LARGE, DISTRIBUTED COMPUTING INFRASTRUCTURES – OPPORTUNITIES & CHALLENGES

Dominique A. Heger Ph.D.
DHTechnologies, Data Nubes
Austin, TX, USA







www.dhtusa.com
www.datanubes.com

WORLD IS DEALING WITH MASSIVE DATA SETS

World-Wide Digital Data Volume (Source IDC 2012)

- 2000 -> ~800 Terabytes
- 2006 -> ~160 Exabytes
- 2012 -> ~2.7 Zettabytes
- 2020 -> ~35 Zettabytes
- 40% to 50% growth-rate per year

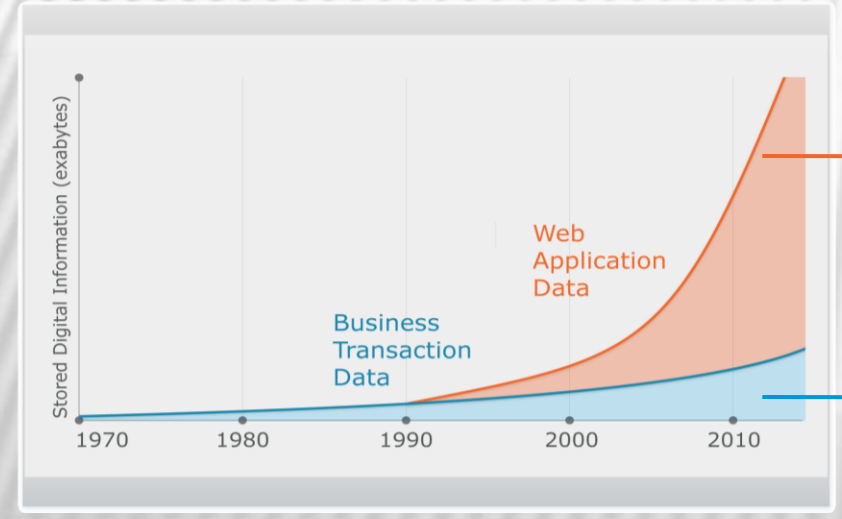
Name	Abbr.	Usage (Decimal)	Number of Bytes (Decimal)
1 megabyte	MB	10^6	1,000,000
1 gigabyte	GB	10^9	1,000,000,000
1 terabyte	TB	10^{12}	1,000,000,000,000
1 petabyte	PB	10^{15}	1,000,000,000,000,000
1 exabyte	EB	10^{18}	1,000,000,000,000,000,000
1 zettabyte	ZB	10^{21}	1,000,000,000,000,000,000,000
1 yottabyte	YB	10^{24}	1,000,000,000,000,000,000,000,000

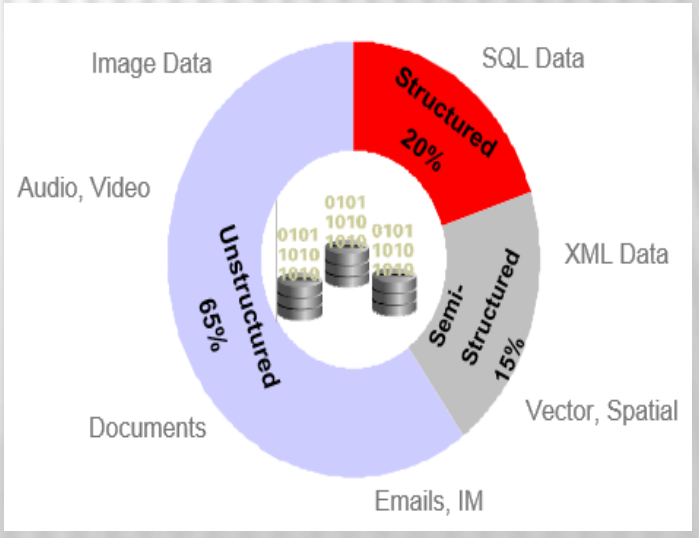
- Storing and managing 1PB of data may cost a company between $500K - $1M/year

Source: IDC 2012

STRUCTURED VERSUS UNSTRUCTURED DATA

- All *systems generated* data has structure!
- 70% to 80% of the digital data volume is labeled as unstructured
- Currently, most companies make all their business decisions solely based on their structured data pool …
- 56% of companies are overwhelmed by their data management requirements
- 60% of companies state that timely capturing & analysis of the data is not optimal
- ~2,700 EB of new information in 2012 with Internet as primary driver





Source: Gartner & IDC (2012)

DATA AS AN *ASSET* TODAY

Just as the Oil Industry Circa 1900 ….

After the refining process, one barrel of crude oil yielded more than 40% gasoline and only 3% kerosene, creating large quantities of waste gasoline for disposal.

"Book: The American Gas Station"

There are many Fortune 1000+ companies today with massive write-once & read-none data sets ….

BIG DATA – BIG CHALLENGES

- Big Data implies that the size of the data sets themselves become part of the problem
- Traditional techniques and tools to process the data sets are running out of steam
- A company does not have to be big to have Big Data problems
- Big Data Analytics & Predictive Analytics
- Data Management moves from batch to real time processing (Intel 2012)
- Cloud IT delivery model supports Big Data projects

HOW TO APPROACH A BIG DATA PROJECT

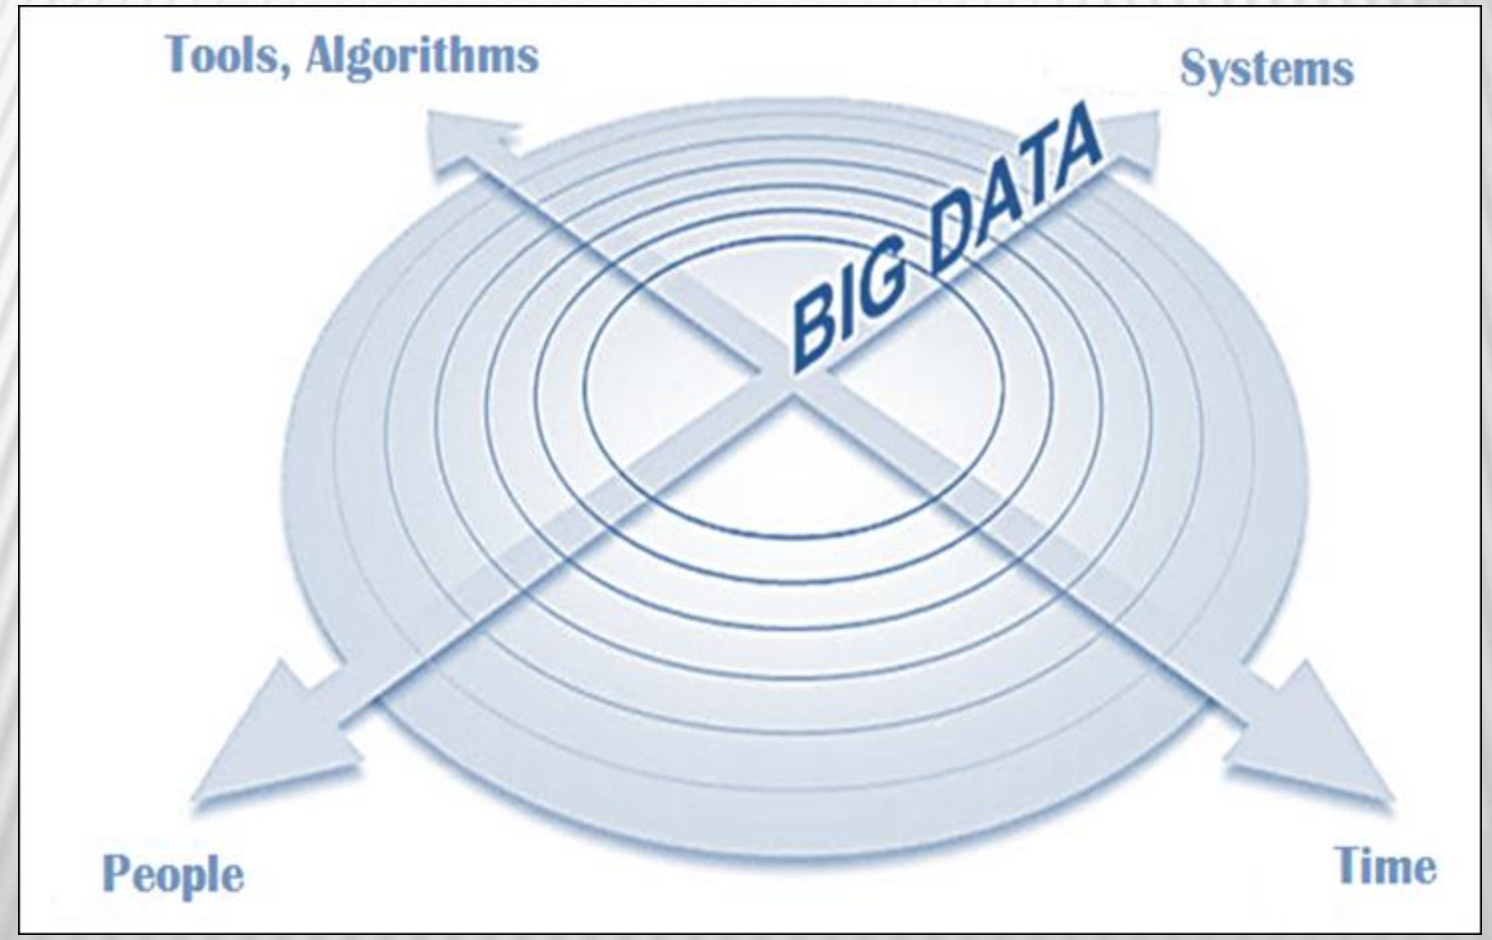
1. First, treat Big Data project as a business mandate and NOT as an IT challenge!
2. Define the top 3 most critical business questions that provide insight that will change the company's dynamic
3. Quantify the current time to answer (TTA) as well as the quality of the answer for these questions
4. Now the Big Data project goals and objectives can be defined as *"reduce the time to answer the following business questions from X number of hours down to Y number of minutes"*
5. Discuss the technology, people, tools, and project management opportunities required to realize these goals & objectives. Always do a POC!

PROBLEM DEFINITION

- Given the Big Data goals and a budget, provide a solution (supported by algorithms and an analysis framework) that guarantees that the quality of the answers meets the time and business objectives while data is accumulating over time.
  - This can only be achieved by implementing a scalable system infrastructure that fuses human intelligence with statistical and computational design principles (science and engineering)
  - Requires the 3 dimensions (systems, tools/algorithms, people) working together to improve the data analysis framework while meeting the goals and objectives
  1. Systems -> Design scalability into the IT solutions (Cloud)
  2. Algorithms -> Assess/Improve scalability, efficiency, and quality of the algorithms
  3. People -> Train & leverage human activity and intelligence (Data Scientist, CDO)

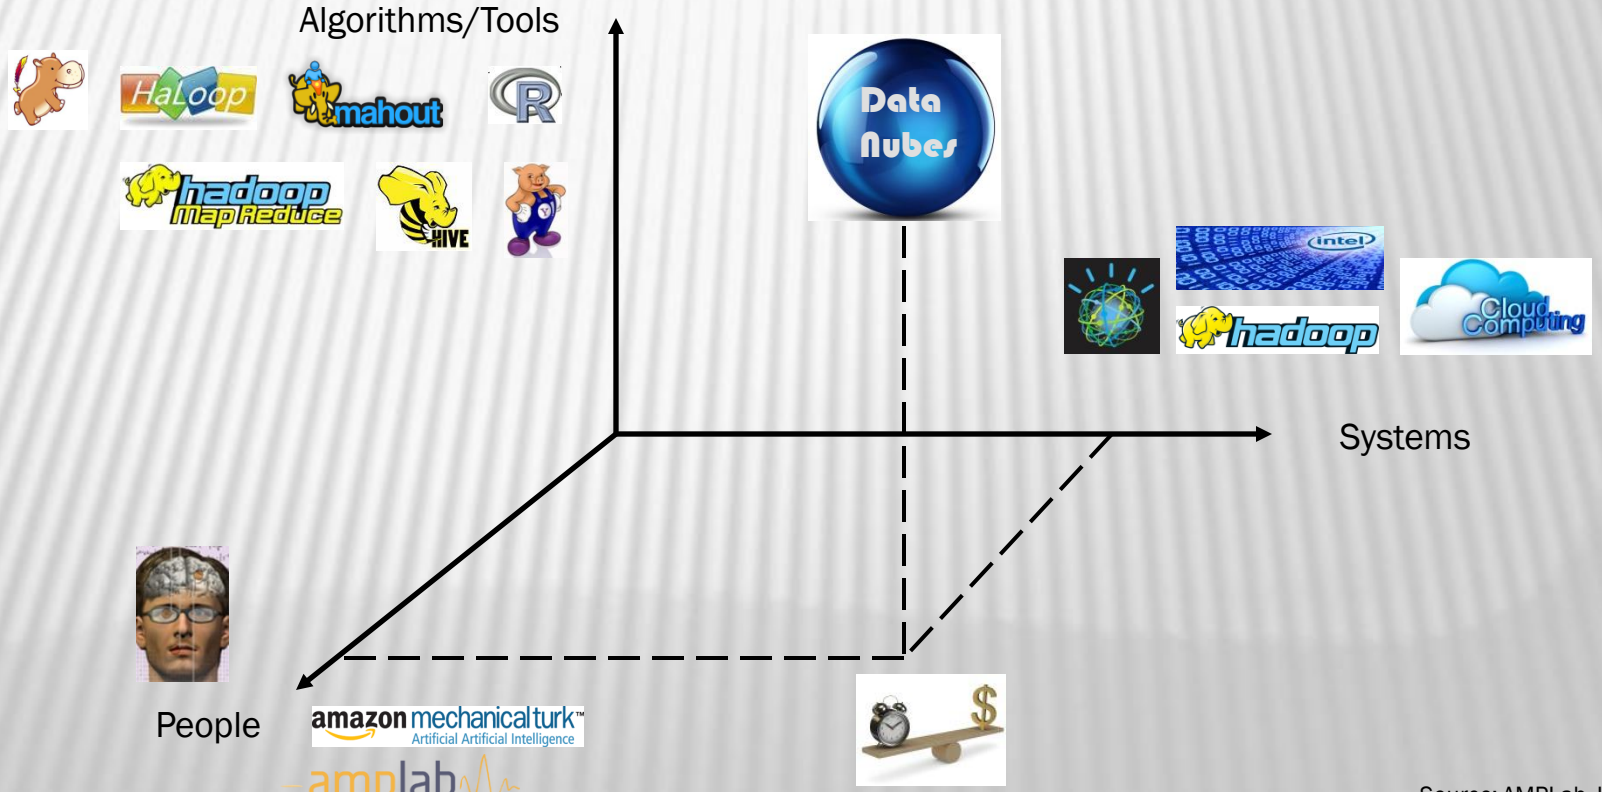
STATUS QUO

- Today's solutions reflect fixed points in the solution space



TARGET SOLUTION

- What is required are techniques to dynamically choose the best-possible operating points in the solution space
- Find answers at scale by tightly integrating algorithms, systems, and people



Source: AMPLab, UCB

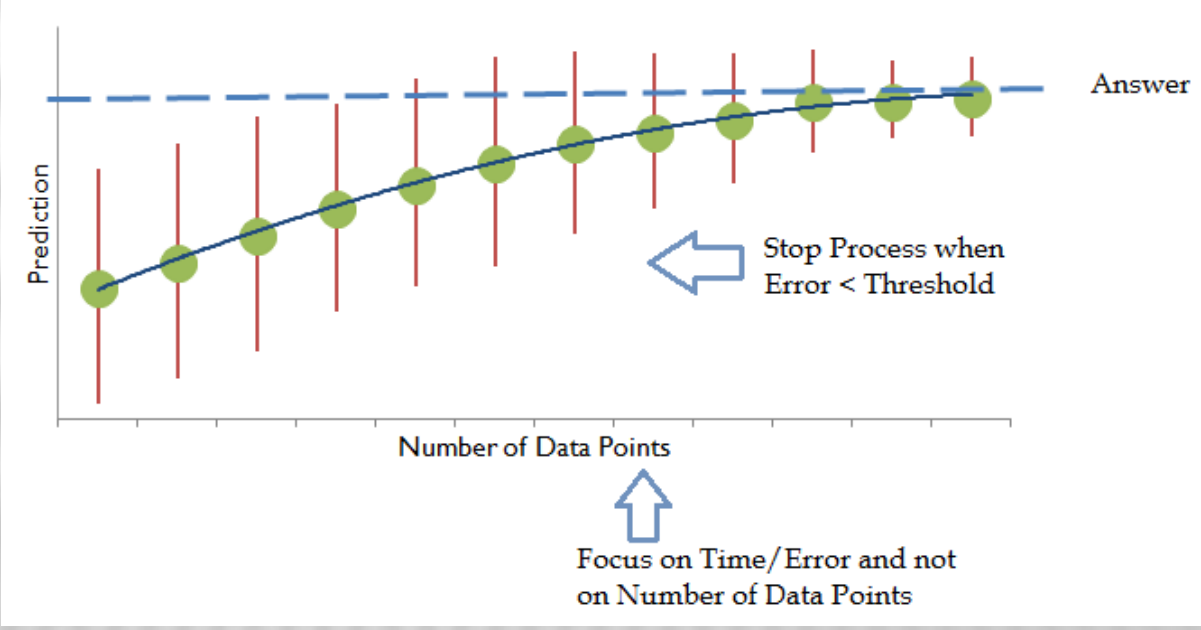
ALGORITHMS & TOOLS

- G1 -> The traditional ML toolsets for machine learning and statistical analysis such as SAS, SPSS, or the R language. They do allow for a deep analysis of smaller data sets (what is considered small is obviously debatable)
- G2 -> 2nd generation ML toolsets such as Mahout or RapidMiner that provide better scalability compared to G1, but may not support the vast range of ML algorithms as the G1 tools
- G3 -> 3d generation toolsets such as Twister, Spark, HaLoop, Hama, R over Hadoop, or GraphLab that provide deeper analysis cycles of big data sets

- Most current ML algorithms do not scale well to large data sets
- Sometimes unreasonable to process all data points and expect an answer within the specified time-frame (project goal)

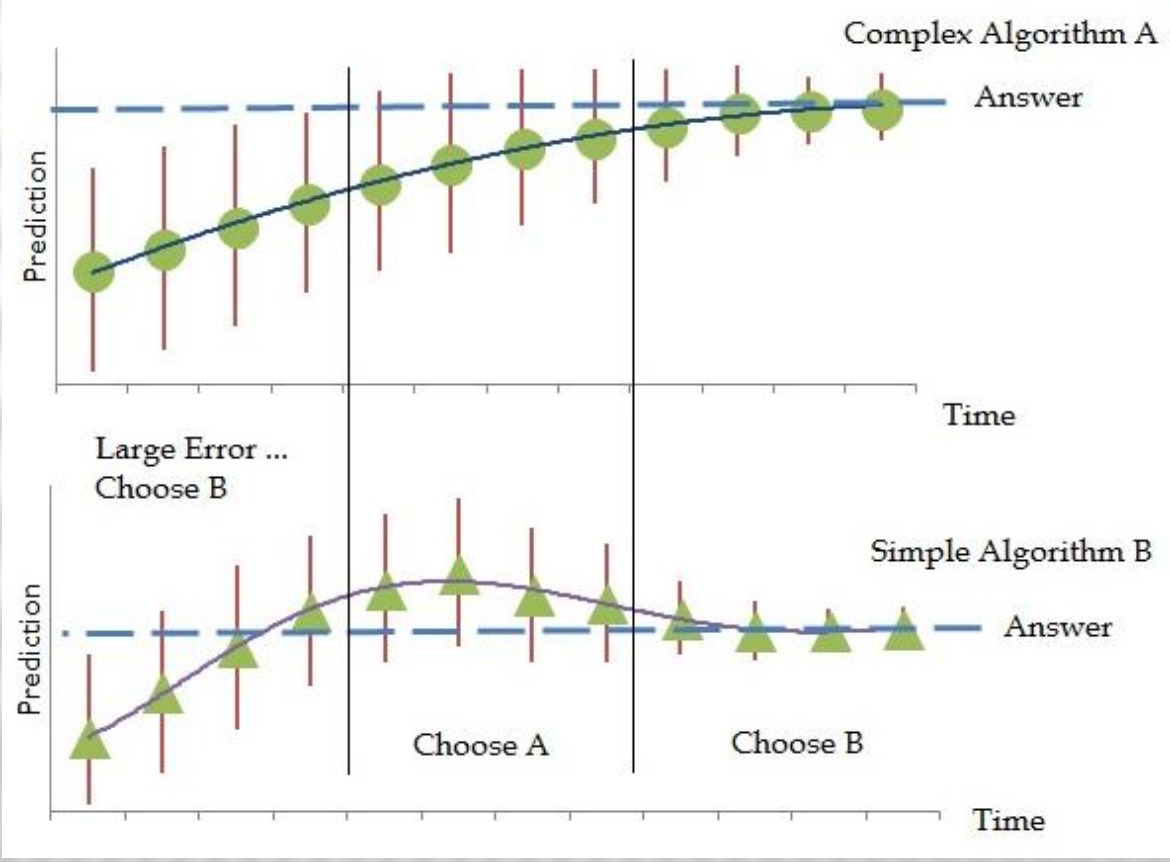
BIG DATA ANALYSIS - SUGGESTED APPROACH

- Given a question to be answered, a time-frame, and a budget, design and implement the system to obtain immediate answers while perpetually improving the quality of the results
- Calibrate the answers and provide error statistics
- Stop the process when the error < given threshold



FLEXIBILITY FOR A DYNAMIC SYSTEM

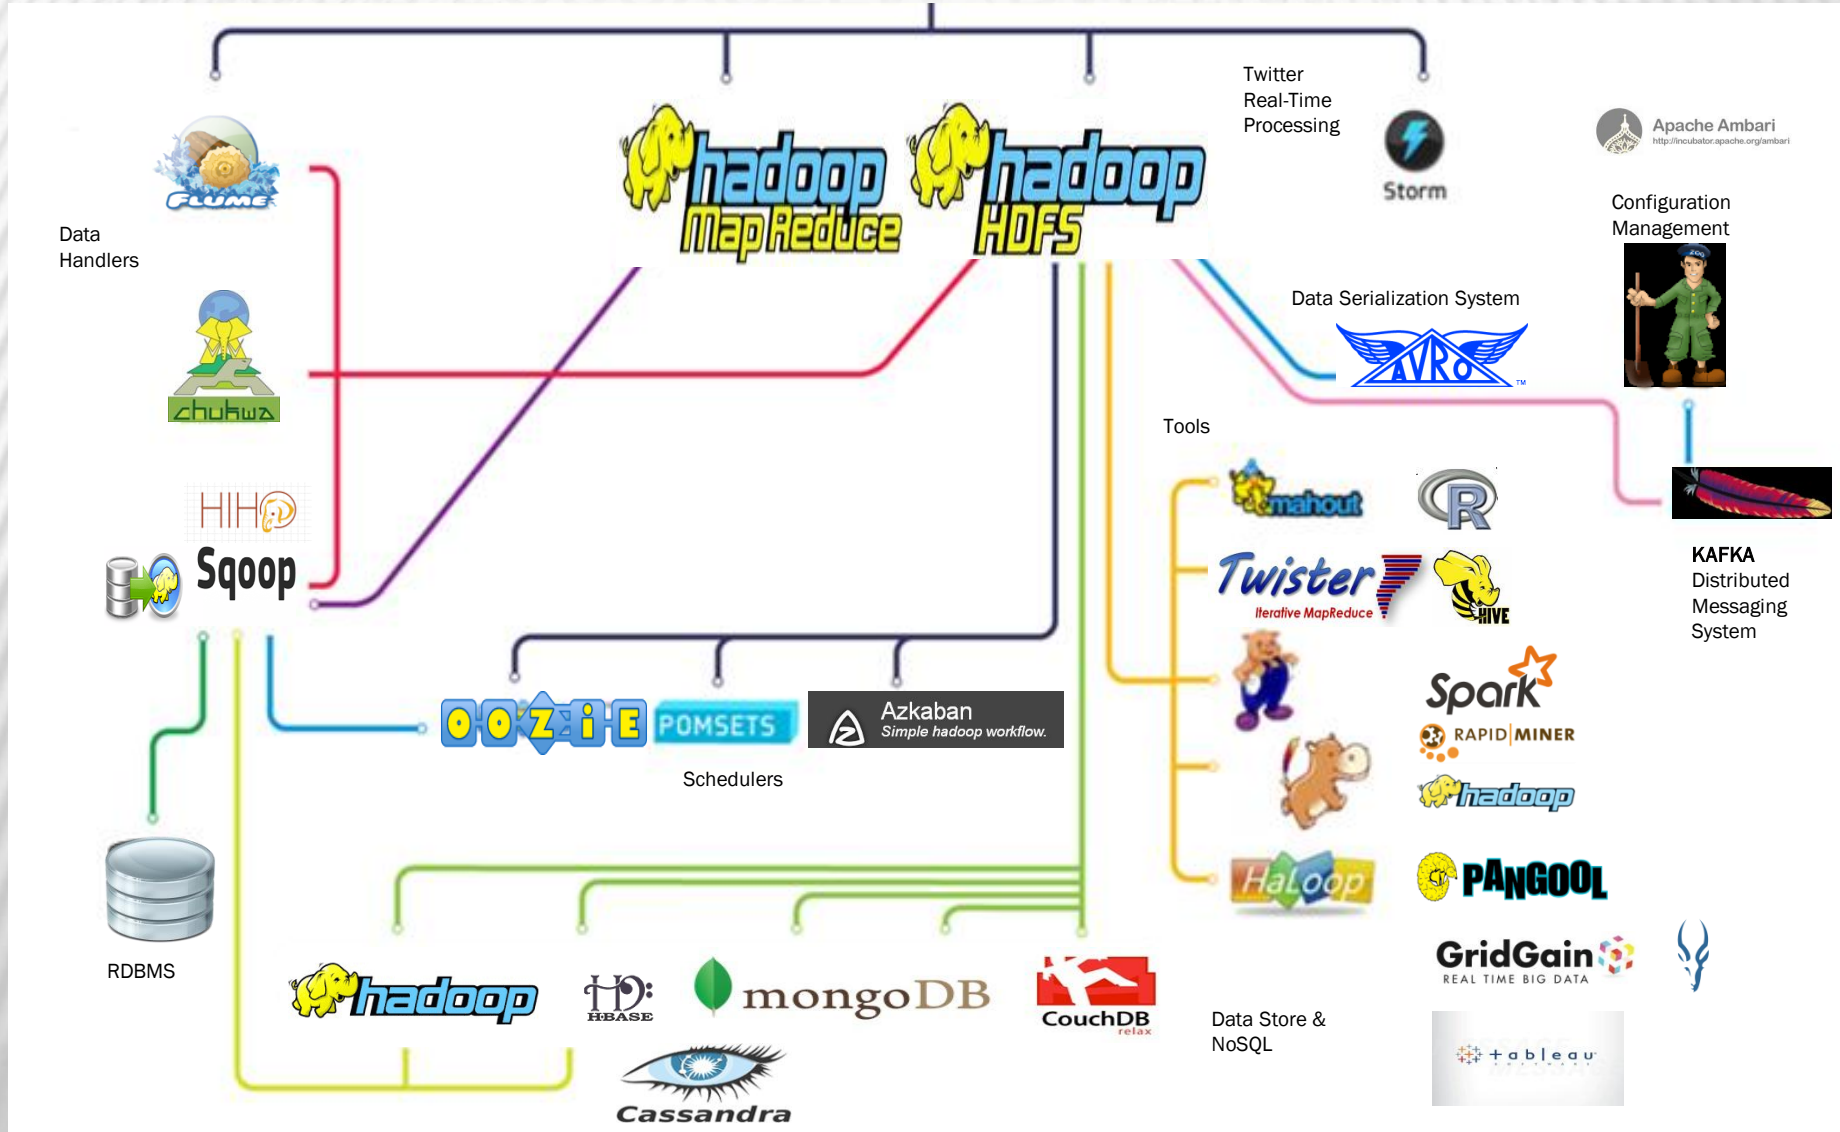
- Given a question to be answered, a time-frame, and a budget, automatically choose the best possible algorithm
- Example: Nearest Neighbor verses Learning Vector Quantization Classifier



SYSTEMS – HADOOP

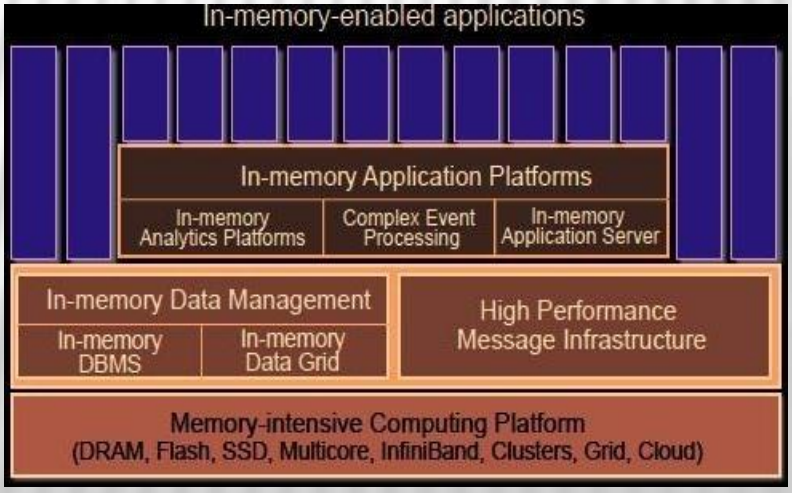
- Hadoop – Java based distributed computing framework that is designed to support applications that are implemented via the MapReduce programming model
- Hadoop Design Strategy – Move the actual computation to the data
- Old Strategy – Move the data to the computation (SAN)
- The traditional Hadoop performance focus is on aggregate data set (batch read) performance and NOT on any individual latency scenarios. The current focus though is more and more on Real Time processing!
- How to extract value from Big Data? ML!

HADOOP ECOSYSTEM (PARTIAL VIEW)



SYSTEMS – IN-MEMORY COMPUTING (IMC)

- IMC represents a set of technology components that allow storing data in system memory (DRAM) and/or Non-Volatile NAND flash memory rather than on traditional hard disks
- Core based systems and memory prices are coming down. Latency delta between NAND flash memory (ns) and HD's (ms) is significant while scaling the workload
- IMDG and IMCG products are available now and are solid
  - Case Study: 177M Tweets/day, 512 bytes each, data-set -> 2 weeks
  - Cluster (Intel Quad, 64GB Ram) with 1TB RAM -> ~$30,000 (20 parallel Quad nodes)
- In-Memory Hadoop available now (GridGain)
- Non-Volatile Phase-Change RAM (PCRAM) or Resistive RAM (RRAM) technologies may supersede NAND flash soon
- Establish an In-Memory Computing roadmap (Due-Diligence & Feasibility Study)



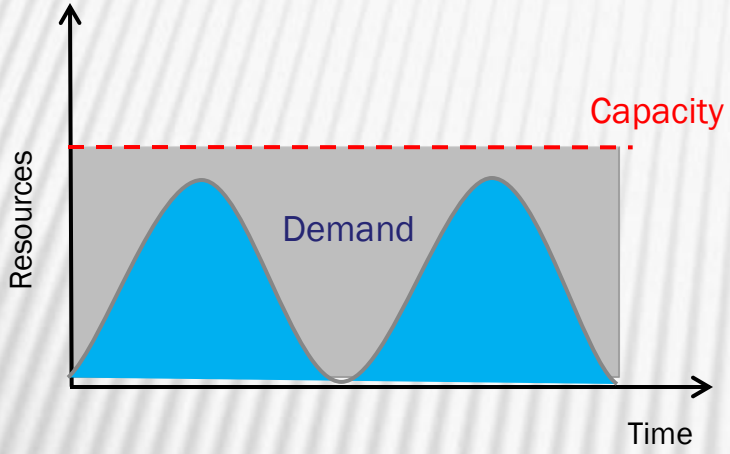
Source: Gartner, 2012

BIG DATA SYSTEMS FOCUS

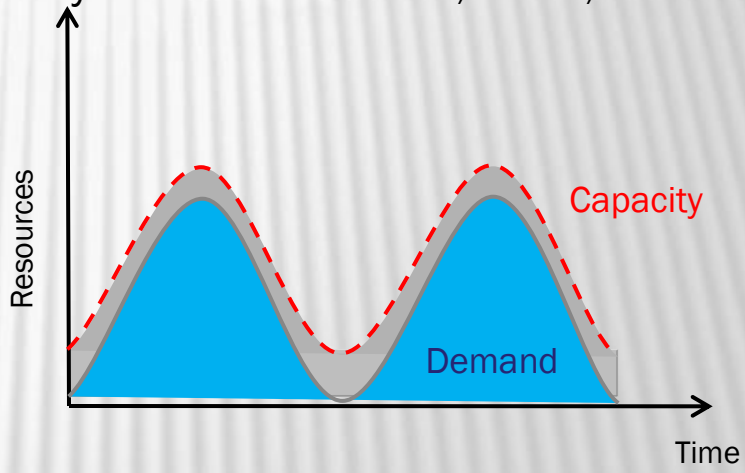
- Convert data center into a (Hadoop) processing unit
  - Commodity HW, Intel Core, Interconnect, Local Disks, No SAN
  - Support existing cluster computing applications (via Cassandra, Hive, Pig, or Hbase)
  - Support interactive and iterative data analysis (ML)
  - Support predictive, insightful query languages (Hive, Pig)
  - Support efficient and effective data movement among RDBMS and column oriented data stores (Sqoop)
  - Support distributed maintenance and monitoring of the entire IT infrastructure (Ganglia, Nagio, Chukwa, Ambari, White Elephant)
  - Scalability, robustness, performance, diversity, analytics, data visualization, and security aspects have to be designed into the solution
  - Make it all happen in a Cloud environment

BIG DATA & CLOUD COMPUTING

- Pay by use instead of provisioning for peak
- Risk of over-provisioning: underutilization
- Heavy penalty for under-provisioning (lost revenue, users)
- Big Data -> Analytics as a Service (AaaS), may be based on IaaS, PaaS, SaaS



Traditional Data Center



Cloud Based Data Center

Unused Resources

PEOPLE – BIG DATA

- Assure that people are an integrated (integral) part of the solution system
  - Leverage human activity
  - Leverage human intelligence
  - Leverage croudsourcing (online community)
  - Curate and clean dirty data (Data Cleaner, Data Wrangler)
  - Address imprecise questions
  - Design, validate, and improve algorithms

- After the business objectives are set, address any data at scale project by tightly integrating algorithms, systems, and people

PEOPLE – MASSIVE DEMAND & SMALL TALENT POOL

- US alone is facing an estimated shortage of approximately 190,000 scientist with deep analytical skills by 2018 (Source McKinsey, 2011)
- By 2018, US alone is facing an estimated shortage of approximately 1.5 million managers and analysts that have the know-how to leverage the results of big data studies to make effective business decisions (Source McKinsey, 2011)
- The Hadoop Ecosystem & Cloud Computing in general is powered by Linux.
- 91.4% of the top 500 supercomputers are Linux-based (Source TOP500)
  - A 2013 job report compiled by Dice showed that 93% of the contacted US companies (850 firms) are hiring Linux professionals this year.
  - The same study revealed that 90% of the firms stated that it is very difficult at the moment (2013) to even find Linux talent in the US. This number is up from 80% for the 2012 study.
  - According to Dice, the average salary increase for a Linux professional in the US is approximately 9% this year. At the same time, the average IT salary increase in the US is approximately 5%.

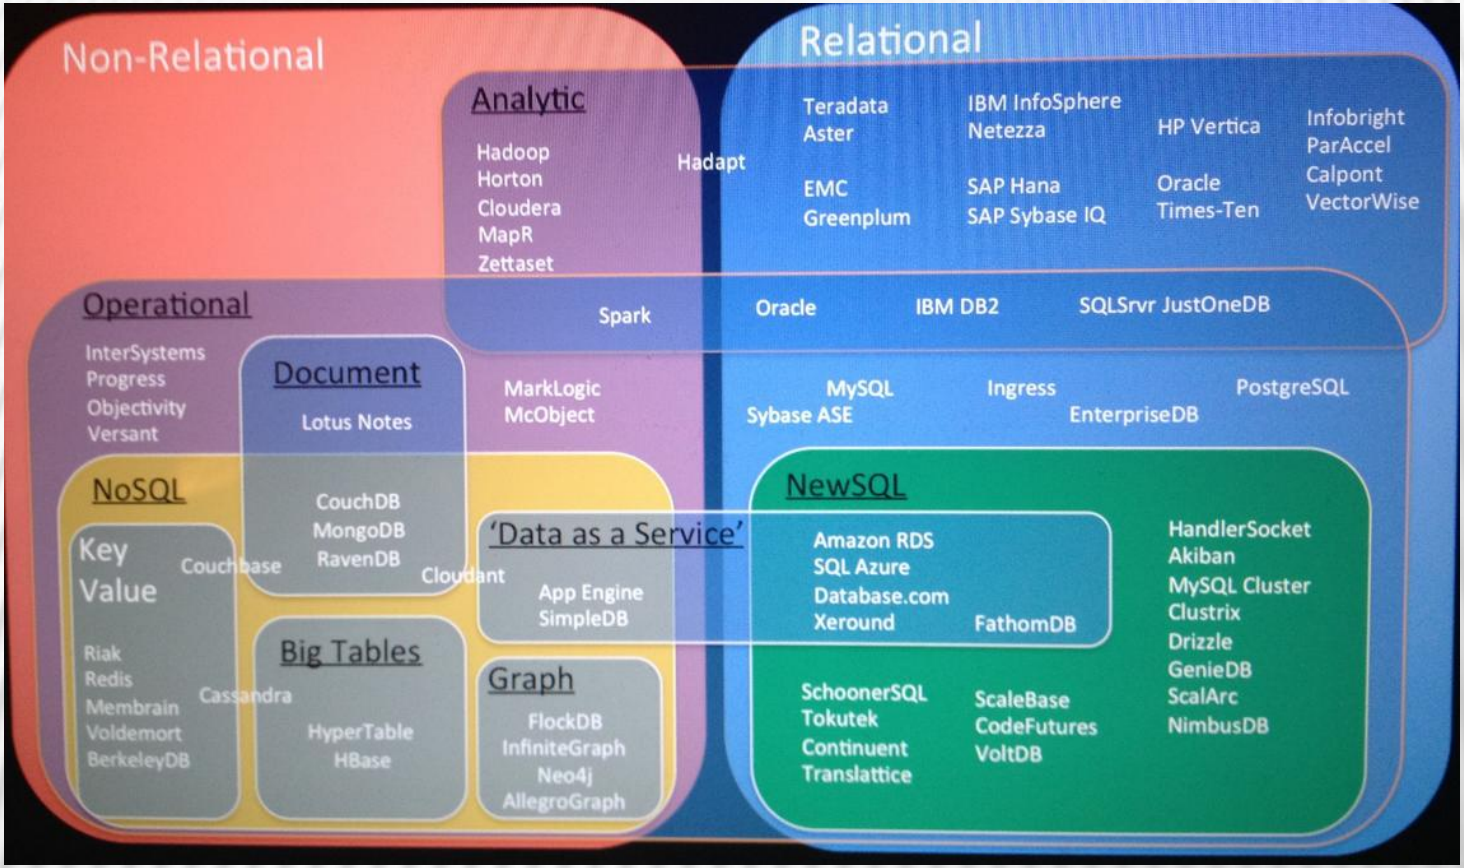
BIG DATA 2020

- Approach Big Data problems first as a business case (not an IT project) and strive for results that provide *the right quality at the right time* answers.
- Big Data projects require the fusion of algorithms/tools, systems, and people.
- In-Memory Computing (IMC), Complex Event Processing (CEP), as well as Quantum Computing reflect powerful options for Big Data projects
- Massive research opportunities across many domains exist, but the main objectives are:
  - Create a new generation of Big Data scientists (cross-disciplinary talent)
  - Machine Learning has to become an engineering discipline
  - Develop competency centers for the Big Data ecosystem
  - Develop centers of excellence for Linux & SW engineering
  - Leverage Cloud computing for Big Data, evaluate IMC/CEP now
  - Plan for IMC, CEP, Cloud, and the Big Data SW/HW infrastructure at the top company level and not the IT department
  - Leverage and be active in the Open Source community

THANKS MUCH!



SQL, NoSQL & NewSQL Framework



NewSQL is a class of modern relational database management systems that seek to provide the same scalable performance of NoSQL systems for online transaction processing (read-write) workloads while still maintaining the ACID (Atomicity, Consistency, Isolation, Durability) guarantees of a traditional database system

Source: Infochimps (2012)

Column verses Row Data Store – Data Operations

Sales Orders

	Doc Num	Doc Date	Sold-To	Order Value	Status	Sales Org	...
Row 1	95769214	2012-04-21	584	10.24	Closed	US	...
Row 2	95769215	2012-04-21	1215	124.35	Closed	US	...
Row 3	95769216	2012-05-02	584	47.11	Open	GER	...
Row 4	95769217	2012-05-07	454	21.20	Open	UK	...

Column Operations

```
SELECT SUM(Order Value)
FROM Sales Orders
WHERE Doc Date > 2012-01-31
```

Row Operations

```
SELECT *
FROM Sales Orders
WHERE Doc Number = ‚95769216'
```



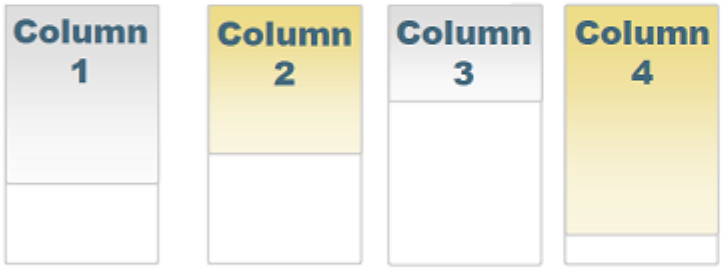
Column verses Row Data Store – Memory Storage

Customer	Country	Product	Amount
100	DE	1	100
100	DE	1	110
200	US	1	120
300	US	2	130

ROW-BASED Storage



COLUMN-BASED Storage



Row based storage in memory would look like this:

100-DE-1-100 100-DE-1-110 200-US-1-120 300-US-2-130

Column based storage in memory would look like this:

100-100-200-300 DE-DE-US-US 1-1-1-2 100-110-120-130