An Introduction to Data Intensive Computing

Amir H. Payberah
payberah@kth.se
28/08/2018

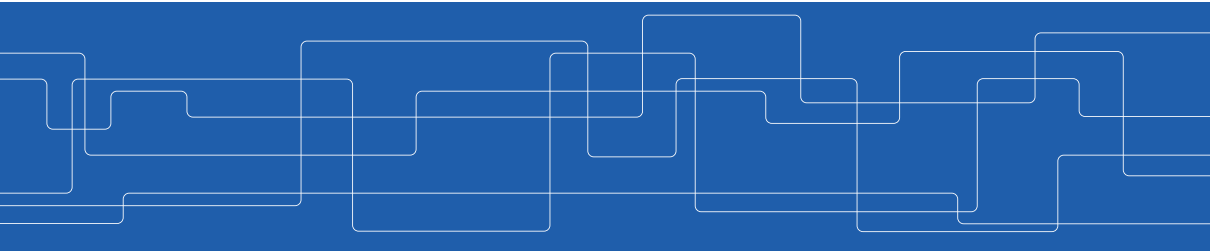

Course Information

Course Objective

- Introduction to main concepts and principles of cloud computing and data intensive computing.
- How to read, review and present a scientific paper.

Topics of Study

- Topics we will cover include:
  - How to store big data?
  - How to process big data?
  - How to manage cluster resources?

The Course Material

- Mainly based on research papers.
- We also cover the following books.

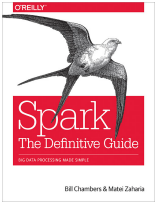

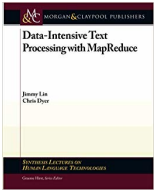

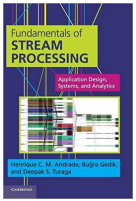

The Course Web Page

https://id2221kth.github.io

The Course Grading

- Four lab assignments: 30%
- Three reading assignments: 15%
- One project: 15%
- The final exam: 40%

The Lab Assignments and The Project

- Self-selected groups of two
- Labs will include Scala/Java programming
  - Lab1: HDFS, HBase, and MapReduce
  - Lab2: Spark and Spark SQL
  - Lab3: Kafka, Spark Streaming, and Cassandra
  - Lab4: GraphX
- Project
  - A self-defined project
  - Demonstrated as a demo and short report

The Reading Assignments

- Three reading assignments.
- Write a review for each paper (at most three pages).
- For each paper you should identify, the motivation, the contribution, the solution, and positive/negative aspects of the solution/paper.
- Students will work in groups of two.

How to Submit the Assignments?

- Through the Canvas site.

The Course Overview

Cloud Computing and Big Data

- The main trends:
  - Computers not getting any faster
  - Internet connections getting faster
  - More people connected to the Internet

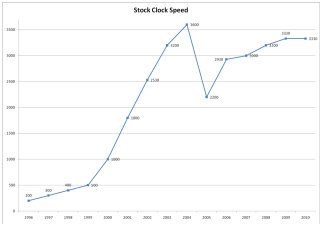

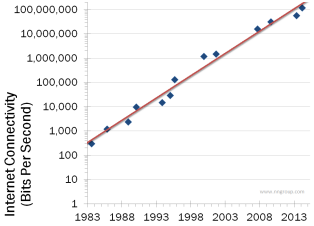

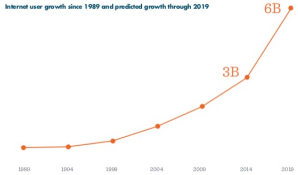

Conclusion

Move the computation and storage of big data to the cloud!

Cisco predicts that by 2020, 92% of IT market workloads will be processed by cloud data centers, while only 8% will be processed by traditional data centers.

Cloud Computing

Cloud Computing Definition

- Cloud Computing refers to both:
  1. The applications delivered as services over the Internet
  2. The hardware and systems software in the datacenters that provide those services

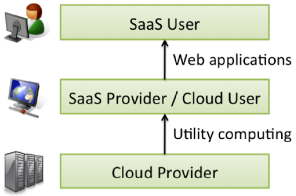

Cloud Computing Definition

- Cloud Computing refers to both:
  1. The applications delivered as services over the Internet
  2. The hardware and systems software in the datacenters that provide those services
- The services: called Software as a Service (SaaS)

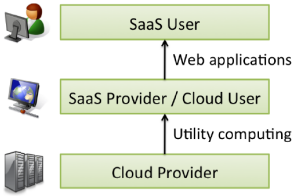

Cloud Computing Definition

- Cloud Computing refers to both:
  1. The applications delivered as services over the Internet
  2. The hardware and systems software in the datacenters that provide those services
- The services: called Software as a Service (SaaS)
- The datacenter hardware and software is called cloud

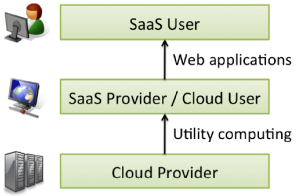

- The NIST definition:
  - Five characteristics
  - Three service models
  - Four deployment models

National Institute of Standards and Technology
Technology Administration, U.S. Department of Commerce

Cloud Characteristics

Cloud Characteristics

On-demand self-service

Ubiquitous network access

Location transparent resource pooling

Rapid elasticity

Measured service with pay per use

[http://aka.ms/532]

Cloud Characteristics - On-demand Self-Service

- A consumer can independently provision computing capabilities without human interaction with the service provider.

On-demand
self-service

Cloud Characteristics - Ubiquitous Network Access

- Available over the network
- Accessed through mobile phones, laptops, ...

Ubiquitous network access

Cloud Characteristics - Resource Pooling

- Provider's computing resources are pooled to serve consumers
- Location transparent

Location transparent resource pooling

Cloud Characteristics - Rapid Elasticity

- Capabilities can be rapidly and elastically provisioned, in some cases automatically.

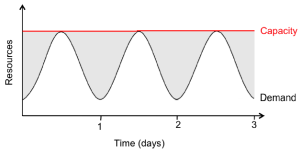

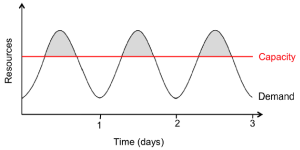

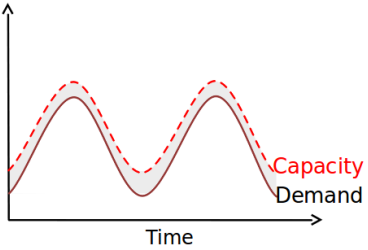

Cloud Characteristics - Measured Service

- Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer.

Measured service with pay per use

Cloud Service Models

Cloud Service Models

SaaS

PaaS

IaaS

[http://aka.ms/532]

- Assume, you just moved to a city and you are looking for a place to live.

- What is your choice?

- What is your choice?
  - Built a new house?
  - Buy an empty house?
  - Live in a hotel?

- Let's built a new house!

- Let's built a new house!
- You can fully control everything your like your new house to have.
- But that is a hard work.

- What if you buy an empty house?

- What if you buy an empty house?
- You can customize some part of your house.
- But never change the original architecture.

- How about live in a hotel?

- How about live in a hotel?
- Live in a hotel will be a good idea if the only thing you care is enjoy your life.
- There is nothing you can do with the house except living in it.

Let's translate it to Cloud Computing

Service Models

- Infrastructure as a Service (IaaS): similar to build a new house.
- Platform as a Service (PaaS): similar to buy an empty house.
- Software as a Service (SaaS): similar to live in a hotel.

IaaS

- Vendor provides resources, e.g., processing, storage, network, ...
- Consumer is provided customized virtual machines.
- Example: Amazon Web Services (EC2 instances and S3 storage)

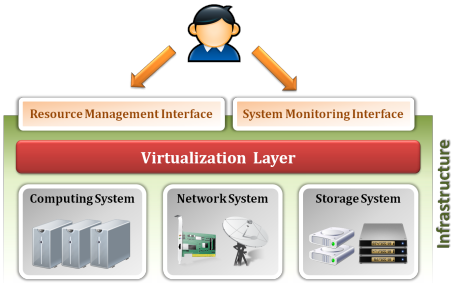

PaaS

- Vendor provides hardware and development environment.
- Example: Google app engine

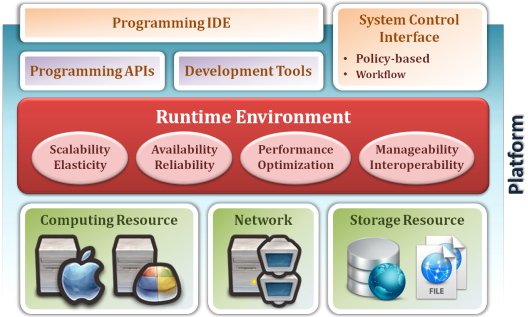

SaaS

- Vendor provides applications accessed over the network.
- Example: Gmail, Github

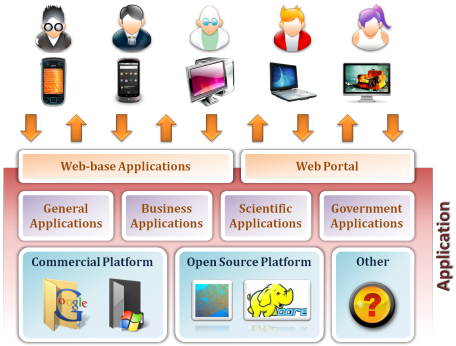

IaaS - PaaS - SaaS

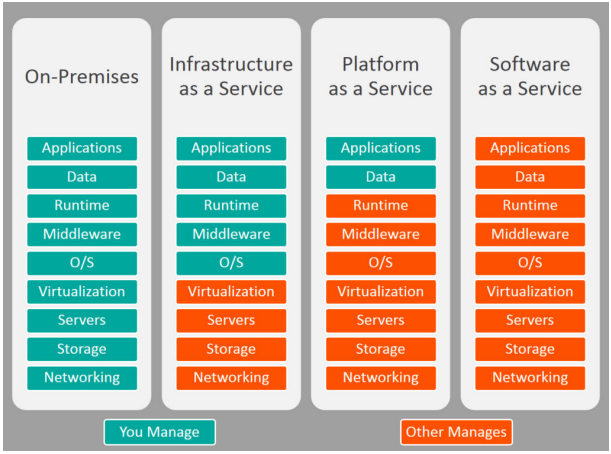

[https://goo.gl/xMko1z]

Deployment Models

Deployment Models

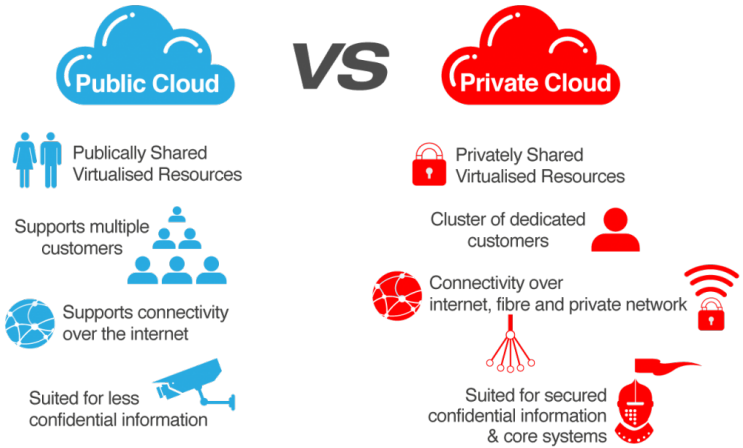

[https://goo.gl/fWmcGK]

Public Cloud Infrastructure Vendors

- Amazon Web Services (AWS)
- Microsoft Azure
- Google Cloud Platform
- IBM Bluemix
- ...

Main Services

- Computing
- Storage
- Database
- Big data analytics
- ...

Computing Services

- Virtual machines
- Container services
- Serverless compute

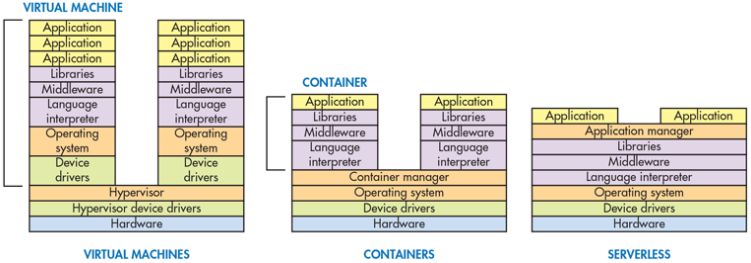

Storage Services

- File storage
- Block storage
- Object storage

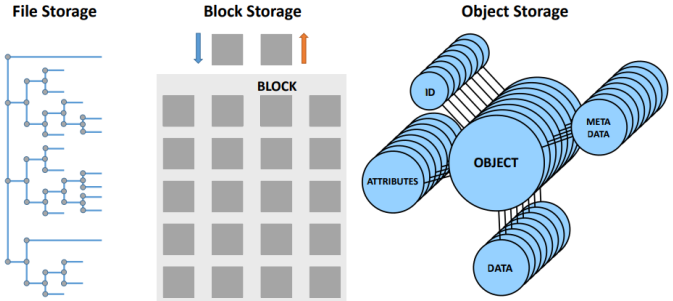

Database Services

- Relational Database Management Services (RDBMS)
- NoSQL databases
- In-Memory data services

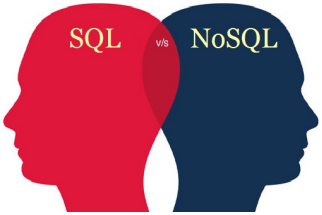

Big Data Analytics

- Big Data Managed Cluster-as-a-Service
- Data warehouse
- Data streaming
- Data queuing

Big Data

Big Data

... everyone talks about it, nobody really knows how to do it, everyone thinks everyone else is doing it, so everyone claims they are doing it.

- Dan Ariely

Big Data

Big data is the data characterized by 4 key attributes: volume, variety, velocity and value.

ORACLE®

Big Data

Big data is the data characterized by 4 key attributes: volume, variety, velocity and value.

Buzzwords

ORACLE

Big Data in Simple Words

DevOps Borat
@DEVOPS_BORAT

Small Data is when is fit in RAM. Big Data is when is crash because is not fit in RAM.

2/6/13, 8:22 AM




The Four Dimensions of Big Data

- Volume: data size
- Velocity: data generation rate
- Variety: data heterogeneity
- This 4th V is for Vacillation: Veracity/Variability/Value

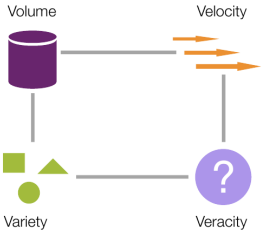

Big Data Sources

How Much Data?

How To Store and Process Big Data?

Problem

- Traditional platforms fail to show the expected performance.
- Need new systems to store and process large-scale data

Scale Up vs. Scale Out (1/2)

- Scale up or scale vertically: adding resources to a single node in a system.
- Scale out or scale horizontally: adding more nodes to a system.

Scale Up vs. Scale Out (2/2)

- Scale up: more expensive than scaling out.
- Scale out: more challenging for fault tolerance and software development.


APACHE
HBASE
mahout
hadoop
GraphX
kafka
Storm
Dato
Spark
HIVE
S4 distributed stream computing platform
cassandra
Google Cloud Platform

Big Data Stack

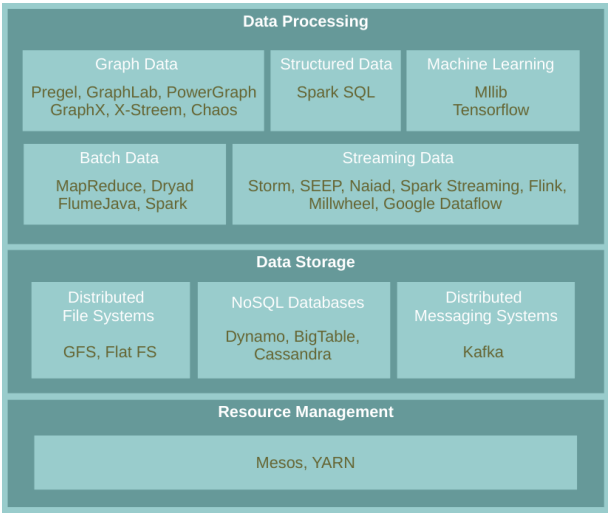

Resource Management

- Manage resources of a cluster
- Share them among the platforms
- Mesos, YARN, Borg, ...

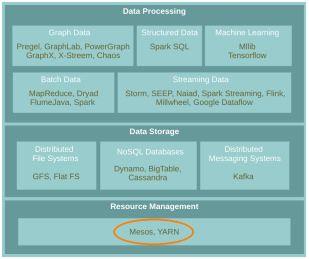

Data Storage - Distributed File Systems

- Store and retrieve files on/from distributed disks
- GFS, HDFS, FlatFS, ...

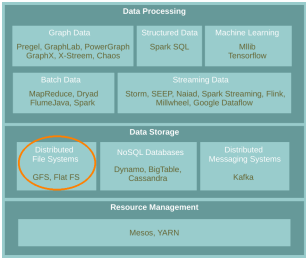

Data Storage - NoSQL Databases

- BASE instead of ACID
- BigTable, Dyanamo, Cassandra, ...

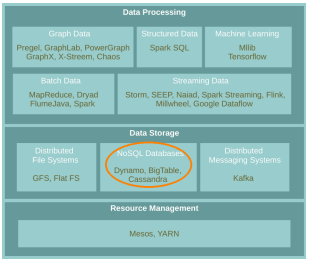

Data Storage - Messaging Systems

- Store streaming data
- Kafka, Flume, ActiveMQ, ...

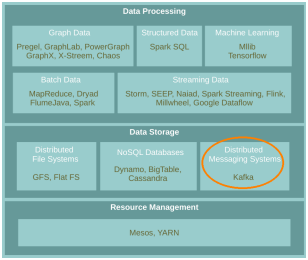

Data Processing - Batch Data

- Process data-at-rest
- Data-parallel processing model
- MapReduce, FlumeJava, Spark, ...

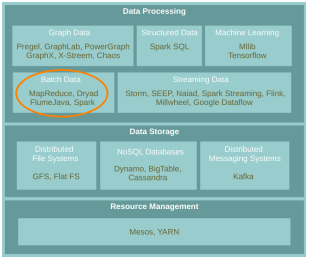

Data Processing - Streaming Data

- Process data-in-motion
- Storm, Flink, Spark Streaming, ...

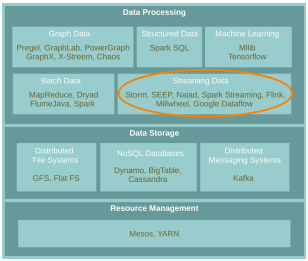

Data Processing - Linked Data (Graph)

- Graph-parallel processing model
- Vertex-centric and Edge-centric programming model
- Pregel, GraphLab, GraphX, ...

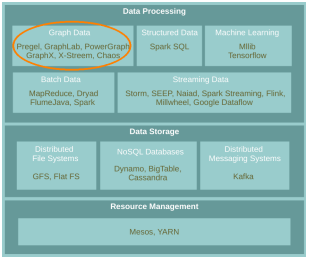

Data Processing - Structured Data

- Take advantage of schemas in data to process
- Hive, Spark SQL, ...

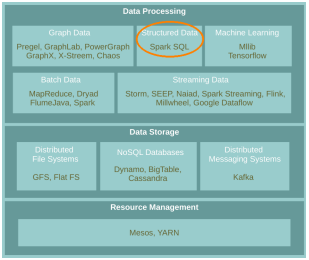

Data Processing - Machine Learning

- Data analysis, e.g., supervised and unsupervised learning
- Mahout, Tensorflow, MLlib, ...

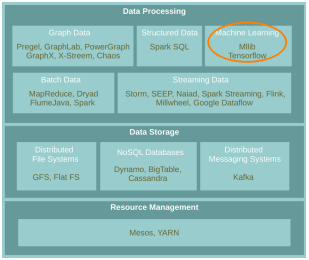

Spark Processing Engine

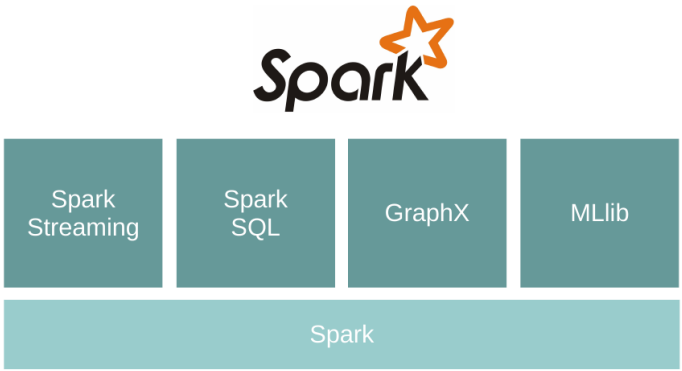

Summary

Summary

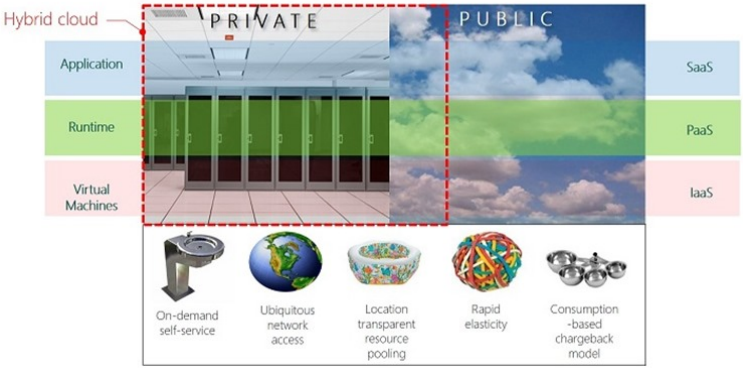

[http://aka.ms/532]

Summary

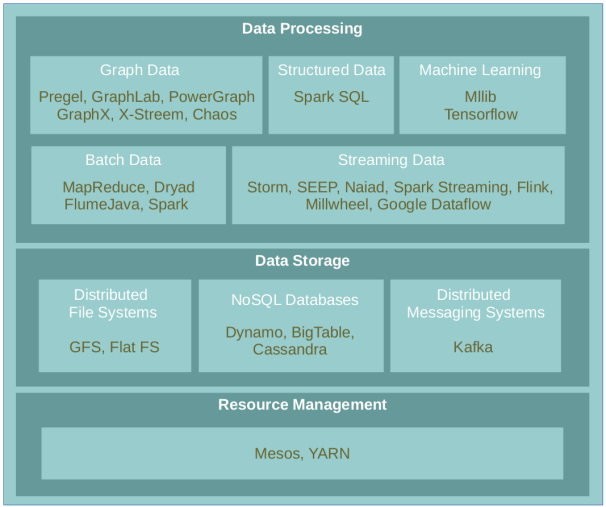
Data Processing
Graph Data
Pregel, GraphLab, PowerGraph
GraphX, X-Streem, Chaos
Structured Data
Spark SQL
Machine Learning
Mllib
Tensorflow
Batch Data
MapReduce, Dryad
FlumeJava, Spark
Streaming Data
Storm, SEEP, Naiad, Spark Streaming, Flink,
Millwheel, Google Dataflow
Data Storage
Distributed
File Systems
GFS, Flat FS
NoSQL Databases
Dynamo, BigTable,
Cassandra
Distributed
Messaging Systems
Kafka
Resource Management
Mesos, YARN

References

- D. Sikeridis et al., A Comparative Taxonomy and Survey of Public Cloud Infrastructure Vendors, arXiv preprint arXiv:1710.01476, 2017.
- A. Fox et al., Above the clouds: A berkeley view of cloud computing, UCB/EECS 28.13 (2009): 2009.
- P. Mell et al., The NIST definition of cloud computing, 2011.

Questions?