advances ADD 2.0 to ADD 3.0.
7. BDD method embodies best practice of complexity mitigation by utilizing quality attribute driven design strategies, reference architectures, technology catalogue (with ratings) and other design concepts databases for knowledge-based design and agile orchestration of technology.

Implications

- Disruptive Innovation Management
- Software Engineering Education

MAHALO & ALOHA!!!