FACULTY DEVELOPMENT PROGRAMME
ON
END TO END PROCESSES IN DATA ANALYTICS
NOVEMBER, 02 - 06, 2020
ORGANIZED BY
DEPARTMENT OF INFORMATION TECHNOLOGY
SSN COLLEGE OF ENGINEERING, CHENNAI.

November 02, 2020

DATA ANALYTICS
KEY ENABLERS AND ASSOCIATED CHALLENGES

Venkatesh Vinayakarao
venkateshv@cmi.ac.in
http://vvtesh.co.in

Chennai Mathematical Institute

Today people are information-rich and time-poor
Marty Neumeier

Data Analytics

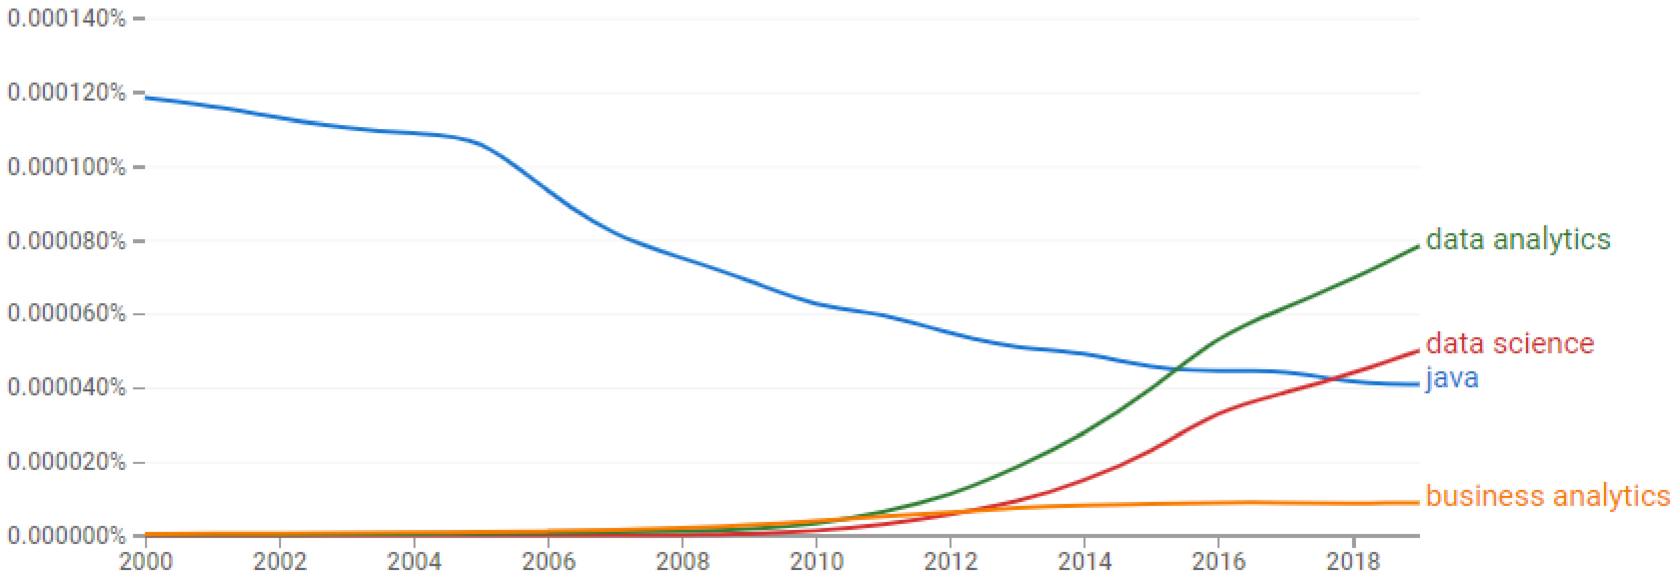

https://books.google.com/ngrams

Source:
https://books.google.com/ngrams/graph?content=java%2Cdata+science%2Cdata+analytics%2Cbusiness+analytics&year_start=2000&year_end=2019&corpus=26&smoothing=3&direct_url=t1%3B%2Cjava%3B%2Cc0%3B.t1%3B%2Cdata%20science%3B%2Cc0%3B.t1%3B%2Cdata%20analytics%3B%2Cc0%3B.t1%3B%2Cbusiness%20analytics%3B%2Cc0#t1%3B%2Cjava%3B%2Cc0%3B.t1%3B%2Cdata%20science%3B%2Cc0%3B.t1%3B%2Cdata%20analytics%3B%2Cc0%3B.t1%3B%2Cbusiness%20analytics%3B%2Cc0

Analytics

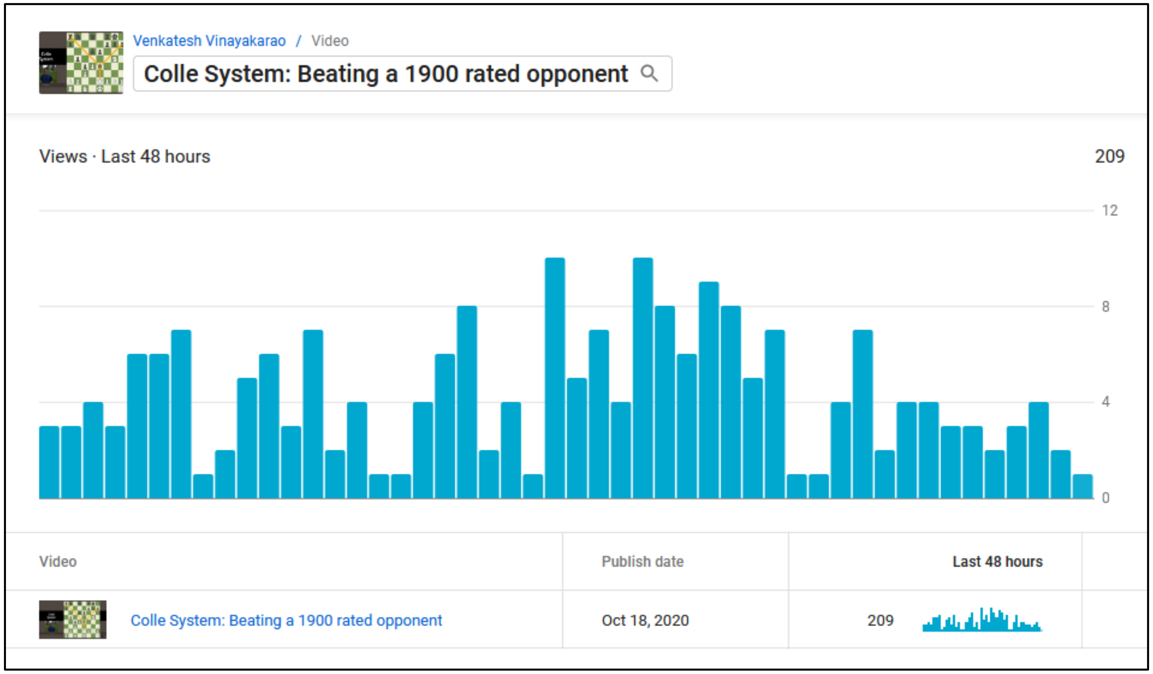

Helps me understand my audience and make better videos

A screenshot from youtube analytics on my chess channel

https://tinyurl.com/chess1900

The Impact of Big Data

Your train is on time thanks to **big data**
TNW - 31-Dec-2019
Thanks to thousands of sensors and **big data** analytics, train ... It's this data that keeps the Dutch rail network moving, and helps NS deliver a ...

The power of **data** in smart city developments
Independent Australia - 03-Jan-2020
Other fascinating **big data** developments that were presented included ... led to the production of the Australian **Cancer** Atlas — an interactive, ...

At HCA Healthcare, Real-Time **Data Saves Lives**
RTInsights (press release) (blog) - 01-Jun-2019
At HCA Healthcare, Real-Time **Data Saves Lives** ... “Our existing **data** infrastructure was designed for **large**-scale business intelligence and ...

Agenda

- Data Analytics
  - What are the key enablers?
    - A journey through the evolution of storage and processing infrastructure
  - What are the associated challenges?
    - Scale! Scale! Scale!

What Comes Next?

byte

kilobyte

megabyte

gigabyte

??

???

????

?????

Sizes

Name	Size
Byte	8 bits
Kilobyte	1024 bytes
Megabyte	1024 kilobytes
Gigabyte	1024 megabytes
Terabyte	1024 gigabytes
Petabyte	1024 terabytes
Exabyte	1024 petabytes
Zettabyte	1024 exabytes
Yottabyte	1024 zettabytes

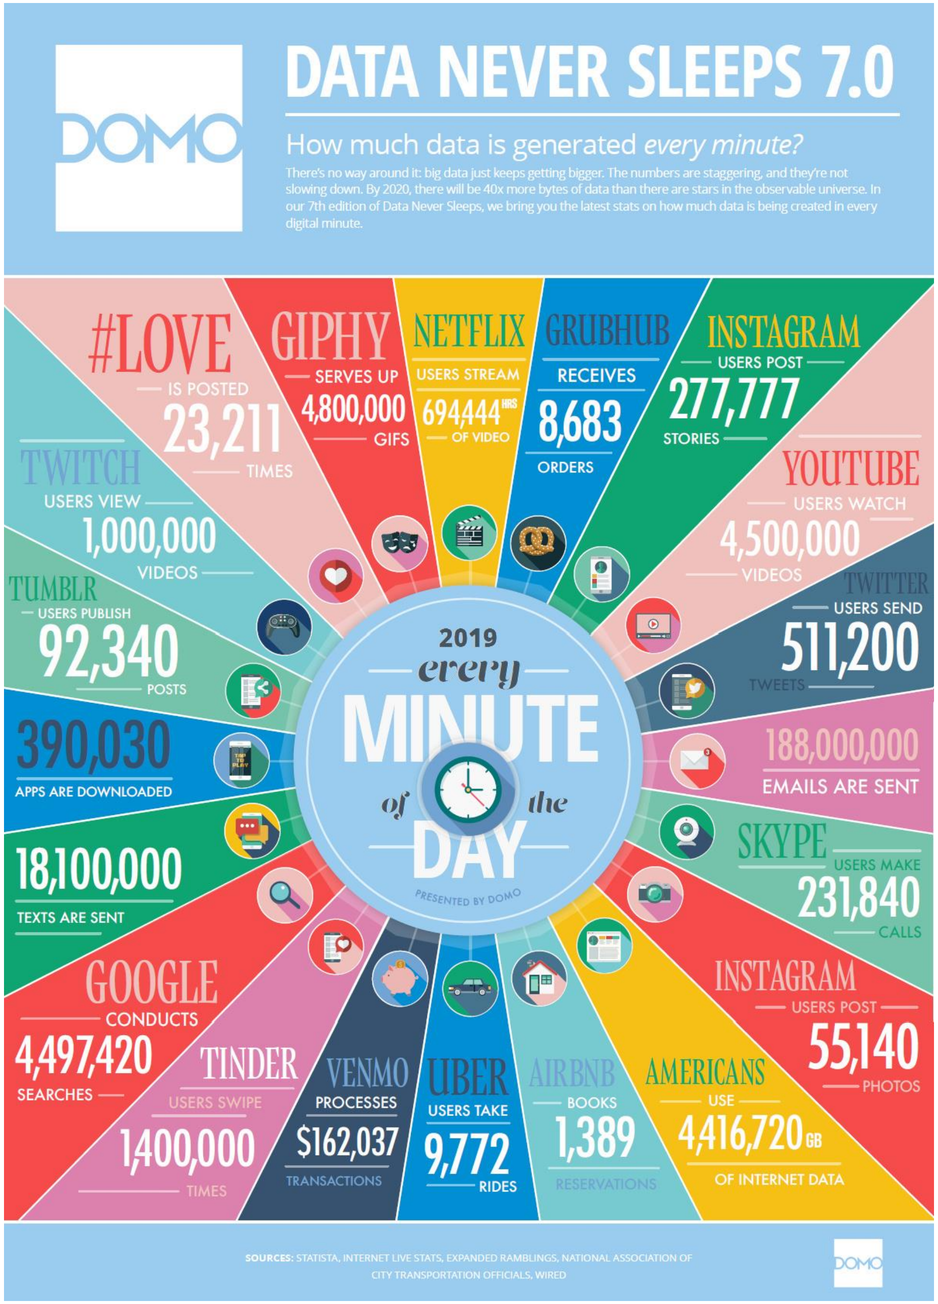

Source: https://www.visualcapitalist.com/big-data-keeps-getting-bigger/

And, It is Growing!

Data Growth

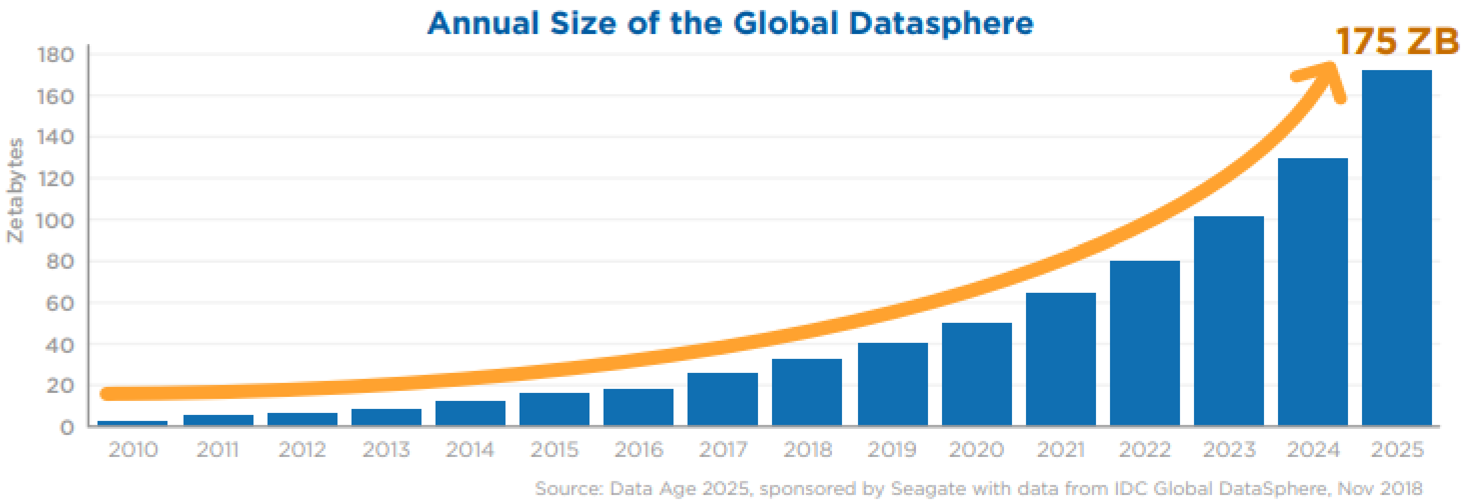

Mankind's quest to digitize the world!
33 ZB (2018) → 175 ZB (2025)
size of global datasphere*

*Source: https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf

Global datasphere is growing!

How have the computers evolved to capture, process and analyze these data?

Data Storage

Overview of hardware & software enablers

Storage Evolution

Seagate **4 TB** external HDD - $89.99

Exadrive DC100 - SSD – **100 TB**

Cloud storage

1950

IBM 350
3.75 MB
Leased out for $3200 per month!

1961

IBM 1301
19 MB
$115K for purchase.

1980

Seagate ST 506
World's first HDD – **5 MB** - $1500.

2000

Trek -**8MB** - $28.
World's first Thumbdrive

2019

(Secondary) Storage Technologies

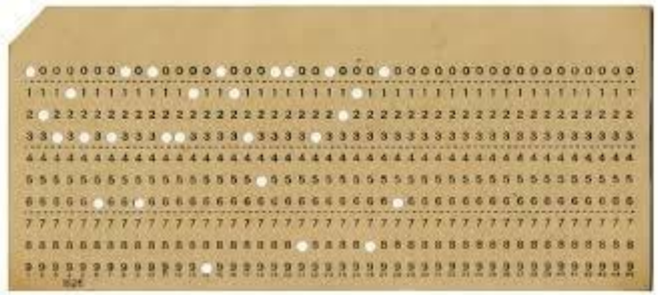

Disk Drive and Access Time

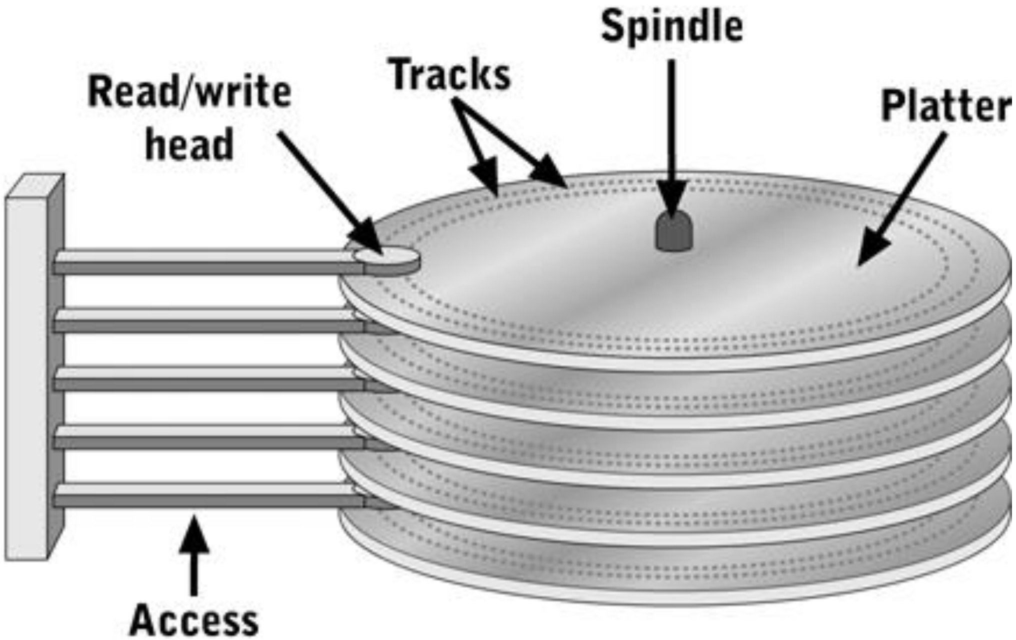

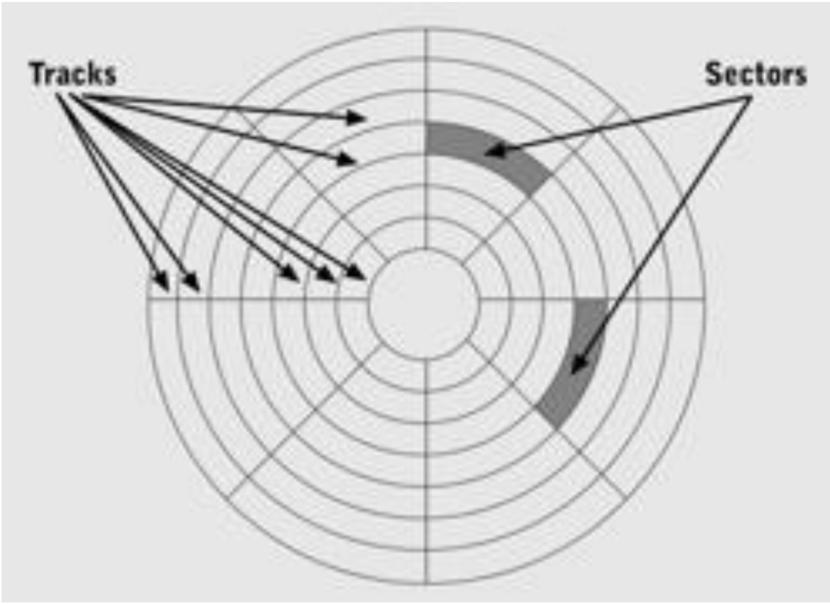

Source: Systems Architecture, Fifth Edition

Roll over image to zoom in

Seagate 500GB SATA Laptop Hard Disk

by Seagate

279 ratings | 493 answered questions

M.R.P.: ~~₹ 2,999.00~~

Price: ₹ 1,433.00 + ₹ 77.00 Delivery charge Details

You Save: ₹ 1,566.00 (52%)

Inclusive of all taxes

Pay on Delivery

10 Days Replacement

Amazon Delivered

1 Year Warranty

In stock.

Delivery by: **Jan 8 - 10** Details

Deliver to Venkatesh - Chennai 600014

Sold by KCM_STORE (3.9 out of 5 stars | 29 ratings).

New (20) from ₹ 1,510.00 + FREE Shipping

- 500 GB capacity
- 5400 RPM spin speed, 16 MB cache buffer
- Designed for durability and low-power consumption
- SATA 3GB interface with native command queuing
- Perpendicular recording technology for increased storage capacity
- Fast performance and whisper quiet acoustics

Average Access Time

- Head switching time is considered negligible (H)
  - Time taken to change from one read/write head to another
- Head seek time (S)
  - Time taken for the head to move from one track to another
- Rotational delay (R)
  - Rotate the disk to reach to the desired sector. Time taken for ½ a rotation (average).
- Read time (T)
  - time to spin an entire sector
- Average Access Time = H + S + R + T

*Sector is a minimum storage unit

Quiz

- If disk spins at 6000 RPM, compute the rotational delay.

Quiz

- If disk spins at 6000 RPM, compute the rotational delay.
  - One turn takes 1/6000 min or 1/100 sec = 10ms
  - ½ a turn takes 5ms.

Read Time

- If the drive spins at 6000RPM and the disk has 20 sectors per track, what is the read time?
  - Time for 1 full spin is $\frac{1}{6000}\text{min} = \frac{1}{100}\text{sec} = 10\text{ms}$
  - Time for 1/20 of a spin is $10\text{ms} \times \frac{1}{20} = 0.5\text{ms}$

Average Access Time

- Drive spins at 7200RPM and has average seek time of 8ms. The disk has 24 sectors per track. What is the average access time?

Head seek time	0.008 sec (Given)
Rotational delay	1/120 * (1/2) = 0.0042 sec
Read time	0.0084 (full spin) / 24 sectors = 0.00035 sec
Avg Seek Time	**= 0.008 + 0.0042 + 0.00035** **= 0.01255 sec or 12.55 ms**

Characteristics

Attribute	Description
Speed	Time to read/write
Volatility	Data persistence even when powered off
Access Method	Serial, Parallel
Portability	Internal, External
Capacity	Volume of data storage

How to store, retrieve and process data?

What were the (software) enablers?

ETL

- Extract, Transform and Load (ETL)
  - general procedure followed to address data variety
- Variety of data sources
  - Tabs, Sensors, Desktops, Bots, Multiple databases, Files,…
- Variety of data formats
  - Text, PDFs, XML, JSON, Images, Videos, …

Fast OLTP

- Online Transaction Processing (OLTP)
- Real-time/Near Real-time Performance. Finds application in:
  - Banking
  - Railway Reservations
  - Stock Market Trading
    - Handle transactions in milliseconds.
      - VoltDB, MemSQL, …

Data Warehouse

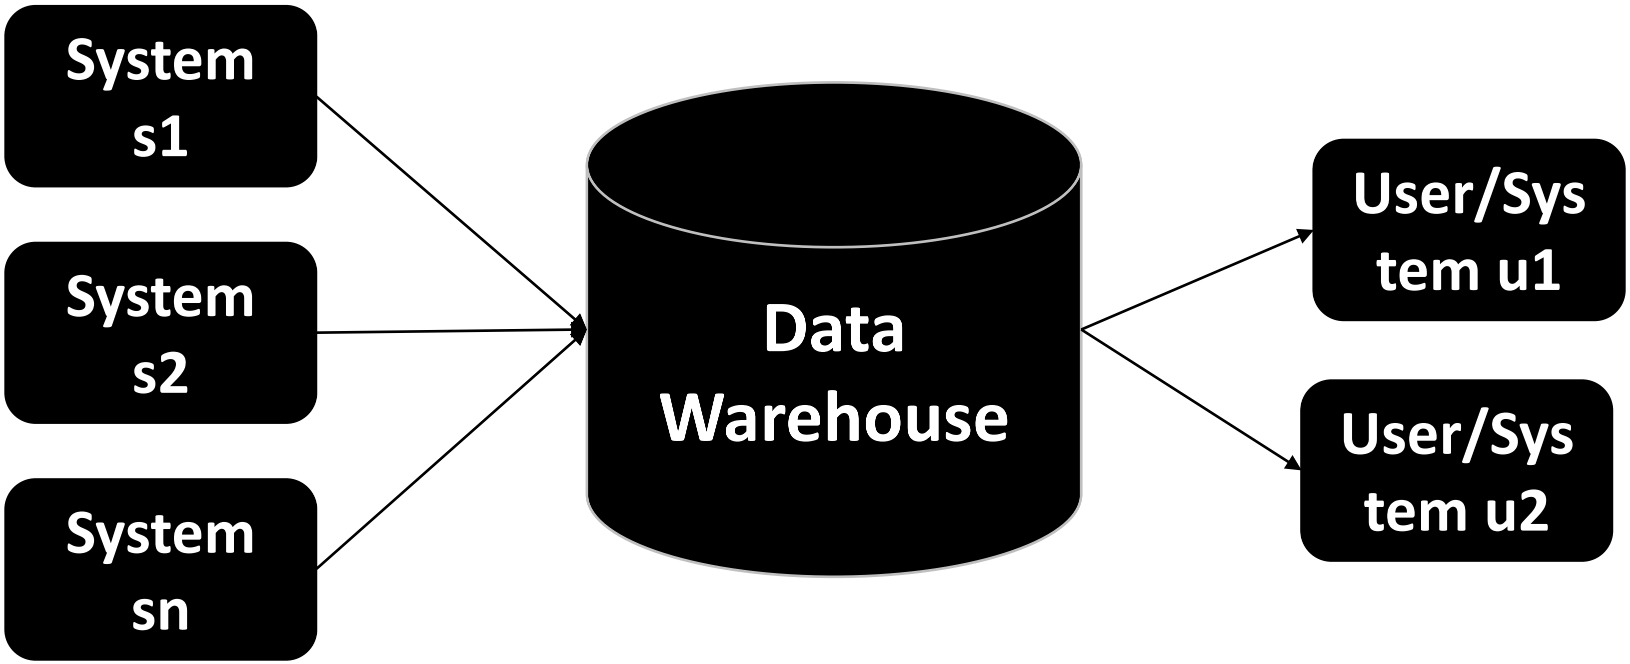

Data Lakes

- No schema definition.
- Store everything
  - often without or with very little pre-processing, /cleaning.
- Use ML, analytics to query, or gather insights.

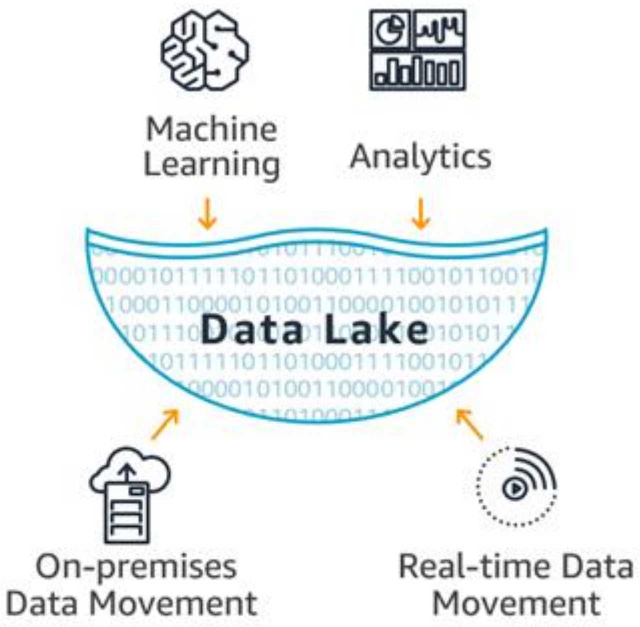

[Source: https://aws.amazon.com/big-data/datalakes-and-analytics/what-is-a-data-lake/](https://aws.amazon.com/big-data/datalakes-and-analytics/what-is-a-data-lake/)

Microsoft's Azure Data Lake

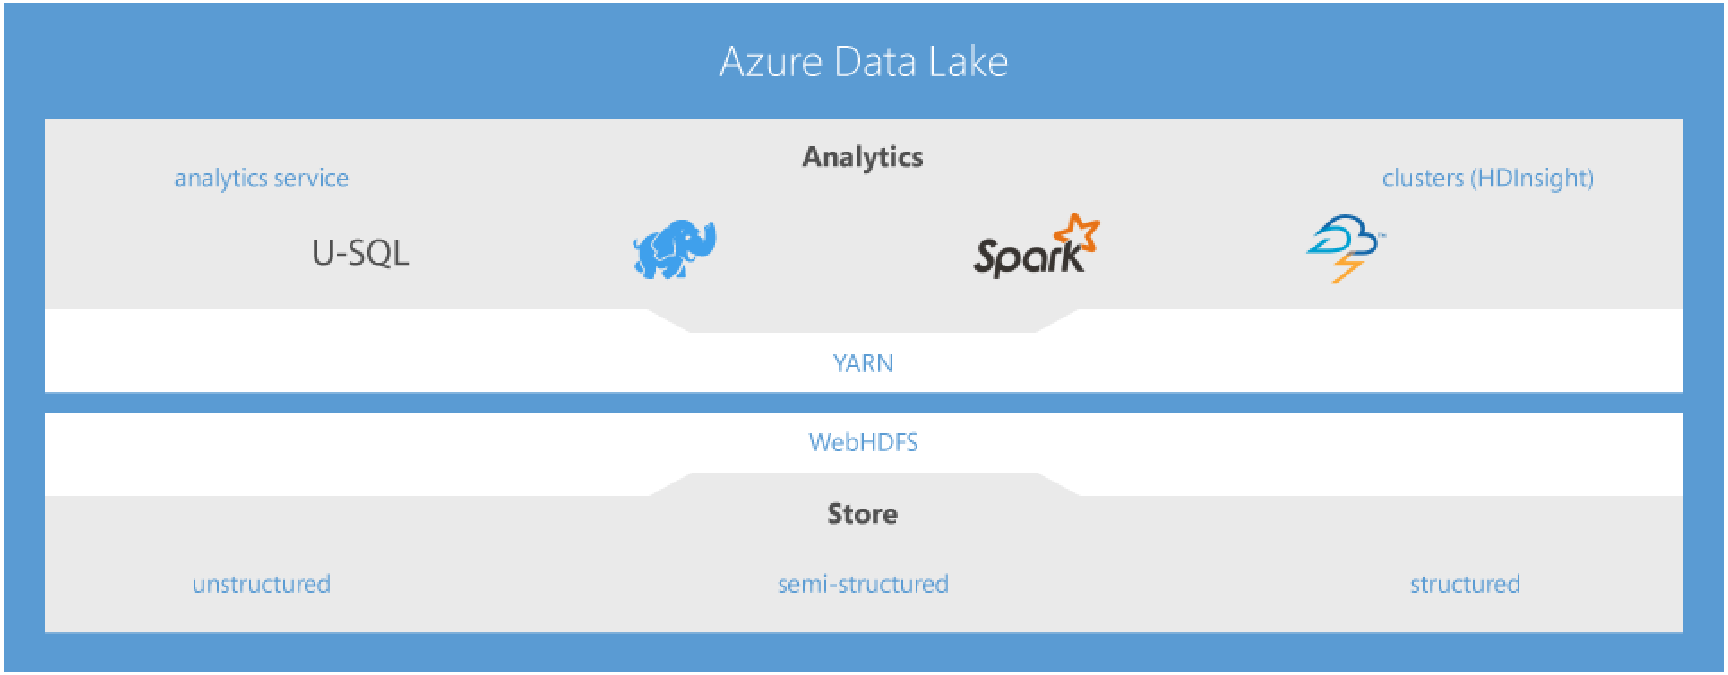

More details at https://azure.microsoft.com/en-us/resources/videos/azure-data-lake-making-big-data-easy/

Storage as a Service (STaaS)

- What is it?
  - A business model in which a company rents space in their storage infrastructure to another company or individual.
- How does it work?
  - STaaS provider rents space
  - cost-per-gigabyte-stored and cost-per-data-transfer basis.
- Benefits
  - Shifting from Capital Expenditure to Operational Expenditure
  - Scale up/down at will (temporarily)

Cloud Storage

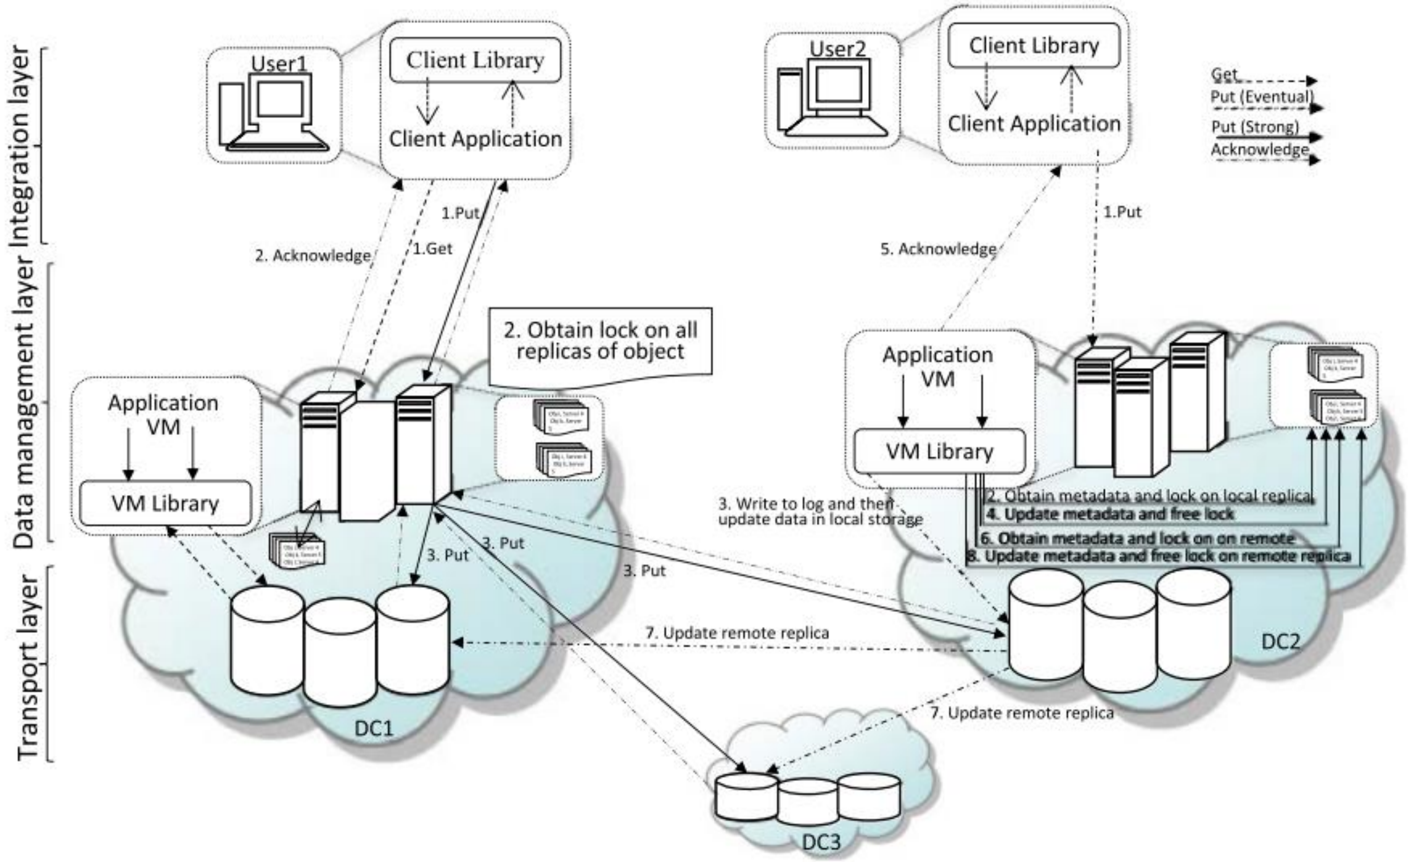

Mansouri et al., **Data Storage Management in Cloud Environments**, ACM CSUR 2018.

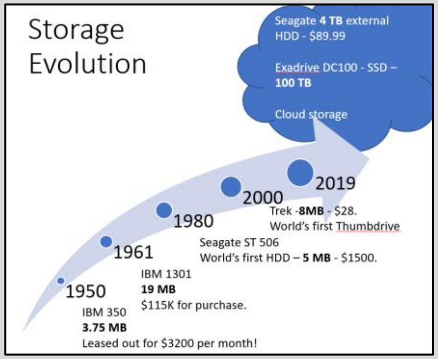

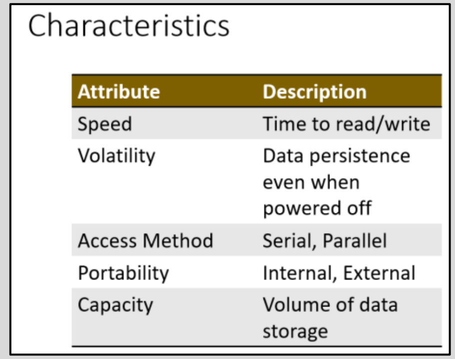
Characteristics

Attribute	Description
Speed	Time to read/write
Volatility	Data persistence even when powered off
Access Method	Serial, Parallel
Portability	Internal, External
Capacity	Volume of data storage

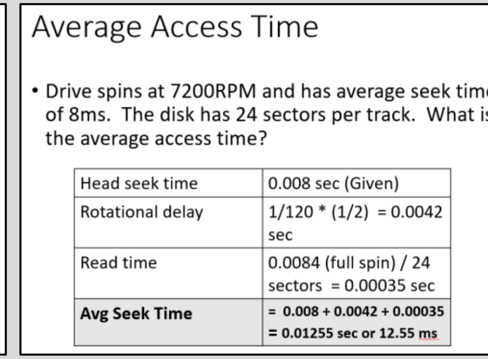
Average Access Time

- Drive spins at 7200RPM and has average seek time of 8ms. The disk has 24 sectors per track. What is the average access time?

Head seek time	0.008 sec (Given)
Rotational delay	1/120 * (1/2) = 0.0042 sec
Read time	0.0084 (full spin) / 24 sectors = 0.00035 sec
Avg Seek Time	**= 0.008 + 0.0042 + 0.00035** **= 0.01255 sec or 12.55 ms**

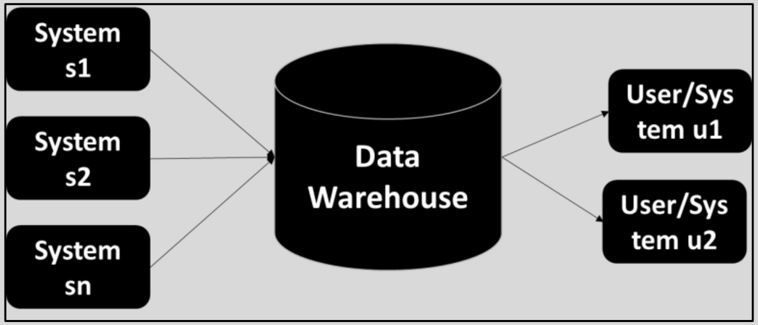

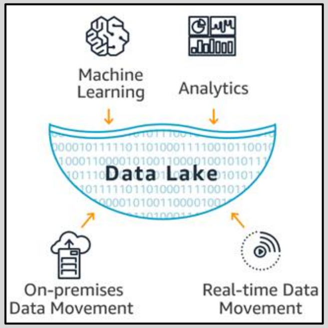

Data Storage

Enablers of Storage – Hardware Summary

Evolution of Computing

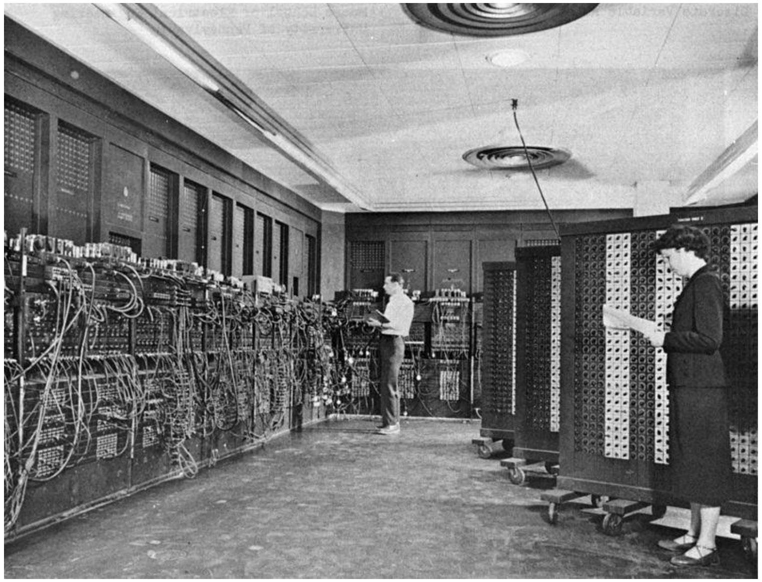

ENIAC
Early 1900s

Stored-program
Von Neumann
Architecture
1940

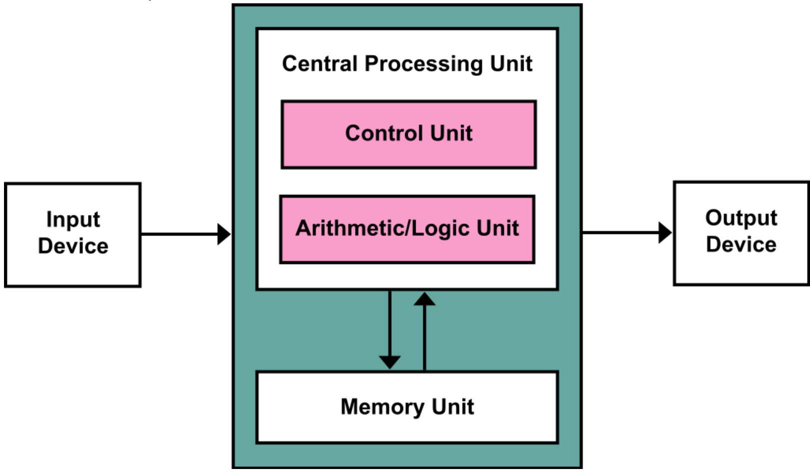

Microprocessors

Processing unit on an integrated circuit

What are ICs made of?

Transistors

- Basic electronic component that alters the flow of current.
- Form the basic building block of an integrated circuit.
- Think of it as an electronic switch

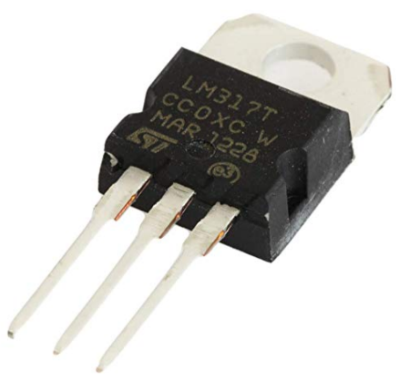

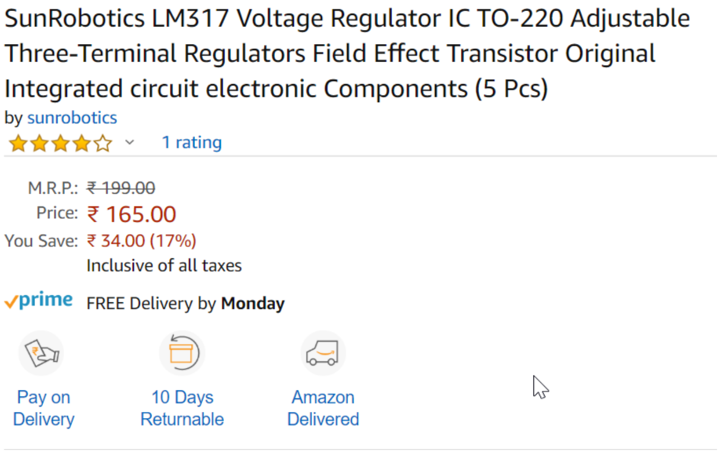

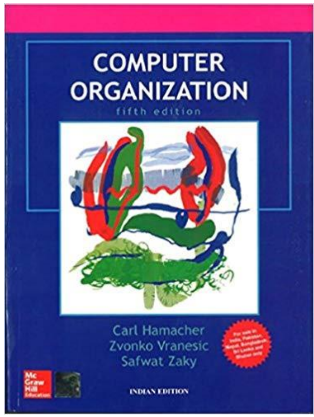

Logic Gates

- Implements Boolean functions (thus performs logical operations)
- Implemented using Transistors

Microprocessors contain millions of logic gates.

Moore’s Law

The number of transistors on a microchip doubles every two years, though the cost of computers is halved.

CES 2019: **Moore's Law** is **dead**, says Nvidia's CEO
CNET - 10-Jan-2019
Intel, for its part, doesn't think **Moore's Law** is **dead**. Companies are just finding new ways to keep it going, like Intel's new 3D chip stacking.

Moores Law is far from death, according to Intel's Jim Keller
TweakTown - 10-Dec-2019
"The death of **Moores Law** is coming and will limit how far we can go ... explanation of what **Moores law** is, from Intel's co-founder, engineer, ...

Moore's Law – The number of transistors on integrated circuit chips (1971-2018)

Our World in Data

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important as other aspects of technological progress – such as processing speed or the price of electronic products – are linked to Moore's law.

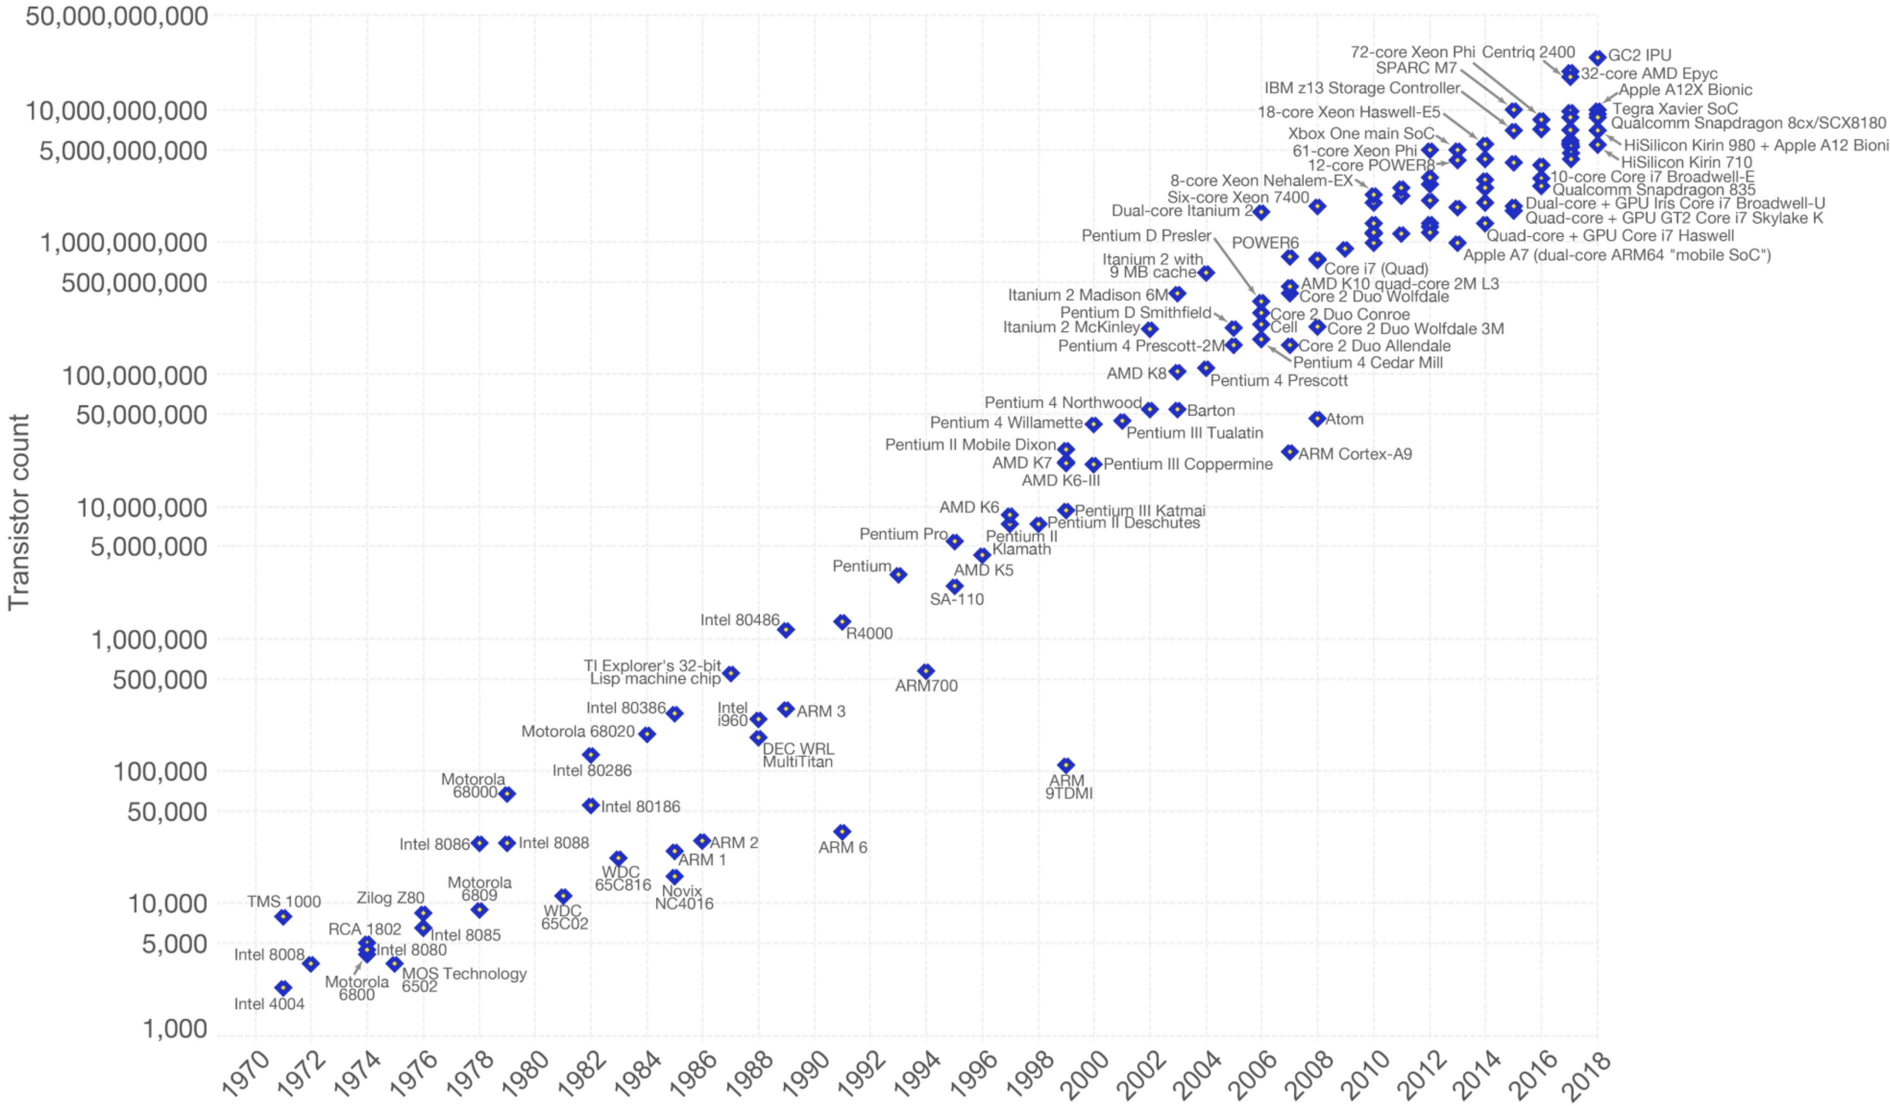

Source: https://www.visualcapitalist.com/visualizing-moores-law-in-action-1971-2019/

Data source: Wikipedia (https://en.wikipedia.org/wiki/Transistor_count)

The data visualization is available at OurWorldinData.org. There you find more visualizations and research on this topic.

Multi-Core Processors

- Two or more separate processing units (called cores)
- Enhances parallel processing

intel core duo
(has 2 cores, 2.66 GHz)

intel core i7
(has 4 cores, 4 GHz)

Quality Up

What do we achieve when we use p processors?

$$\text{Quality Up} = \frac{quality\ on\ p\ processors}{quality\ on\ 1\ processor}$$

Read Section 1.1 of Jan Verschelde's book on "Introduction to Supercomputing".

Can we use multiple processors?

- Amdahl's law
  - Let R be the fraction of the operations which cannot be parallelized. The speedup with p processors is bound by $\frac{1}{R+\frac{1-R}{p}}$.
- Example
  - Say, 10% cannot be parallelized, and we have 8 processors. Best speedup = $\frac{1}{1/10+\frac{1-1/10}{8}}$ ≈ 4.7x.

Multiple Processors for Speedup

- Amdahl's law
  - Let R be the fraction of the operations which cannot be parallelized. The speedup with p processors is bound by $\frac{1}{R+\frac{1-R}{p}}$.

Speed up in terms of problem size.

- Example
  - Say, 10% of the task cannot be parallelized, and we have 8 processors. Best speedup = $\frac{1}{1/10+\frac{1-1/10}{8}}$ ≈ 4.7x.

Speedup

Speed up in terms of time.

- Gustafson's Law
  - If s **is the fraction of serial operations** in a parallel program run on p **processors**, then the scaled speedup is bounded by $p + (1 - p)s$.
- Example
  - Say, all other seven processers are kept idle while one processor completes 5% work, scaled speedup = 8 + (1 – 8) * 0.05 = 7.65.

Our ability to parallelize determines the successful use of multi-core processors.

Supercomputer

- A computing system that provides close to the best currently achievable sustained performance on demanding computational problems.

How do supercomputers achieve such performance levels?

GPUs and GPGPUs

- Graphics Processing Unit (GPU)
  - Massive parallelization
  - Thousands of cores
  - Originally created for the gaming industry
- General Purpose GPU (GPGPU)
  - Architecture allows for programming (Example: Compute Unified Device Architecture (CUDA) on NVIDIA GPGPUs).
- Performance is measured in FLOP (Floating Point Operation)
  - sometimes, FLOPS (floating point operations per second)

CPU vs. GPU

- Say, two floating point operations could be performed in a clock cycle,
  - 3 GHz processor ➔ 6 gigaflop per second.
- Top GPUs achieve petaflop per second.
  - Achieved through an array of cores (V100 has 5120 cores)

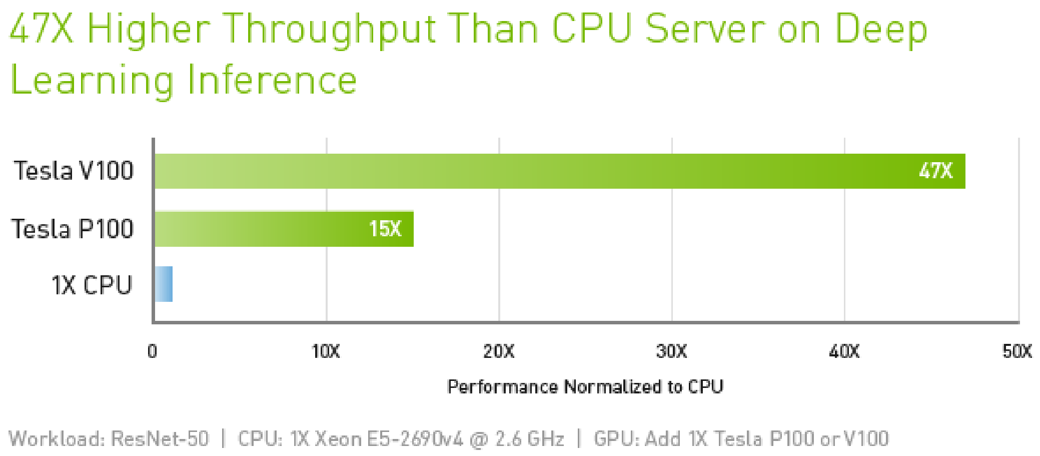

My System

My Lenovo X390 uses
Intel® Core™ i7-8565U CPU @ 1.8 GHz

4 cores only! ☹

Deep Blue

- Beat Chess World Champion Garry Kasparov in 1997

259th most powerful supercomputer.
Achieved 11.38 GFLOPS.

IBM Watson and Jeopardy Game, 2011

- Cluster of **90 servers** each having **3.5GHz eight-core processor** and **16 TB of RAM**.
- Equivalent to 80 Teraflops (a slow supercomputer by today's standards).

Trivia

- Can you name the fastest supercomputer as of date?
  - How much data can it store?
  - How fast is it?

Trivia

- Can you name the fastest supercomputer as of date? IBM SUMMIT
  - How much data can it store? 250 PB
  - How fast is it? 200 petaflops

For more info, see https://www.livescience.com/62827-fastest-supercomputer.html

How fast is 200 petaflops?

Uses NVIDIA Tesla V100 GPU – How fast is its 200 petaflops?

"If every person on Earth completed one calculation per second, it would take the world population 305 days to do what Summit can do in 1 second" - Oak Ridge National Laboratory.

That is 200 quadrillion calculations in one second!

Limitations and Opportunities

- Supercomputers
  - are too expensive
  - still far away from achieving desirable speedups
  - need skilled programming (distributed computing algorithms, parallelizable code)
- But,
  - GPUs are becoming commonplace
  - High Performance Clusters are increasingly available

The Central Question!!!!!

Instead of using supercomputers, can we put commodity hardware into a cluster and achieve speedup?

Computing with Commodity Hardware – Distributed Computing

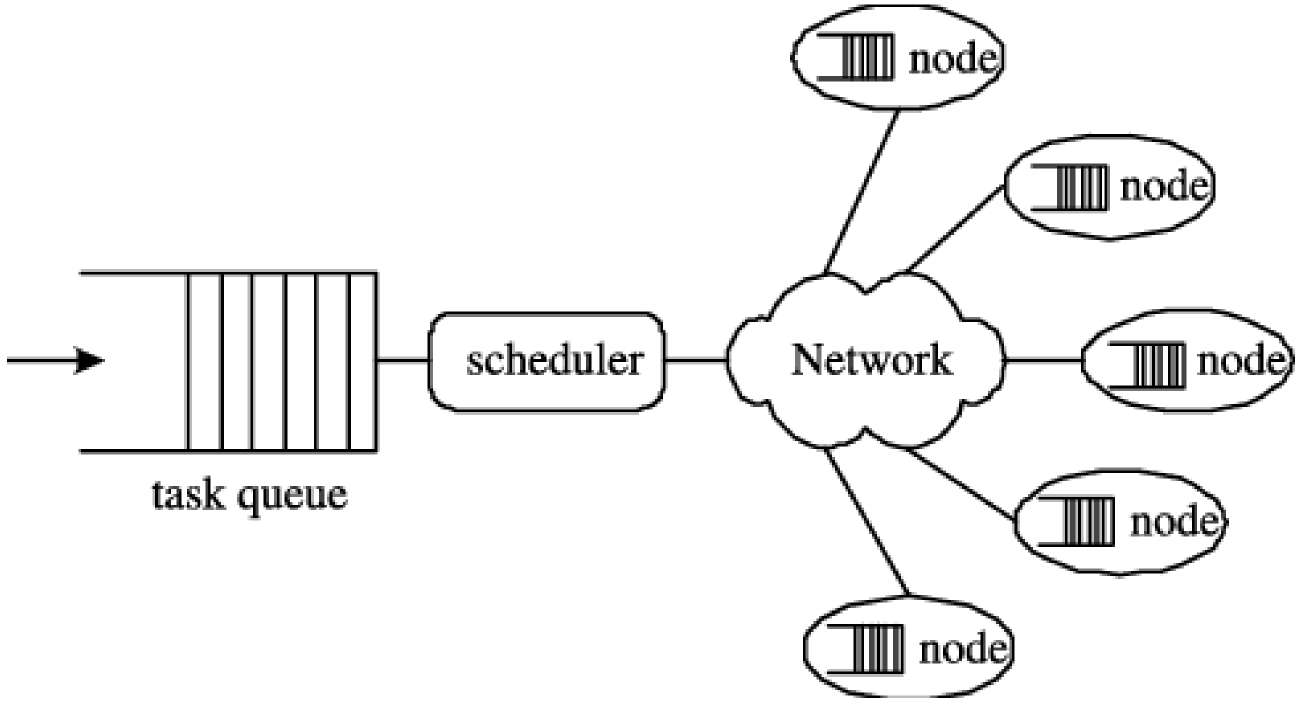

Sun et al., Dynamic Task Flow Scheduling for Heterogeneous Distributed Computing, 2007.

All Roads Lead To.... Cloud

We are in the Big Data era!

Two kinds of Big Data Opportunities

Storage

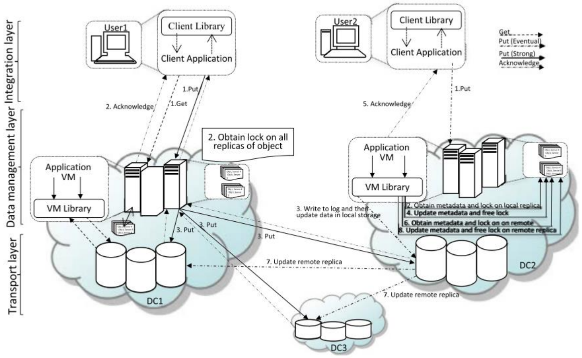

Processing

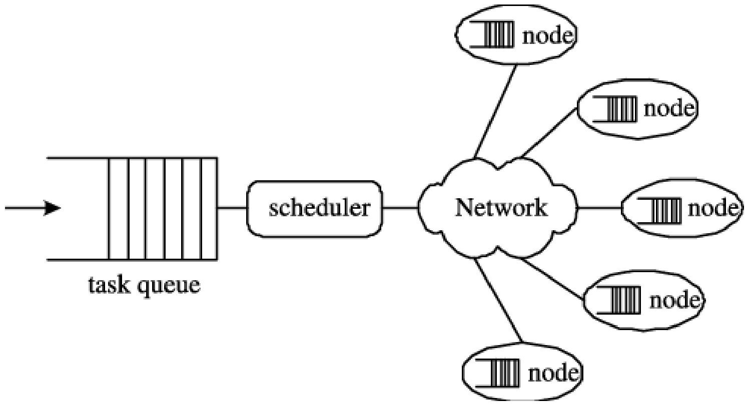

Data Processing

Enablers of Analytics –Summary

CPU → Multi-core processors → GPU → Cloud

Storage

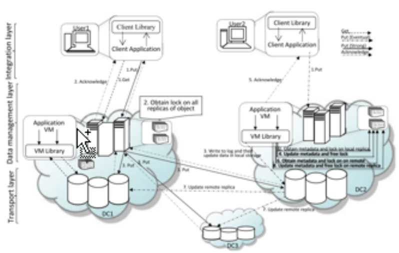

Processing

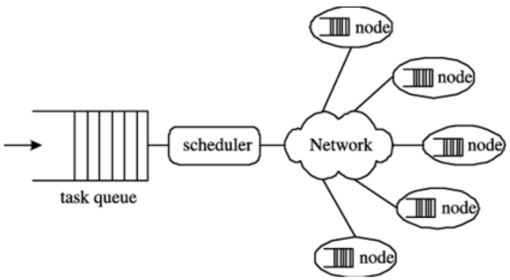

Key Questions

- How to setup and manage such clusters?
- How to achieve reliability, availability, scalability, …?
- How to build services on cloud?

Apache Hadoop

Open source platform - reliable, scalable, distributed processing of large data sets - built on clusters of commodity computers.

Hadoop Cluster

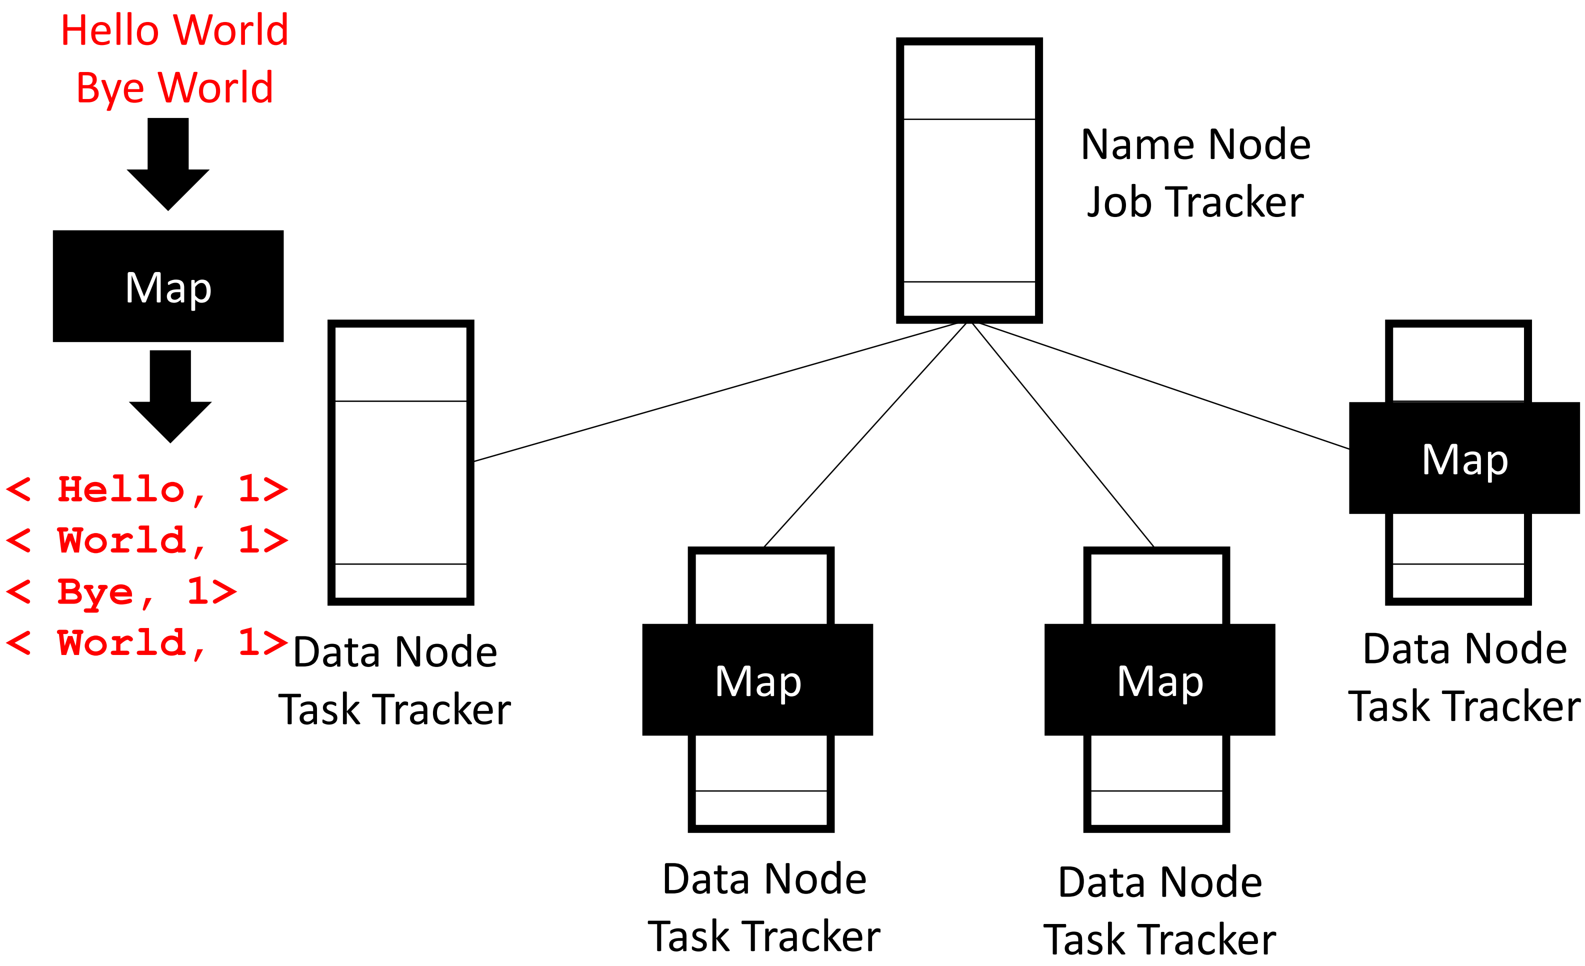

Building Great Apps/Services

- We need products that make certain features easy to implement:
  - Data Acquisition
    - Crawling with Nutch
    - Using Flume Agents
    - Stream Processing with Kafka and Storm
  - A GraphDB – Neo4j
  - Data Visualization with Tableu

There are more challenges and products to deal with them. For example, zookeeper for co-ordination.

Data Acquisition - Crawling with Nutch

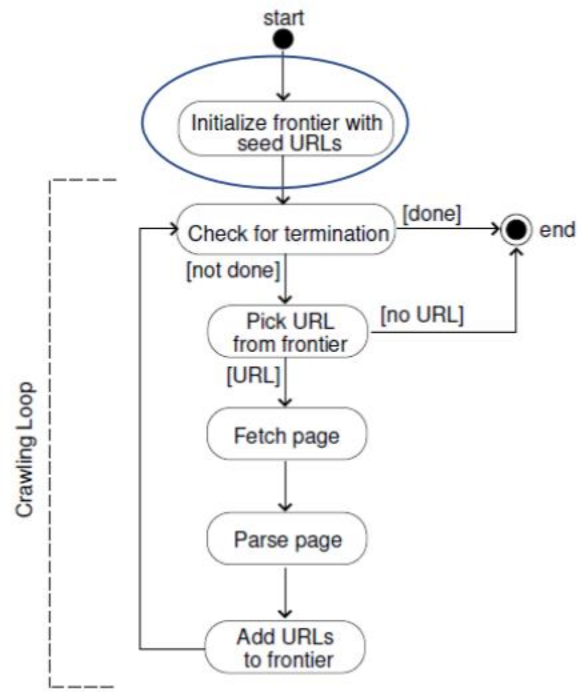

[Solr Integration Image Src: https://suyashaoc.wordpress.com/2016/12/04/nutch-2-3-1-hbase-0-98-8-hadoop-2-5-2-solr-4-1-web-crawling-and-indexing/](https://suyashaoc.wordpress.com/2016/12/04/nutch-2-3-1-hbase-0-98-8-hadoop-2-5-2-solr-4-1-web-crawling-and-indexing/)

Data Acquisition - Flume

Flume

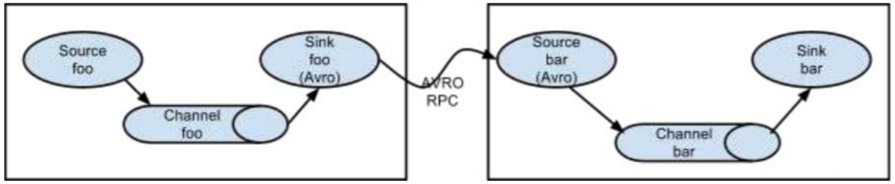


```
usingFlume.sources = usingFlumeSource
usingFlume.channels = memory

usingFlume.sources.usingFlumeSource.type = avro
usingFlume.sources.usingFlumeSource.channels = memory
usingFlume.sources.usingFlumeSource.port = 7877
usingFlume.sources.usingFlumeSource.bind = 0.0.0.0
```

Flume Config Files

Flume event

Stream Processing

- Process data as they arrive.

Stream Processing with Storm

One of these is a master node. “**Nimbus**” is the “job tracker”!

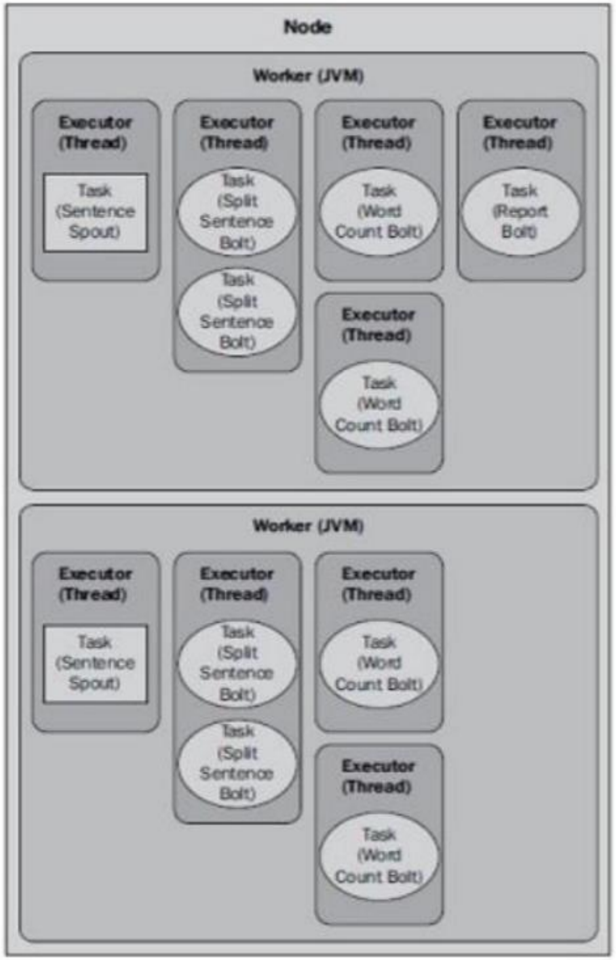

In MR parlance, “**Supervisor**” process is our “task tracker”

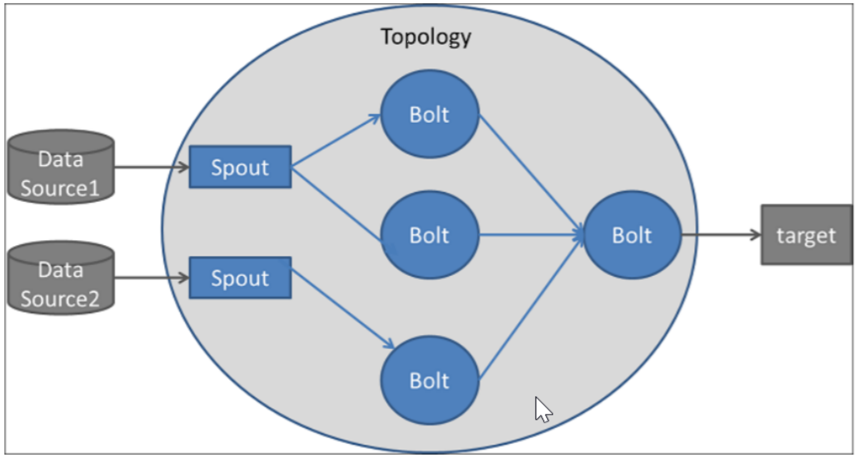

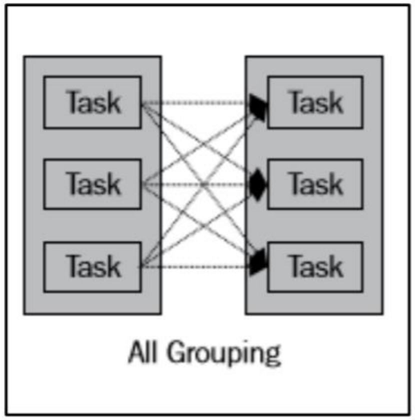

Apache Kafka

- Uses Publish-Subscribe Mechanism

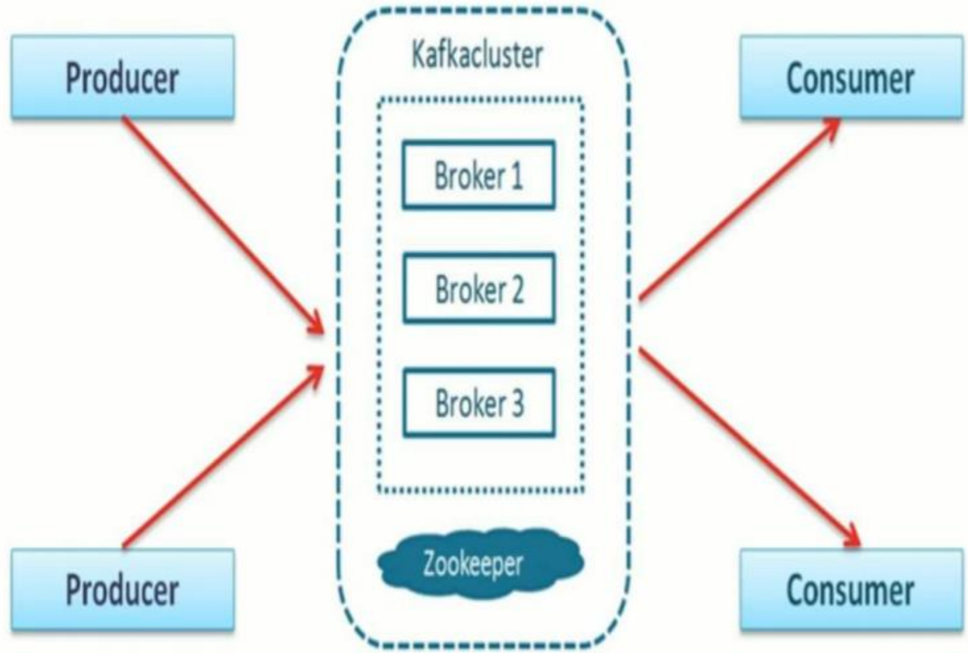

Kafka – Multi-node

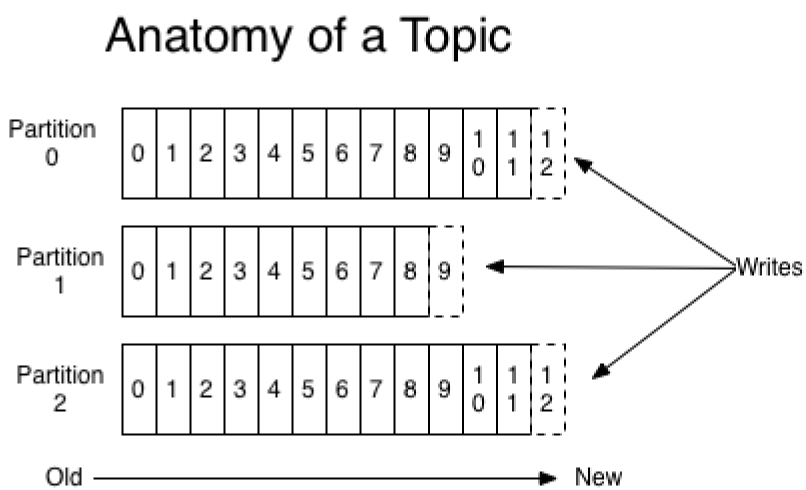

- topic is a stream of records.
- for each topic, the Kafka cluster maintains a partitioned log
- records in the partitions are each assigned a sequential id number called the *offset*

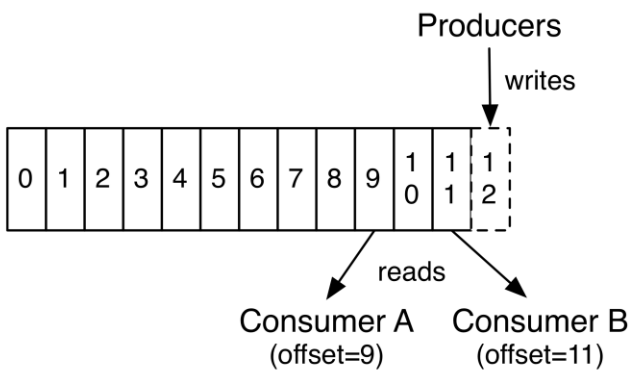

Apache Spark (A Unified Library)

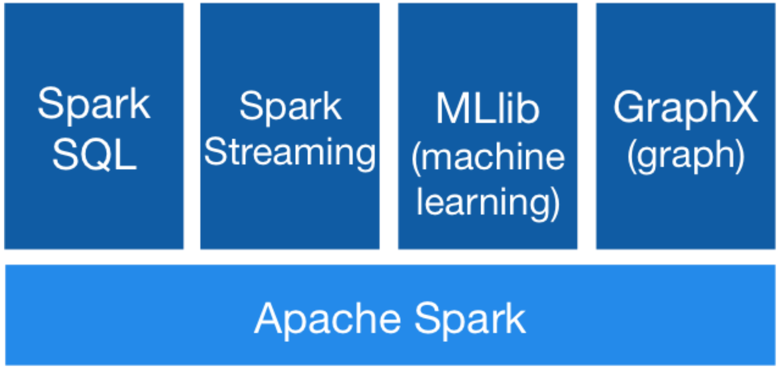


```
df = spark.read.json("logs.json")
df.where("age > 21")
  .select("name.first").show()
```

In spark, use data frames as tables

Spark's Python DataFrame API
Read JSON files with automatic schema inference

https://spark.apache.org/

Resilient Distributed Datasets (RDDs)

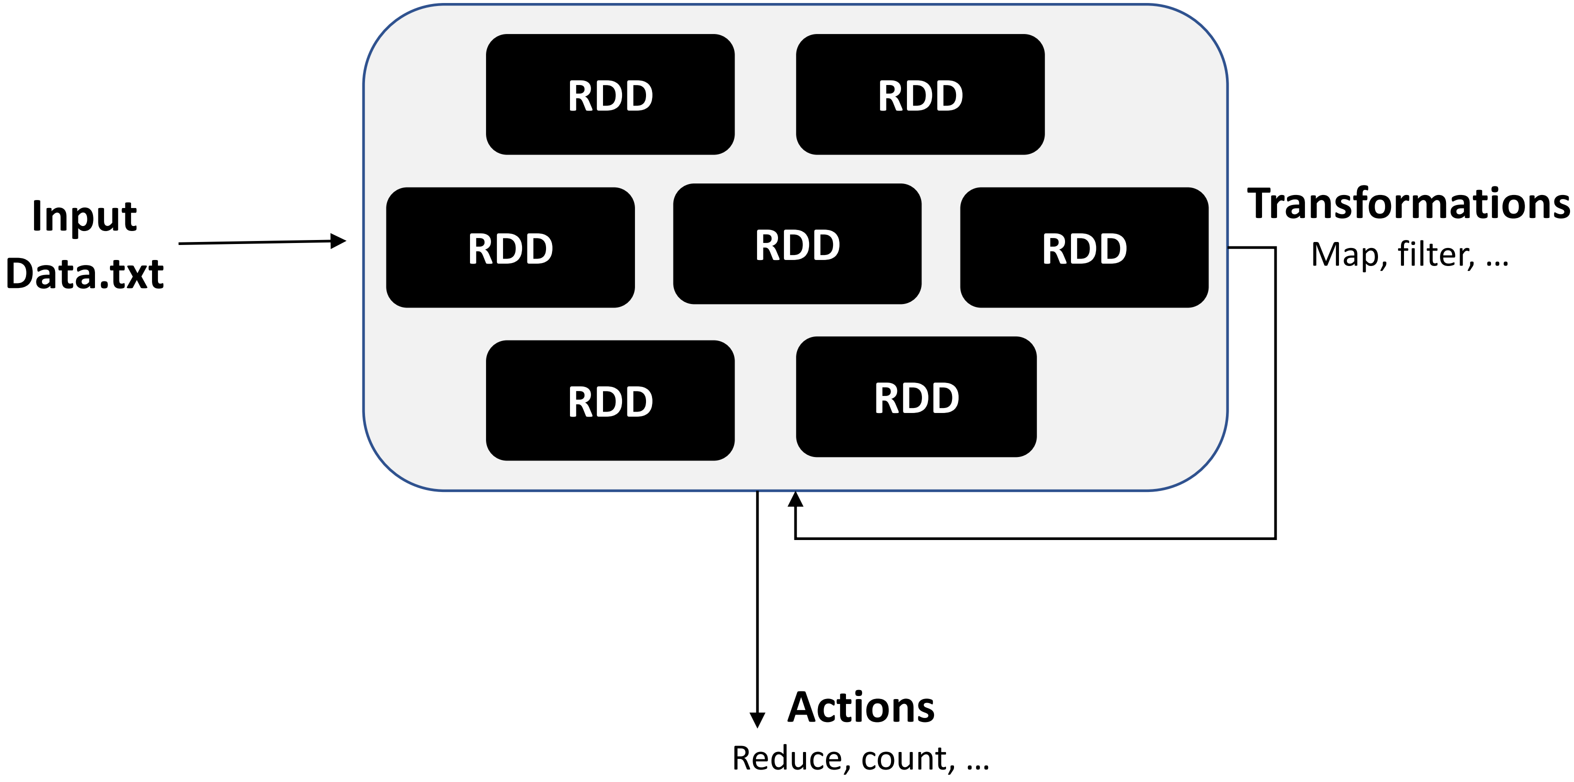

Spark Examples

```
data = [1, 2, 3, 4, 5]
distData = sc.parallelize(data)
```

distributed dataset can be used in parallel

```
"""MyScript.py"""
if __name__ == "__main__":
    def myFunc(s):
        words = s.split(" ")
        return len(words)

    sc = SparkContext(...)
    sc.textFile("file.txt").map(myFunc)
```

passing functions through spark

```
distFile = sc.textFile("data.txt")
distFile.map(s => s.length).
         reduce((a, b) => a + b)
```

Map/reduce

Data Stores on the Cloud

- The NoSQL DBs!

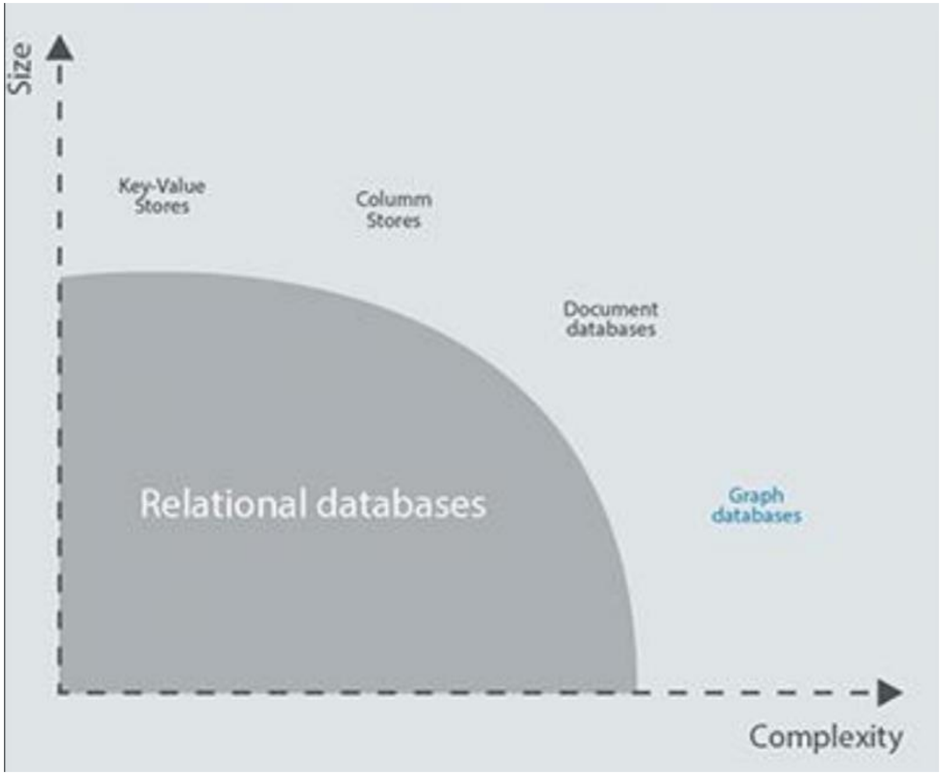

GraphDB – Neo4j

Some data are best seen as graphs – like friends network on facebook

neo4j

neo4j	Cypher	Most famous graph database, Cypher O(1) access using fixed-size array
DSE Graph DATASTAX AURELIUS	Gremlin	Distributed graph system based on Cassandra
ArangoDB	AQL	Multi-model database (Document + Graph)
OrientDB	OQL	Multi-model database (Document + Graph)

Neo4j

- A leading graph database, with native graph storage and processing.
- Open Source
- NoSQL
- ACID compliant

Neo4j Sandbox

https://sandbox.neo4j.com/

Neo4j Desktop

https://neo4j.com/download

Data Model

- create (p:Person {name:'Venkatesh'})-[:Teaches]->(c:Course {name:'BigData'})

Query Language

- Cypher Query Language
  - Similar to SQL
  - Optimized for graphs
  - Used by Neo4j, SAP HANA Graph, Redis Graph, etc.

CQL

- create (p:Person {name:'Venkatesh'})-[:Teaches]->(c:Course {name:'BigData'})
- Don’t forget the single quotes.

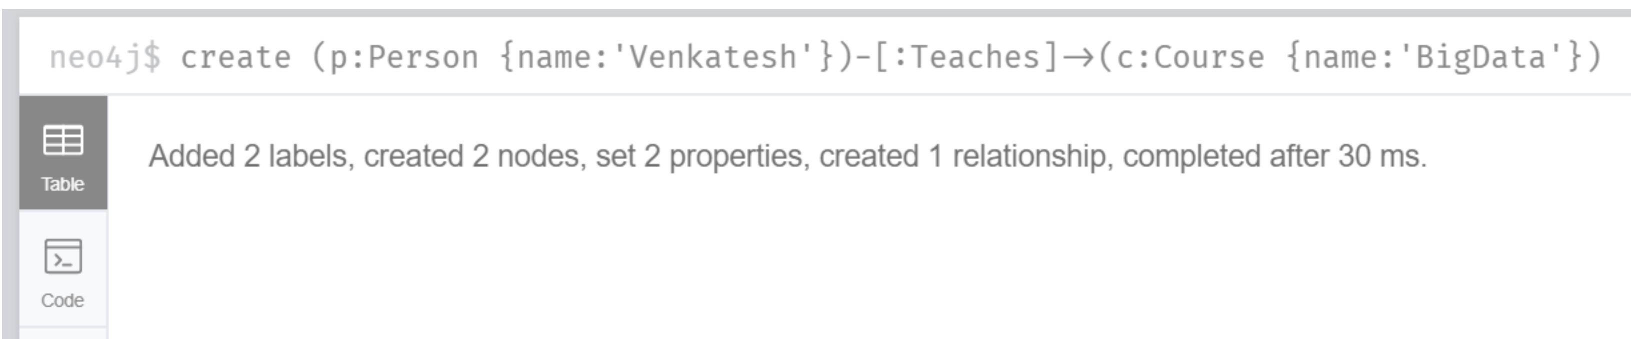

CQL

- Match (n) return n

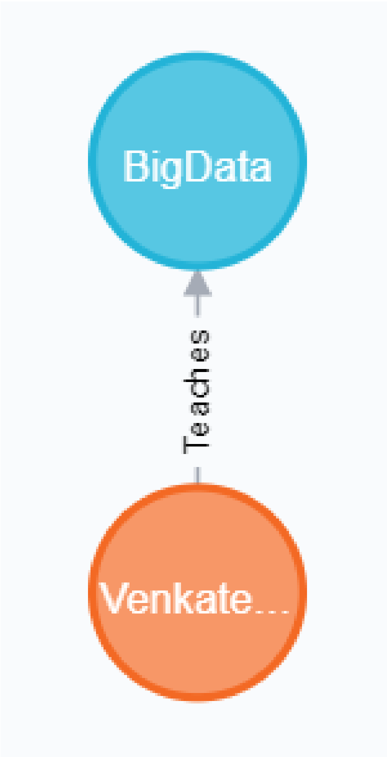

CQL

- match(p:Person {name:'Venkatesh'}) set p.surname='Vinayakarao' return p

CQL

- Create (p:Person {name:’Raj’})-[:StudentOf]->(o:Org {name:’CMI’})
- Match (n) return n

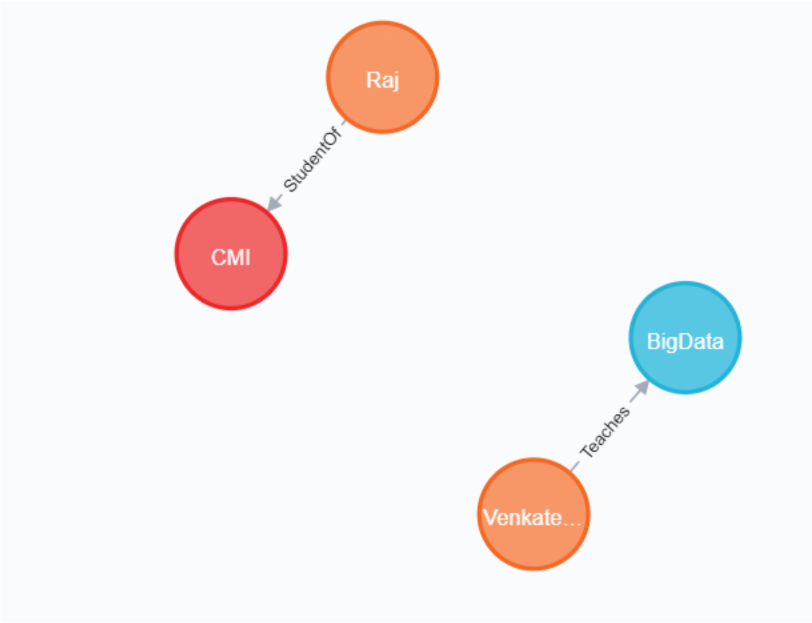

CQL

- create (p:Person {name:'Venkatesh'})-[:FacultyAt]->(o:Org {name:'CMI'})
- Match (n) return n

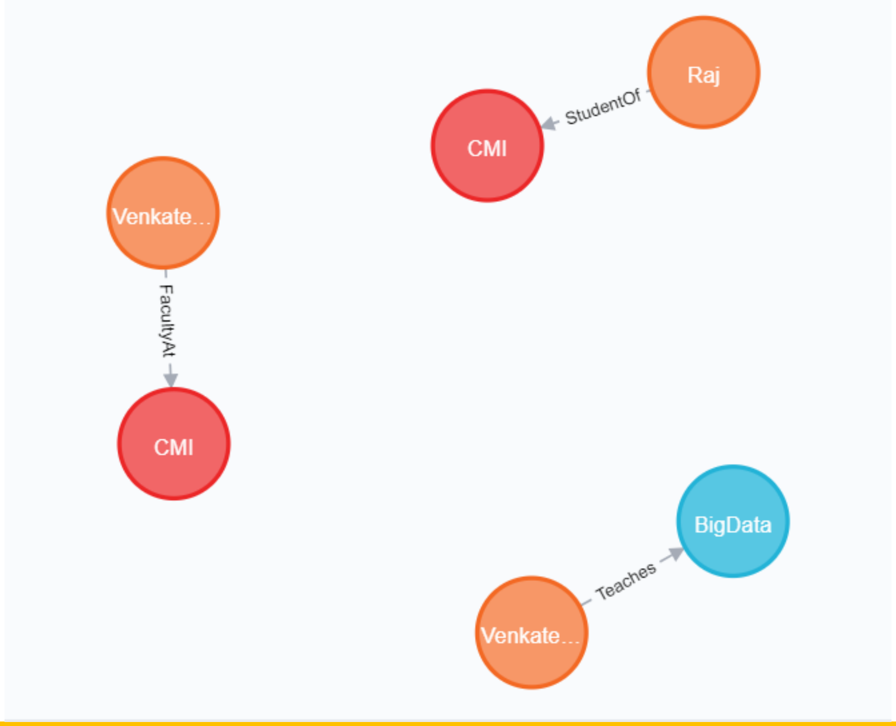

CQL

- MATCH (p:Person {name:'Venkatesh'})-[r:FacultyAt]->()
- DELETE r
- MATCH (p:Person) where ID(p)=4
- DELETE p
- MATCH (o:Org) where ID(o)=5
- DELETE o

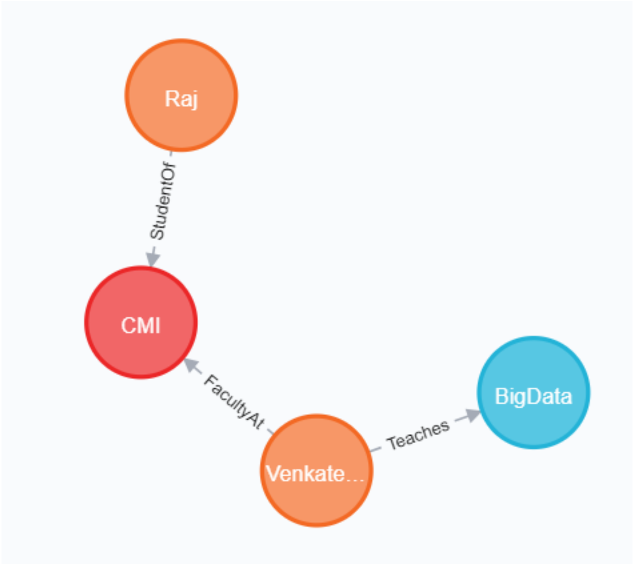

- MATCH (a:Person),(b:Org)
- WHERE a.name = 'Venkatesh' AND b.name = 'CMI'
- CREATE (a)-[:FacultyAt]->(b)

CQL

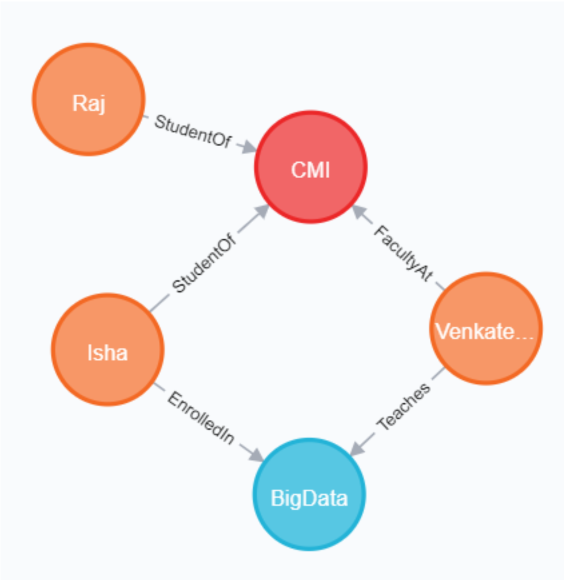


```
create (p:Person {name:'Isha'})

MATCH (a:Person),(b:Course)
WHERE a.name = 'Isha' and b.name = 'BigData'
CREATE (a)-[:StudentOf]->(b)

MATCH (a:Person)-[o:StudentOf]->(b:Course) where a.name = 'Isha' DELETE o

MATCH (a:Person),(b:Org)
WHERE a.name = 'Isha' and b.name = 'CMI'
CREATE (a)-[:StudentOf]->(b)

MATCH (a:Person),(b:Course)
WHERE a.name = 'Isha' and b.name = 'BigData'
CREATE (a)-[:EnrolledIn]->(b)
```

Visualization with Tableau

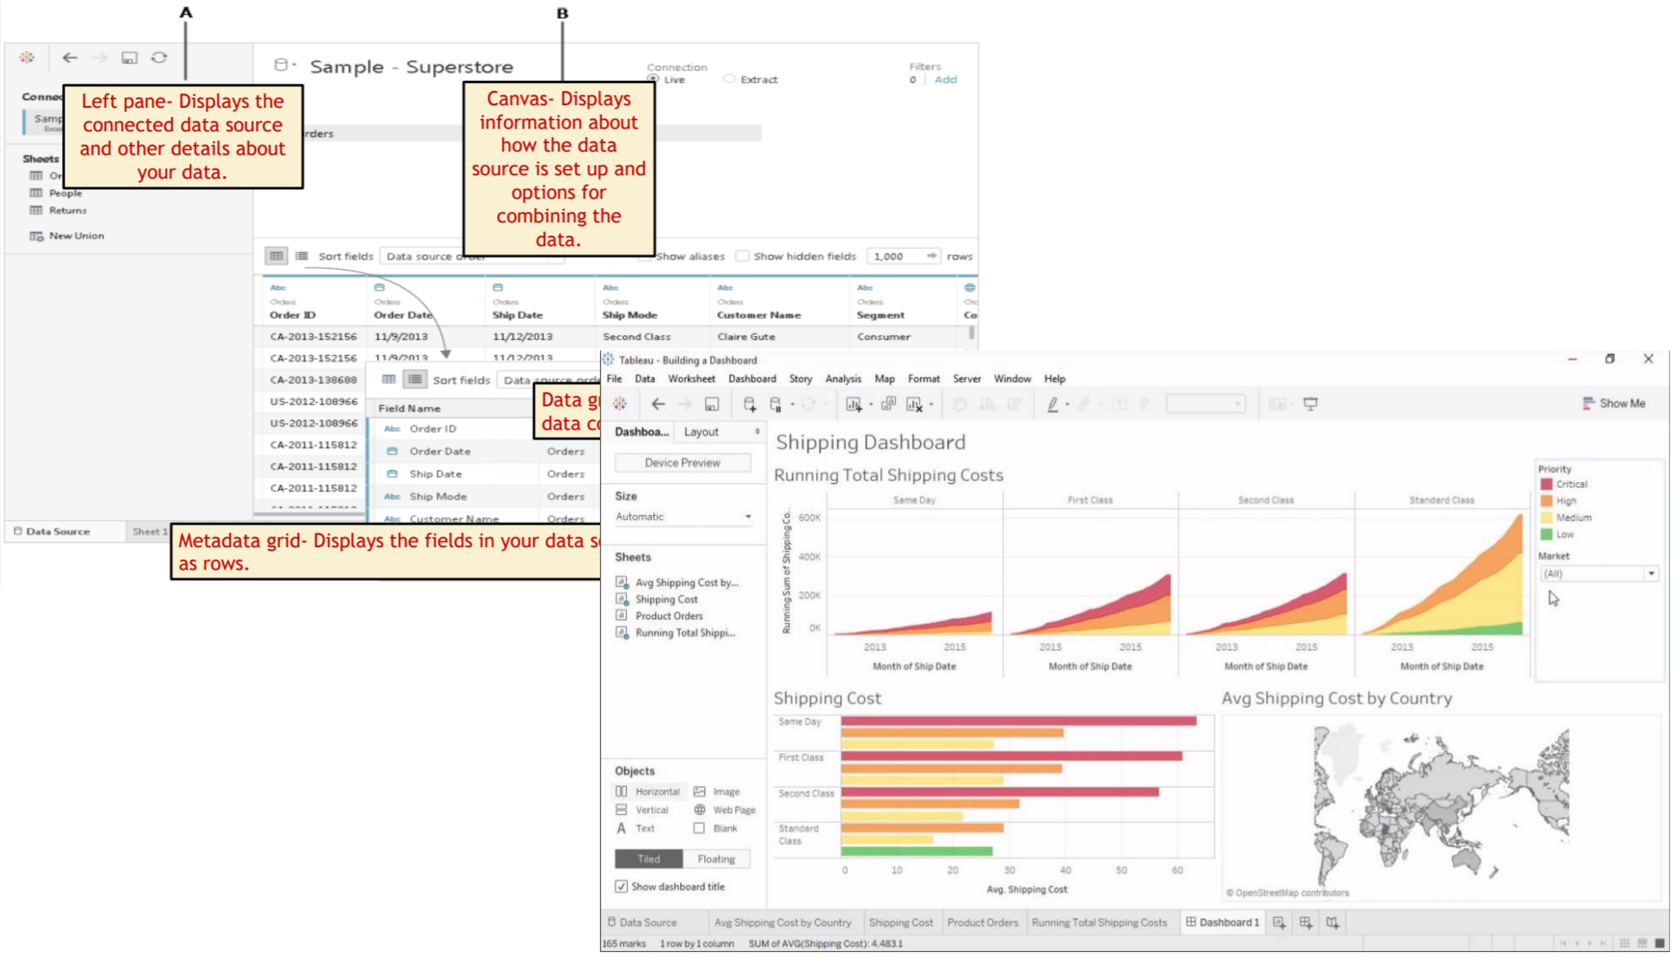

Cloud Platforms

Microsoft Azure

aws

Ref: https://maelfabien.github.io/bigdata/gcps_1/#what-is-gcp

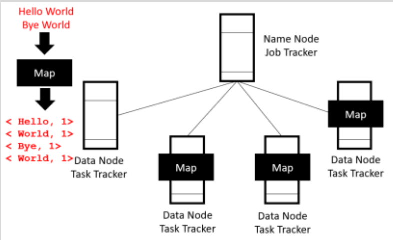

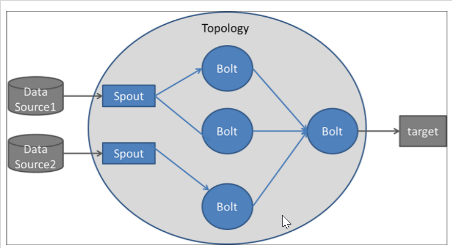

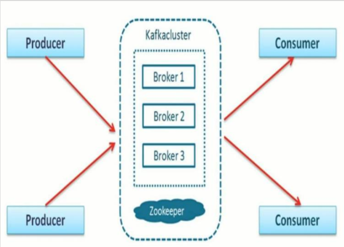

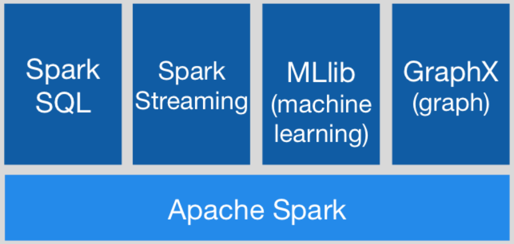

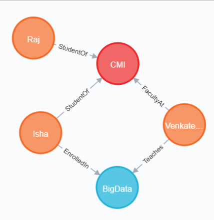

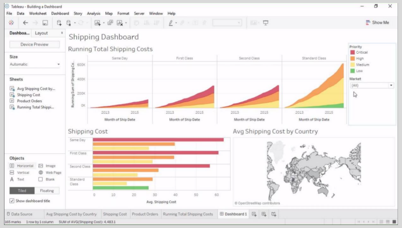

aws

Microsoft Azure

Thanks to these platforms, building effective analytics is now easy!

Data Acquisition & Visualization

Challenges in Analytics –Summary

Storage Evolution

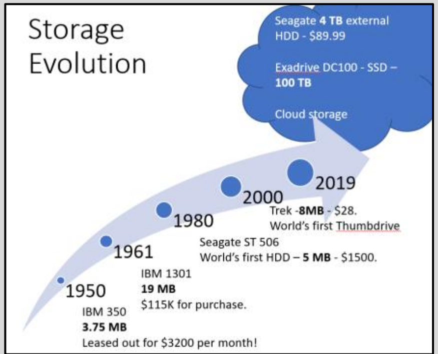

(Secondary) Storage Technologies

Average Access Time

- Drive spins at 7200RPM and has average seek time of 8ms. The disk has 24 sectors per track. What is the average access time?

Head seek time	0.008 sec (Given)
Rotational delay	1/120 * (1/2) = 0.0042 sec
Read time	0.0084 (full spin) / 24 sectors = 0.00035 sec
Avg Seek Time	**= 0.008 + 0.0042 + 0.00035 = 0.01255 sec or 12.55 ms**

Characteristics

Attribute	Description
Speed	Time to read/write
Volatility	Data persistence even when powered off
Access Method	Serial, Parallel
Portability	Internal, External
Capacity	Volume of data storage

Microsoft® SQL Server®

MySQL®

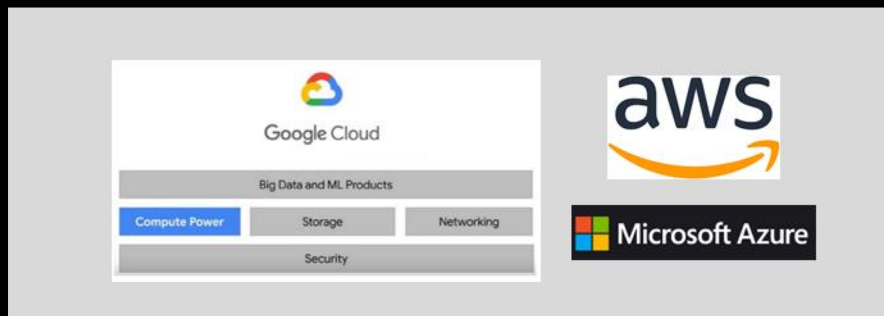

CPU → Multi-core processors → GPU → Cloud

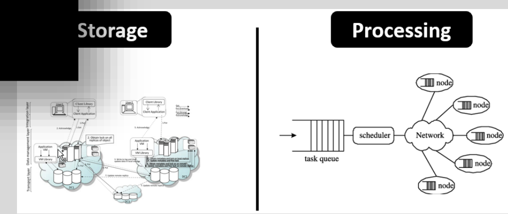

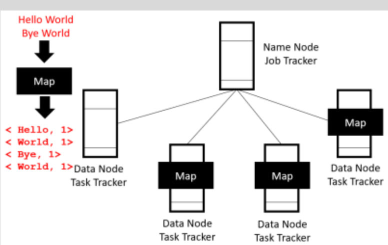

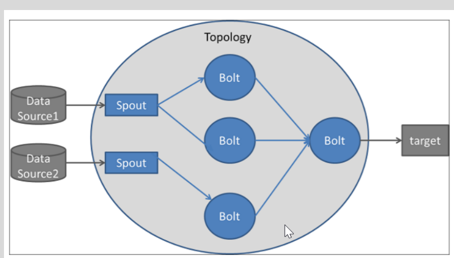

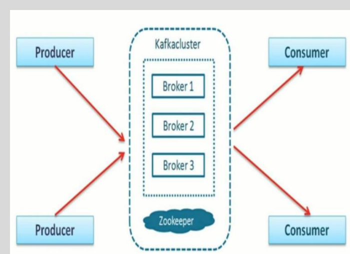

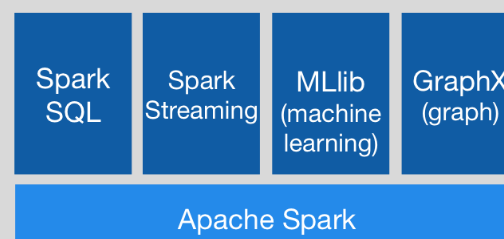

Thank You

venkateshv@cmi.ac.in

Slides are available at http://vvtesh.co.in