

*"...computation may someday be organized as a public utility..."*

John McCarthy, Professor at Stanford, 1961

# Lecture A.2: Large-Scale Processing on the Cloud

CS205: Computing Foundations for Computational Science  
Dr. David Sondak  
Spring Term 2020



**HARVARD**  
School of Engineering  
and Applied Sciences



**INSTITUTE FOR APPLIED  
COMPUTATIONAL SCIENCE**  
AT HARVARD UNIVERSITY

Lectures developed by Dr. Ignacio M. Illorente

# Before We Start

## Where We Are

Computing Foundations for Computational and Data Science

How to use modern computing platforms in solving scientific problems

Intro: Large-Scale Computational and Data Science

A. Parallel Processing Fundamentals

A.1. Parallel Processing Architectures

A.2. Large-scale Processing on the Cloud

A.3. Practical Aspects of Cloud Computing

A.4. Application Parallelism

A.5. Designing Parallel Programs

B. Parallel Computing

C. Parallel Data Processing

Wrap-Up: Advanced Topics

# CS205: Contents

## APPLICATION SOFTWARE

APPLICATION  
PARALLELISM

PARALLEL PROGRAM  
DESIGN



Optimization

## PROGRAMMING MODEL

OpenACC

Spark

OpenMP

Map-Reduce

MPI

B. BIG COMPUTE

PLATFORM

C. BIG DATA



CLOUD COMPUTING



Open  
Nebula



FASRC



ODYSSEY  
HARVARD FAS  
RESEARCH COMPUTING



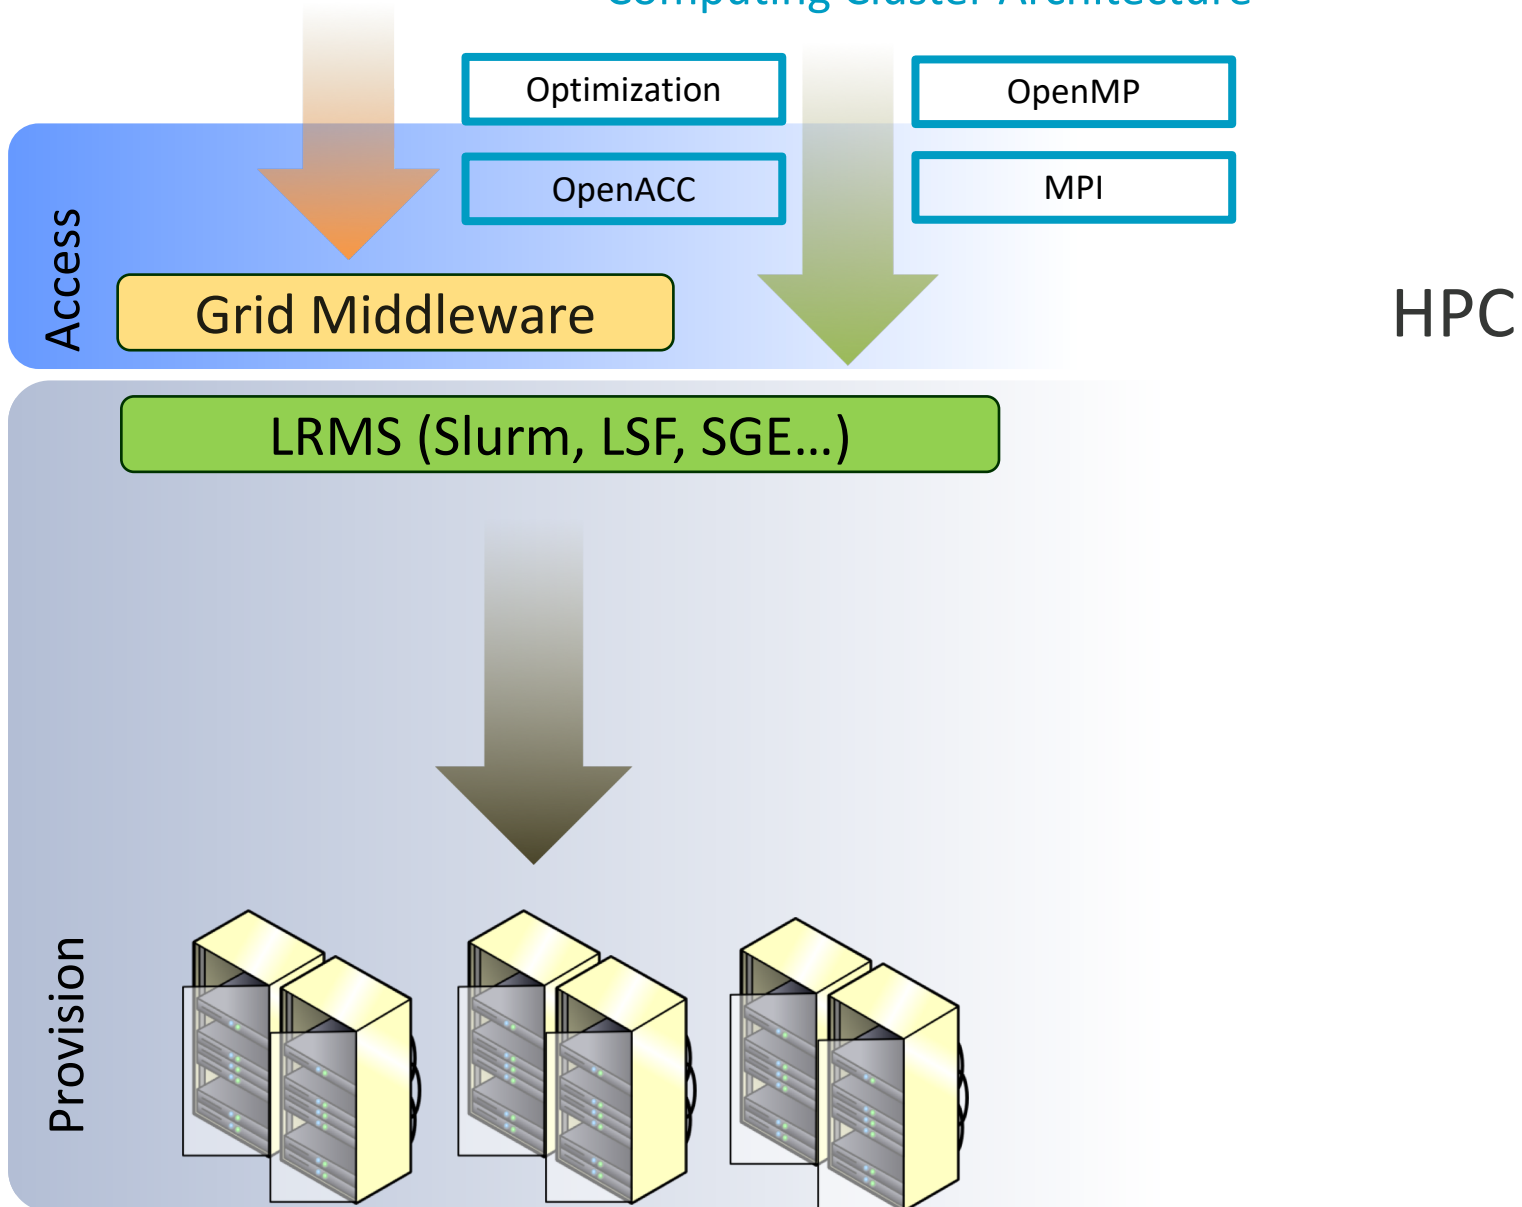
FASRC

PARALLEL ARCHITECTURES



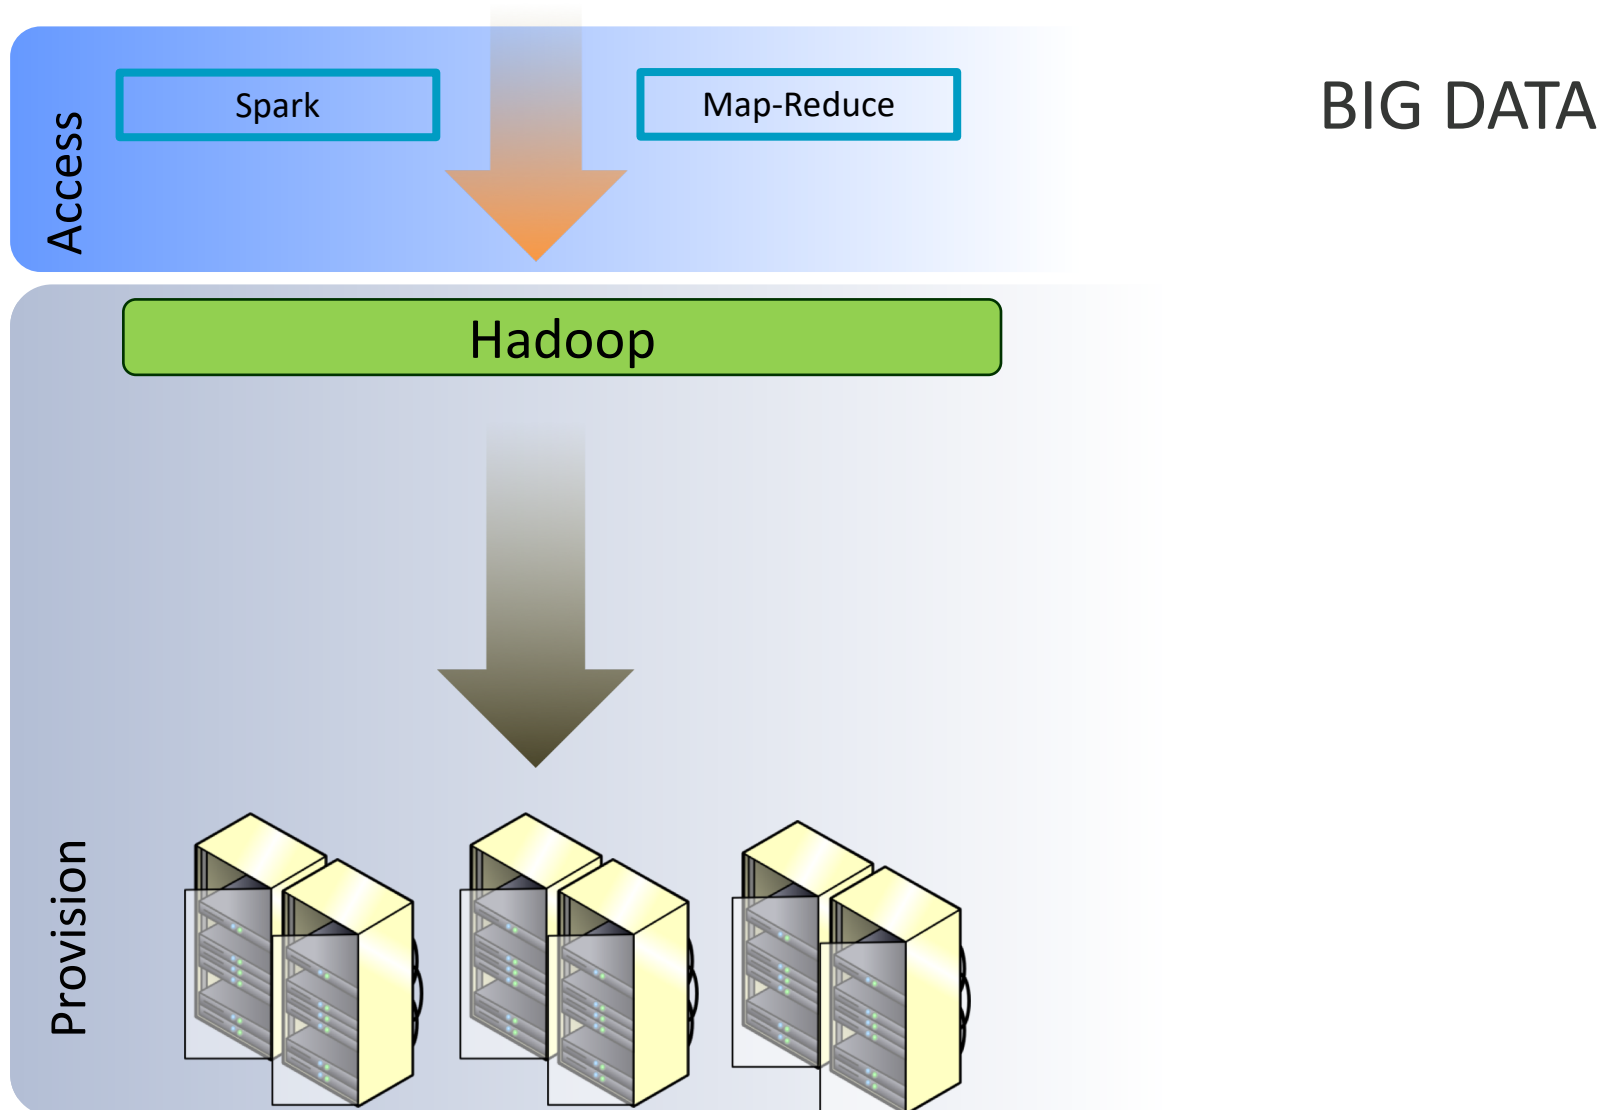
# Context

## Computing Cluster Architecture



# Context

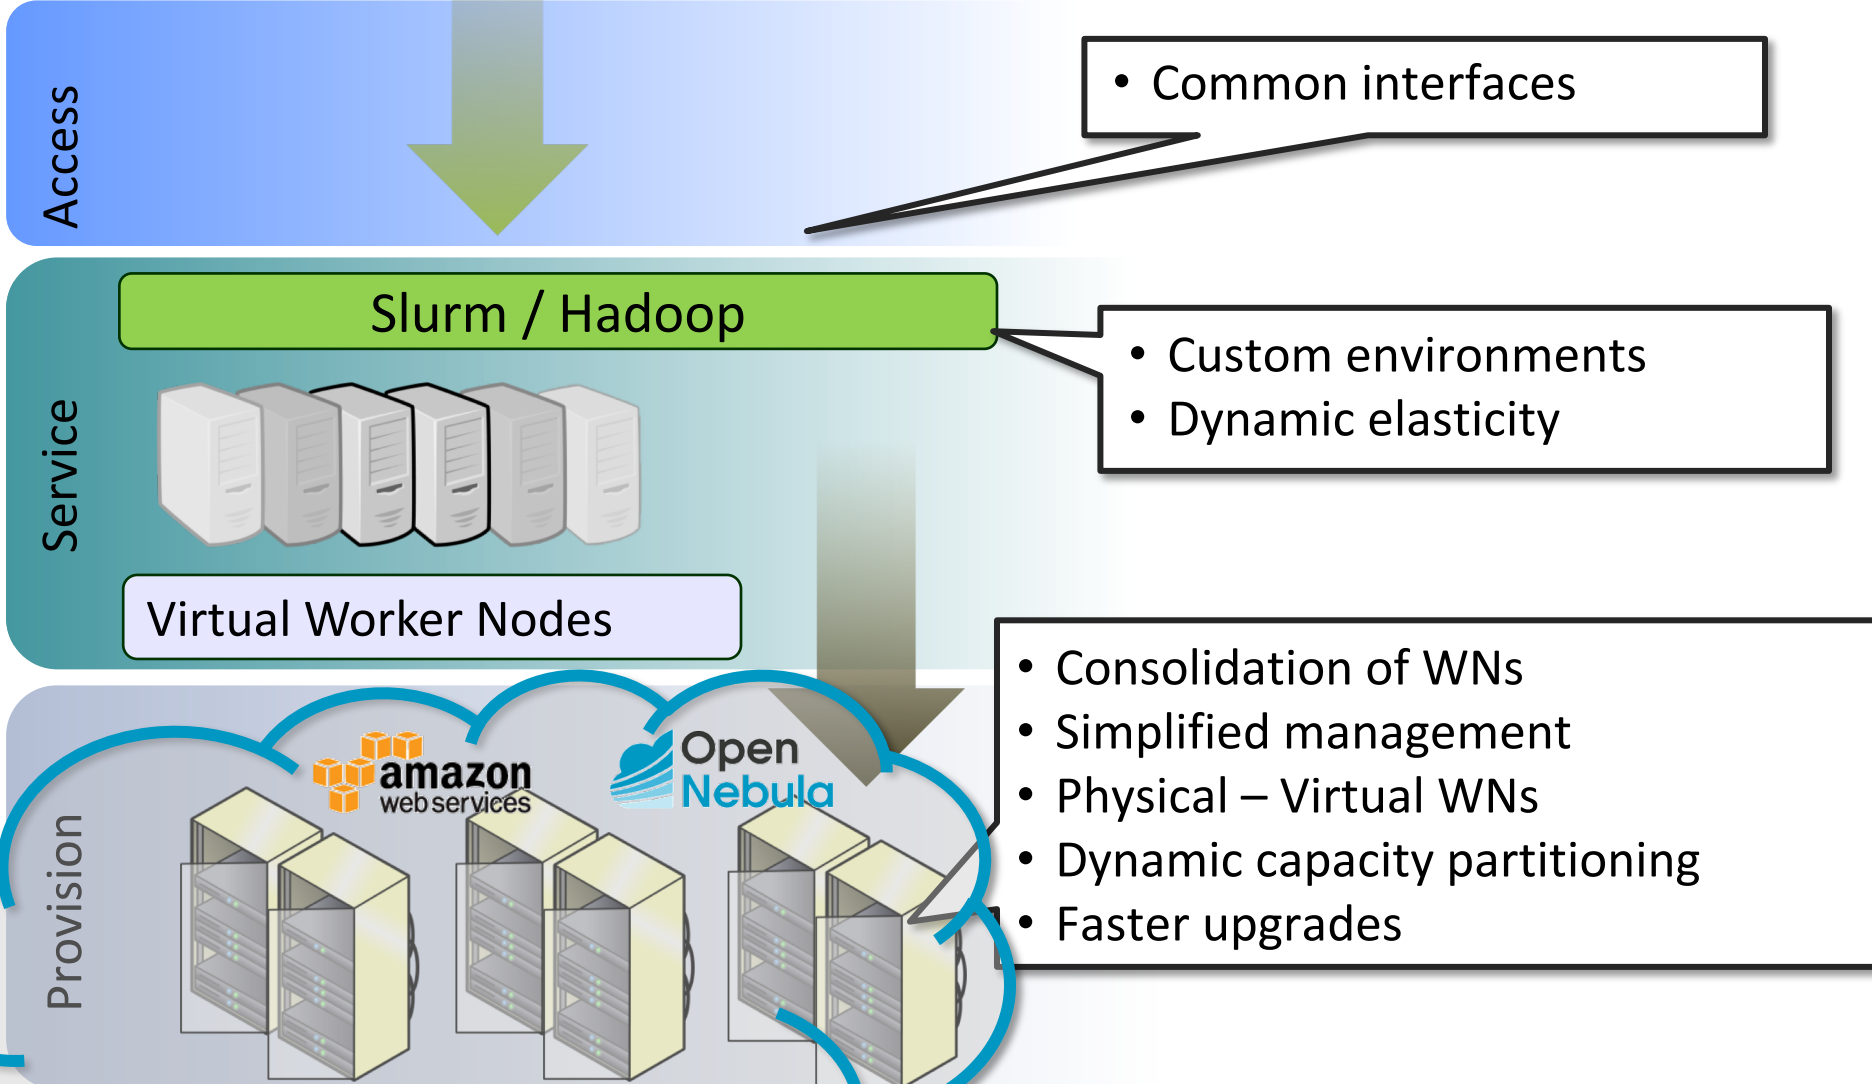
## Computing Cluster Architecture



# Context

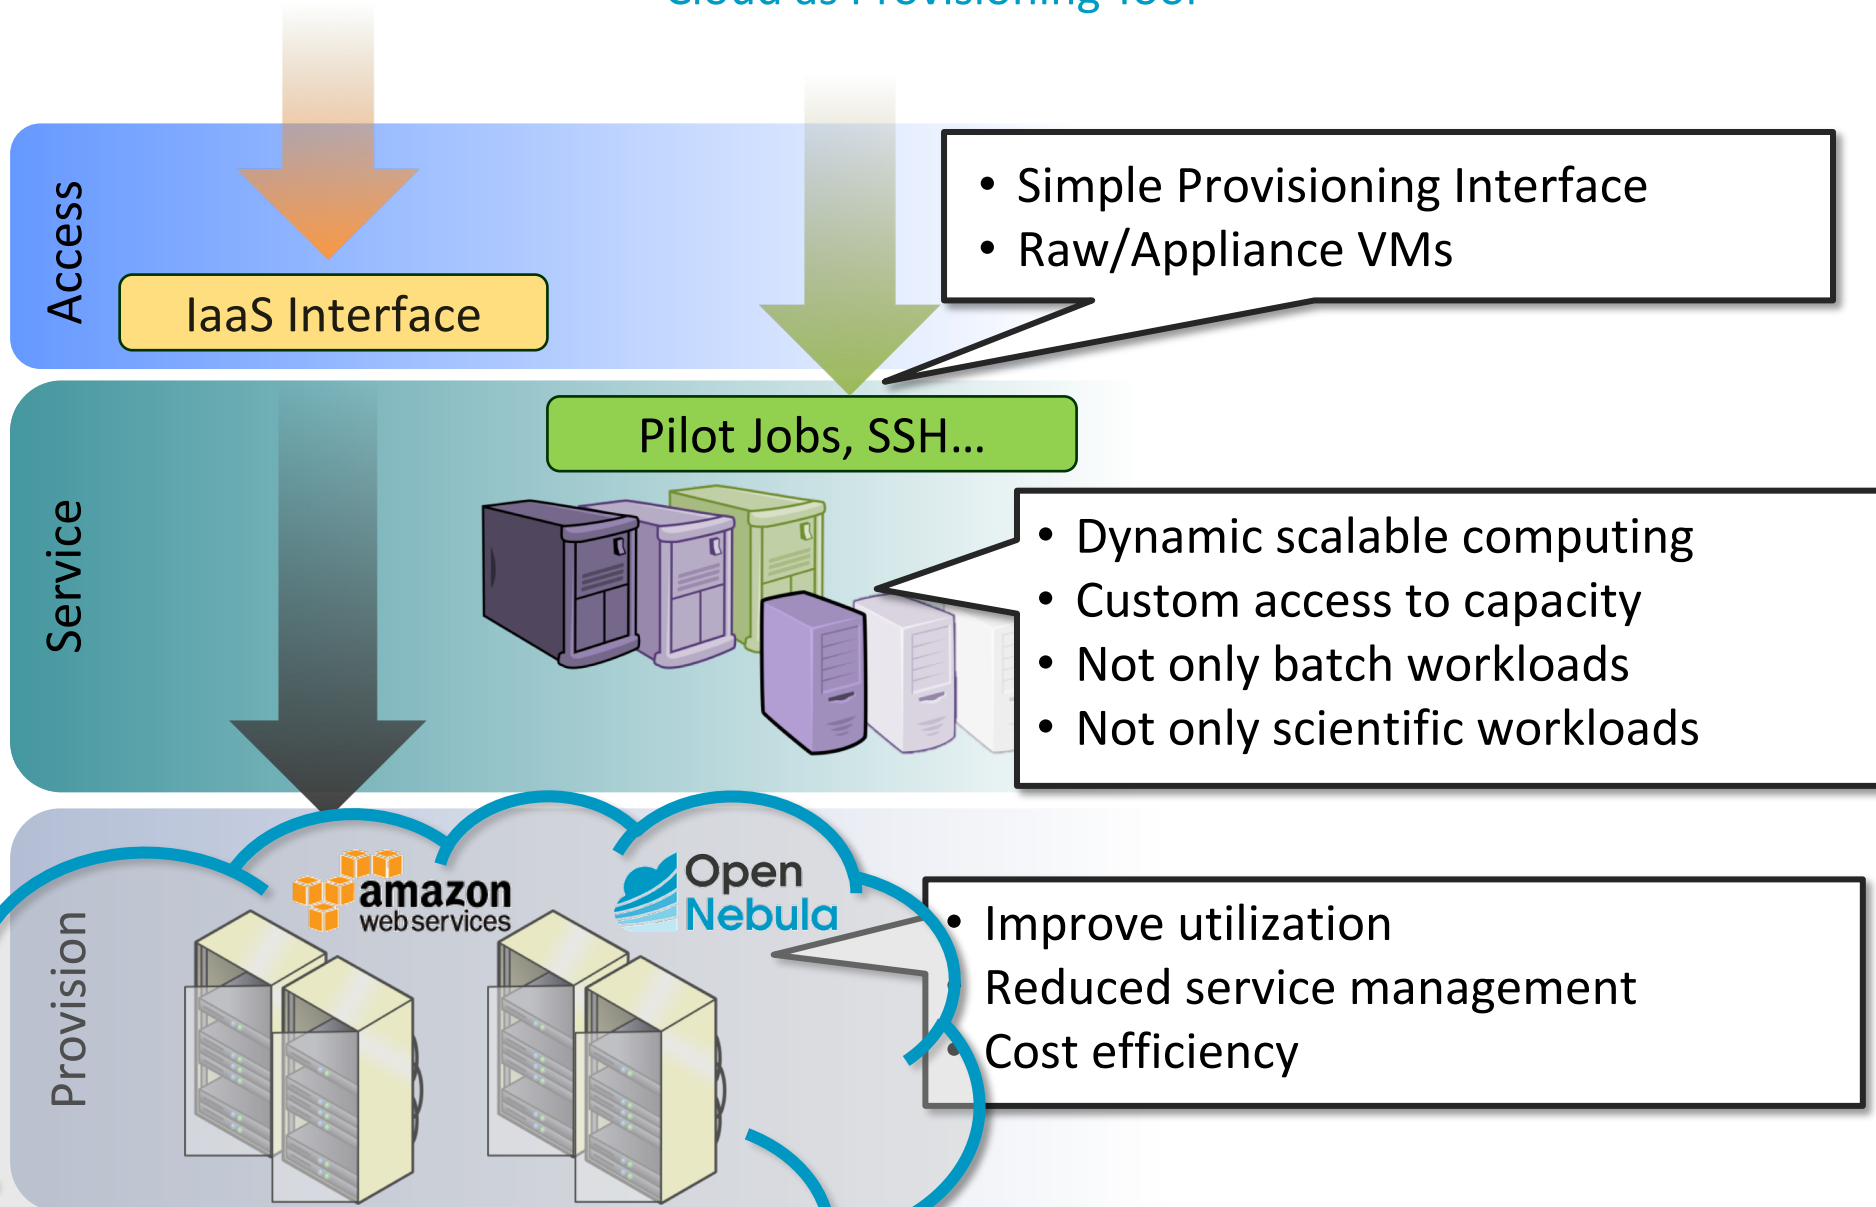
## Cloud as Infrastructure Tool

Service/Provisioning Decoupling



# Context

## Cloud as Provisioning Tool



# Roadmap

## Large-scale Data Processing on the Cloud

### Virtualization

What is Cloud Computing?

Types of Cloud Services

Cloud Services for Parallel Processing

# Virtualization

## The Key Enabler of Cloud Computing



Who has ever used virtualization?

[PollEv.com/davidsondak234](https://PollEv.com/davidsondak234),

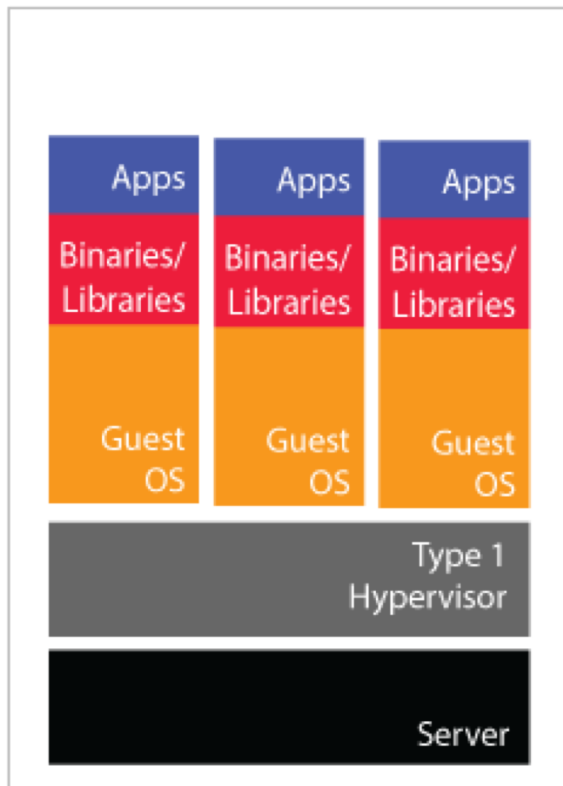
# Have you ever used virtualization?

Yes

No


# Virtualization

## The Key Enabler of Cloud Computing




Type 1 Hypervisors

VMware, Xen Project, Hyper-V



Type 2 Hypervisors

KVM, VirtualBox



OS Containers

LXC, OpenVz

Which one offers the best performance?



# Virtualization

## Pros and Cons

### End User Benefits

- Run multiple OS on the same machine
- Test/development
- Hardware/OS independent
- Replicate environments => Scientific method!

### Operator Benefits

- Backup and recovery
- Live migration and consolidation
- Automation => “Software defined hardware”
- Cost savings

### Drawbacks

- Less efficiency => HPC!

# Virtualization


## Application Containers

### Encapsulation of Application Execution Environments

- Built on top of OS containers
- Provide lightweight, portable, self-sufficient execution environments from any application.
- Same container that a developer builds and tests on a laptop can run at scale, in production, on VMs, bare metal, public clouds...

### Benefits

- Reproducibility and mobility of compute
- Ability to re-use and manage similar applications
- Content and hardware agnostic
- Easy to automate and integrate with other tools



# What is Cloud Computing?

An New Provision Model



Who has ever used a cloud?

[PollEv.com/davidsondak234](https://PollEv.com/davidsondak234)

# Have you ever used a cloud?

Yes

No

# What is Cloud Computing?

An New Provision Model



What are the benefits of the cloud?

[PollEv.com/davidsondak234](https://PollEv.com/davidsondak234)



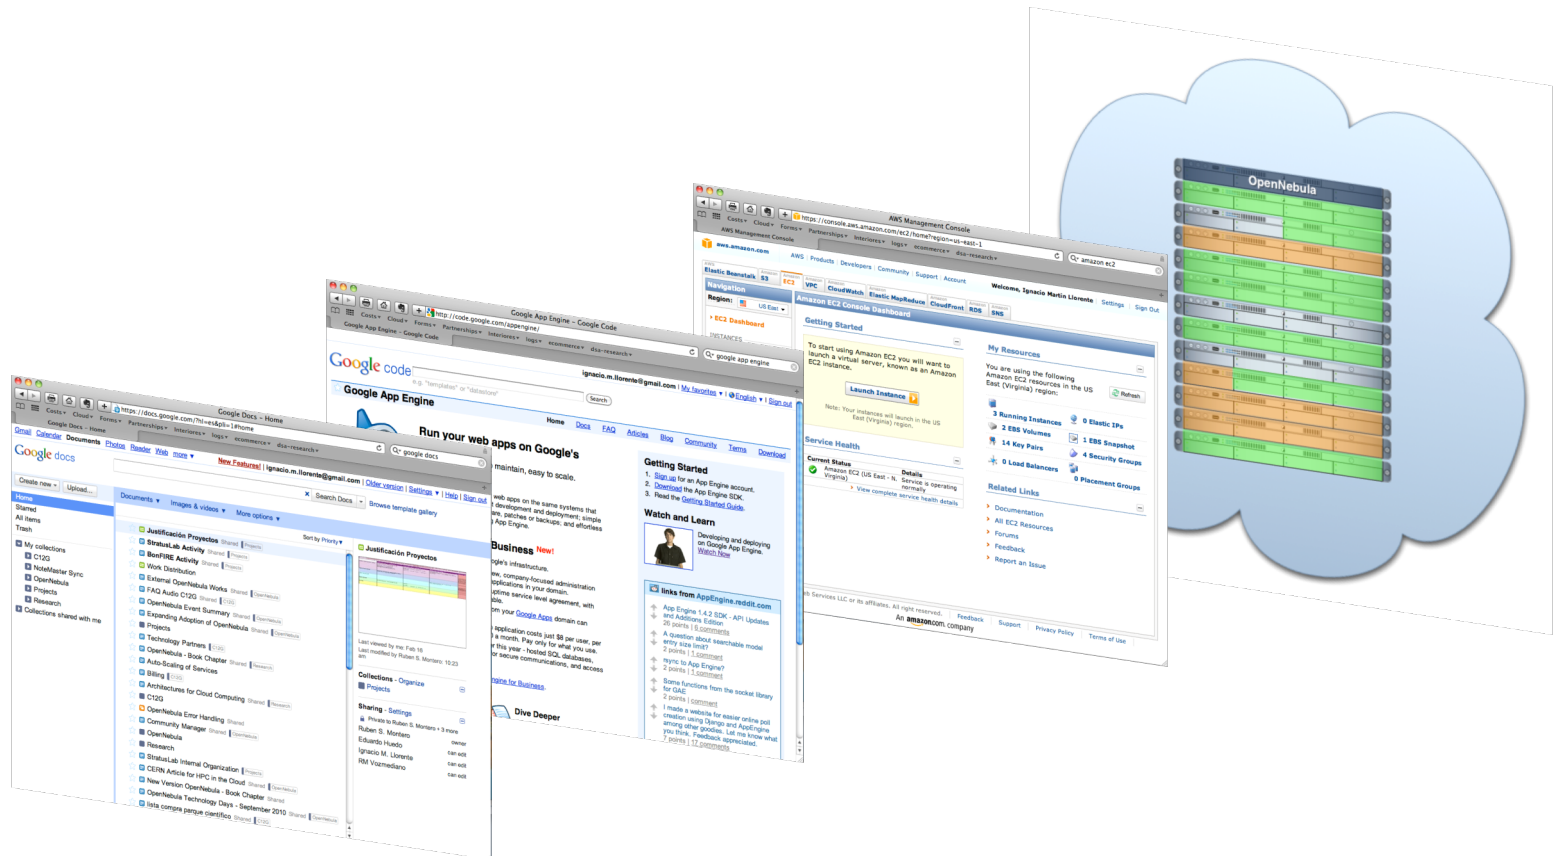
# What are the benefits of the cloud?

# What is Cloud Computing?

## An New Provision Model

### Definition as a New IT Model

- Provision of resources (applications, development platforms and infrastructure) as a service, on-demand, with elastic capacity, through the internet.




# What is Cloud Computing?

## Elasticity on the Cloud


### Vertical Scaling (Scale-Up)

Allocate additional resources to VMs, requires a reboot, no need for distributed app logic, single-point of OS failure



### Horizontal Scaling (Scale-Out)

Application needs logic to work in distributed fashion (e.g. HA-Proxy and Apache Hadoop)






# What is Cloud Computing?

A New Model, but an Old Idea


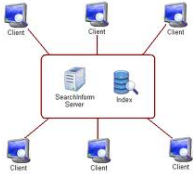

1961, John McCarthy " *...computation may someday be organized as a public utility...* "



[https://www.youtube.com/watch?v=gwsDEgiN6\\_M](https://www.youtube.com/watch?v=gwsDEgiN6_M)  
In 1974 Arthur C Clarke Predicts the Remote Worker  
and Future of the Internet and Computers

# What is Cloud Computing?

## The Third Era of Computing

	Technology	Management	Costs
<b>Mainframe</b> 	Centralized computing and storage based on high performance special-purpose systems, accessed through terminals	Optimized for maximum performance due to their high cost	Very high cost, both for hardware and software
<b>Client/Server</b> 	Distributed computing and storage based on commodity workstations and personal computers	Optimized for maximum agility due to their lower cost	Lower, incremental cost for hardware, and a perpetual license-based model for software
<b>Cloud</b> 	Centralized computing and storage based on commodity systems, accessed through devices	Optimized for maximum performance and agility	Pay per use

# What is Cloud Computing?

## The Future of Resource Provisioning

*“The horse is here to stay but the automobile is only a novelty, a fad”*

President of Morgan Saving Bank to Ford Motor Co. Lawyers,  
1903



# What is Cloud Computing?

## Different Perspectives

### Use Model

- Focus on the result and not its implementation

### Access Model

- Ubiquitous access from any device



### Cost Model

- From Capital expenditures (CAPEX) to operating expenses (OPEX)

### Infrastructure Model

- Virtualized, elastic, on-demand, multi-tenant IT capacity

# What is Cloud Computing?

## Different Perspectives

### Security

- Backup and HA provided by the provider

### Ubiquity

- Access from everywhere and any device

### Coherency

- No need for information synchronization

### Pay per Use

- No licenses, only pay for the time you need it

### Updates

- Software always patched and updated

### Lite Clients

- Does not require powerful client devices

### Collaborative Work








- Natural support for remote collaboration

### Energy Saving

- Allows consolidation of resources

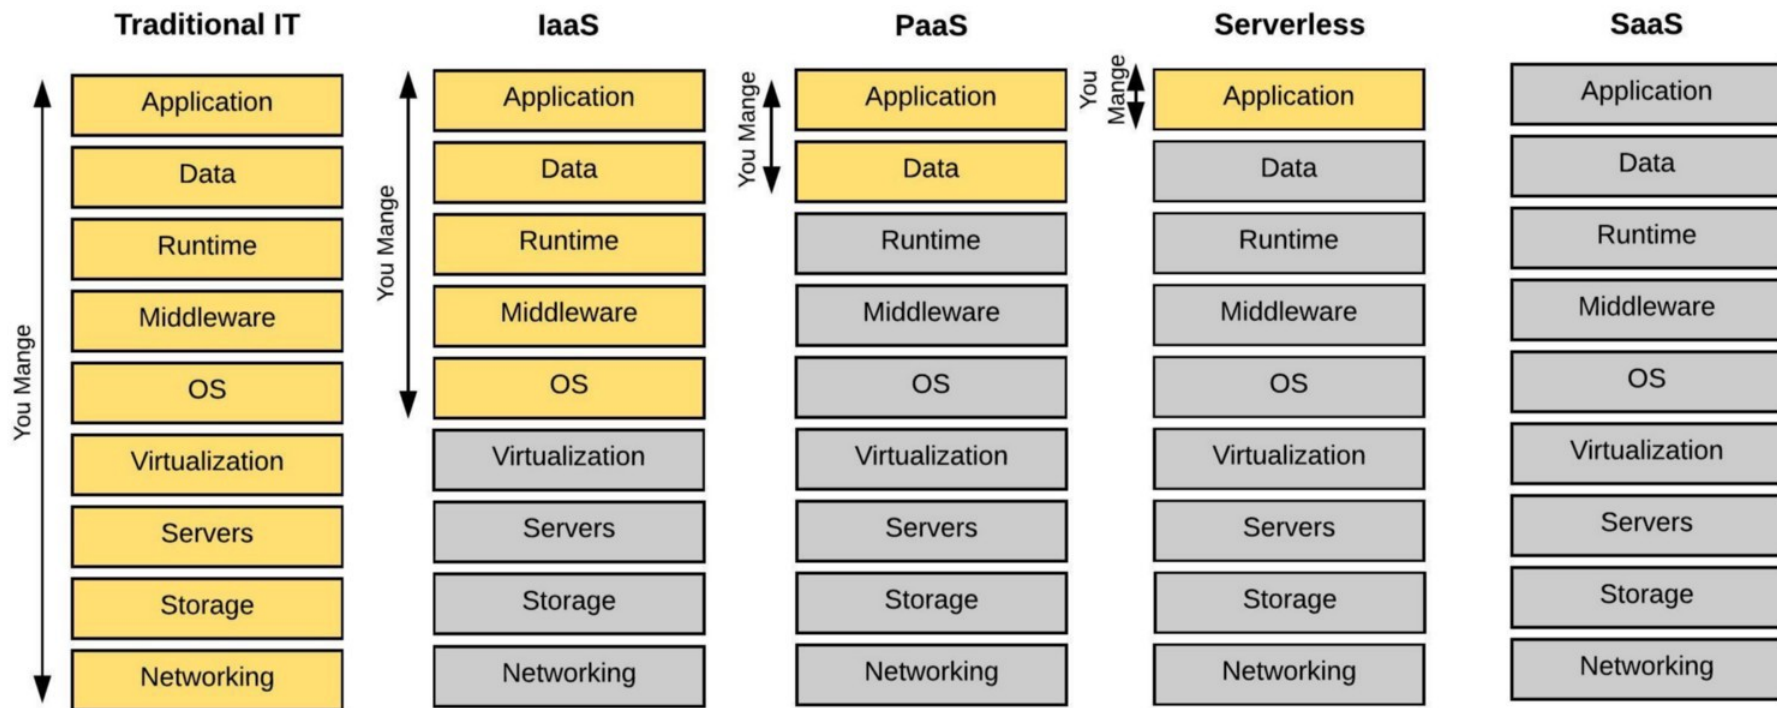
# Types of Cloud Services

## From Infrastructure to Application

	What	Who
<b>Software as a Service</b>	On-demand access to an application	Final user (not interested in sw/hw)
		 
<b>Platform as a Service</b>	On-demand access to a platform to develop and test applications	Developer (not interested in hw)
		 
<b>Infrastructure as a Service</b>	IT resources	System administrators (complete management of the infrastructure)
<b>Physical Infrastructure</b>		  

# Types of Cloud Services

## From Infrastructure to Application






# Types of Cloud Services

## Infrastructure as a Service

### On-demand Access to IT Resources

- Simple REST interface
- Pay per use
- Elastic capacity

\$0.006 (1 vCPU 0.5 GiB) - \$24.48 (64 vCPU 48 GiB) / Hour





# Types of Cloud Services

## Infrastructure as a Service



### On-Demand

With On-Demand instances, you pay for compute capacity by the hour with no long-term commitments or upfront payments. You can increase or decrease your compute capacity depending on the demands of your application and only pay the specified hourly rate for the instances you use.

On-Demand instances are recommended for:

- Users that prefer the low cost and flexibility of Amazon EC2 without any up-front payment or long-term commitment
- Applications with short-term, spiky, or unpredictable workloads that cannot be interrupted
- Applications being developed or tested on Amazon EC2 for the first time

[See On-Demand Pricing](#)

### Spot Instances

Amazon EC2 Spot instances allow you to bid on spare Amazon EC2 computing capacity for up to 90% off the On-Demand price.

[Learn More.](#)

Spot instances are recommended for:

- Applications that have flexible start and end times
- Applications that are only feasible at very low compute prices
- Users with urgent computing needs for large amounts of additional capacity

[See Spot Pricing](#)

### Reserved Instances

Reserved Instances provide you with a significant discount (up to 75%) compared to On-Demand instance pricing. In addition, when Reserved Instances are assigned to a specific Availability Zone, they provide a capacity reservation, giving you additional confidence in your ability to launch instances when you need them.

For applications that have steady state or predictable usage, Reserved Instances can provide significant savings compared to using On-Demand instances. See [How to Purchase Reserved Instances](#) for more information.

Reserved Instances are recommended for:

- Applications with steady state usage
- Applications that may require reserved capacity
- Customers that can commit to using EC2 over a 1 or 3 year term to reduce their total computing costs

### Dedicated Hosts

A Dedicated Host is a physical EC2 server dedicated for your use. Dedicated Hosts can help you reduce costs by allowing you to use your existing server-bound software licenses, including Windows Server, SQL Server, and SUSE Linux Enterprise Server (subject to your license terms), and can also help you meet compliance requirements. [Learn more.](#)

- Can be purchased On-Demand (hourly).
- Can be purchased as a Reservation for up to 70% off the On-Demand price.

# Types of Cloud Services


## Infrastructure as a Service



	vCPU	ECU	Memory (GiB)	Instance Storage (GB)	Linux/UNIX Usage
<b>General Purpose - Current Generation</b>					
t2.nano	1	Variable	0.5	EBS Only	\$0.0059 per Hour
t2.micro	1	Variable	1	EBS Only	\$0.012 per Hour
t2.small	1	Variable	2	EBS Only	\$0.023 per Hour
t2.medium	2	Variable	4	EBS Only	\$0.047 per Hour
t2.large	2	Variable	8	EBS Only	\$0.094 per Hour
t2.xlarge	4	Variable	16	EBS Only	\$0.188 per Hour
t2.2xlarge	8	Variable	32	EBS Only	\$0.376 per Hour
m4.large	2	6.5	8	EBS Only	\$0.1 per Hour
m4.xlarge	4	13	16	EBS Only	\$0.2 per Hour
m4.2xlarge	8	26	32	EBS Only	\$0.4 per Hour
m4.4xlarge	16	53.5	64	EBS Only	\$0.8 per Hour
m4.10xlarge	40	124.5	160	EBS Only	\$2 per Hour
m4.16xlarge	64	188	256	EBS Only	\$3.2 per Hour

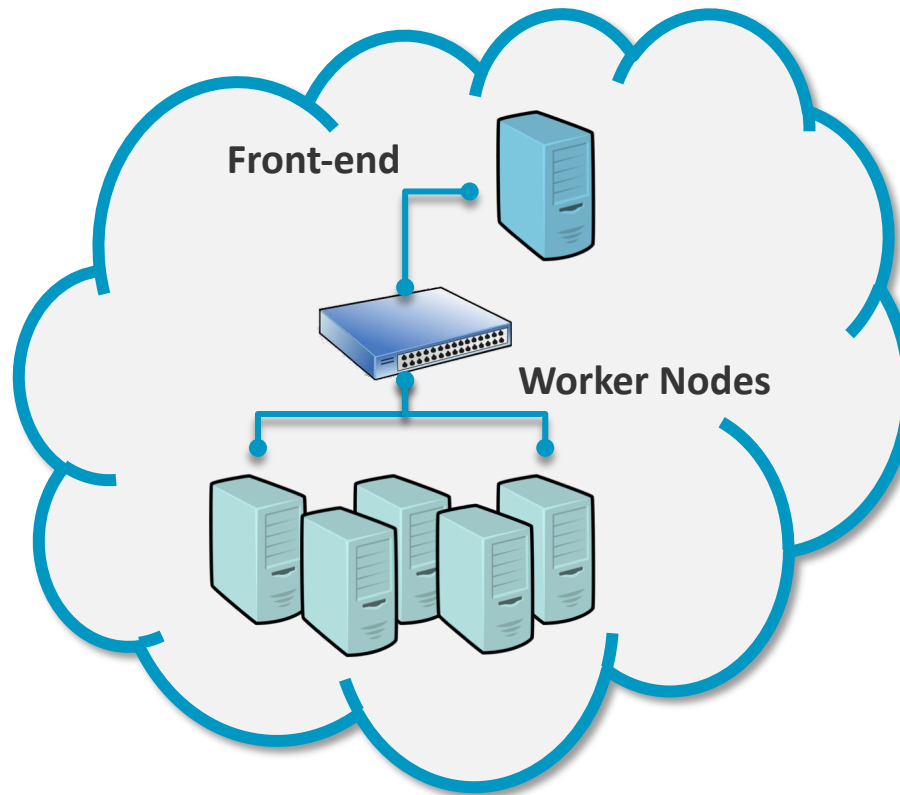
# Types of Cloud Services

## Infrastructure as a Service



# Cloud Services for Parallel Processing

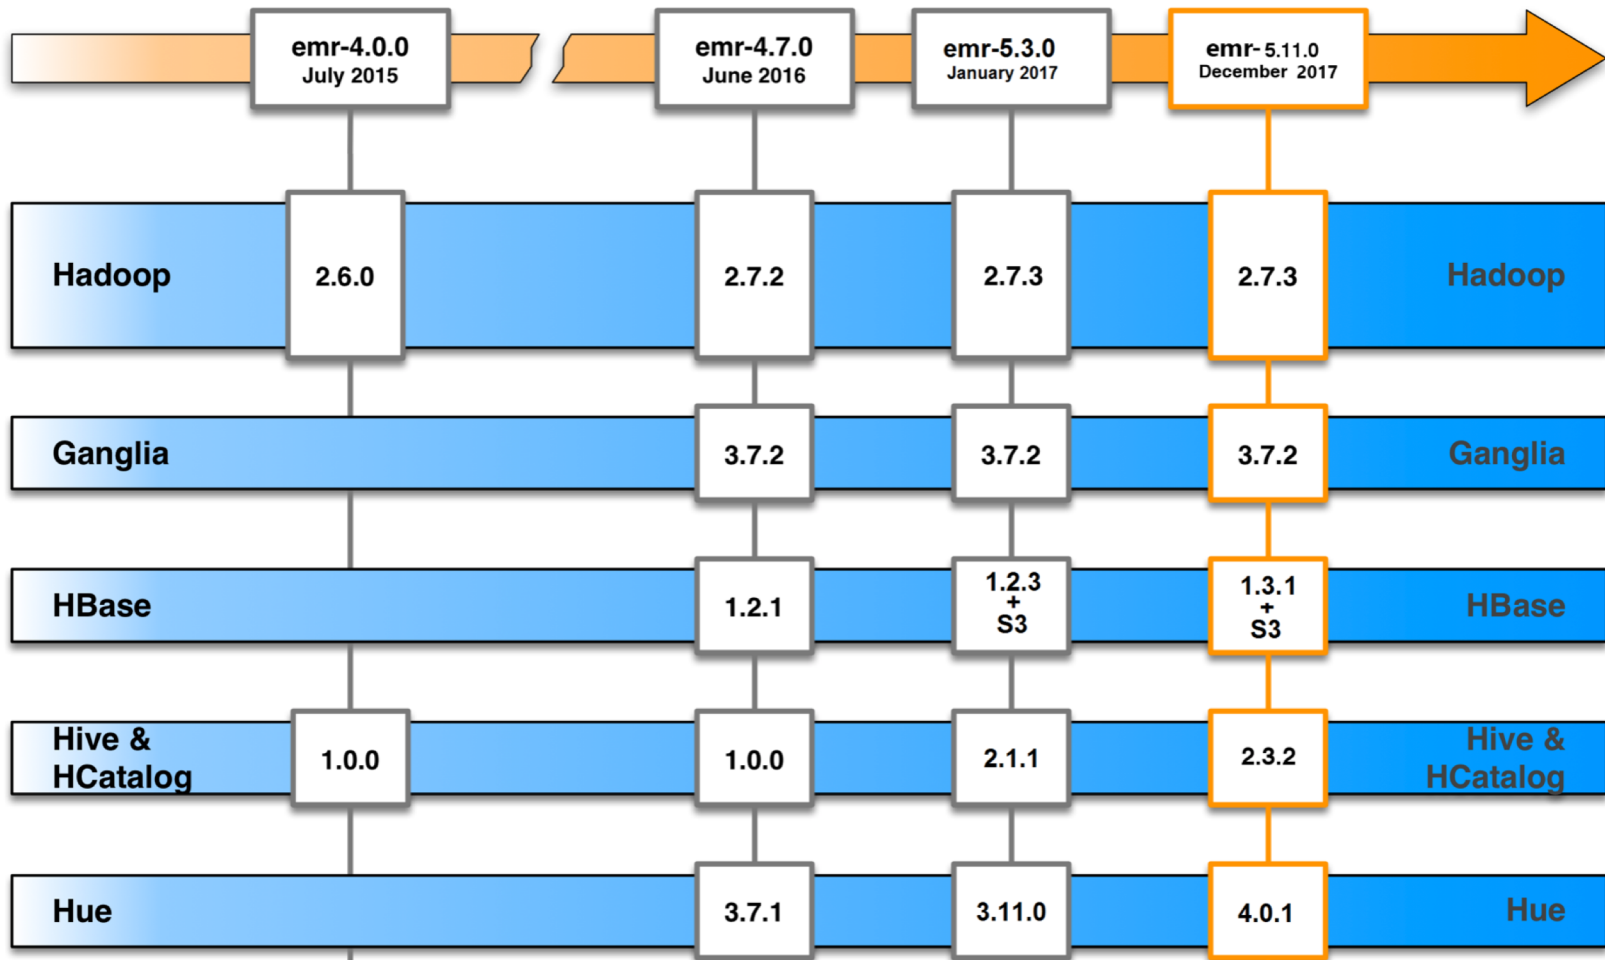
## Cloud as Provisioning Tool: AWS - DIY



# Cloud Services for Parallel Processing

## Cloud as Infrastructure Tool: AWS - EMR

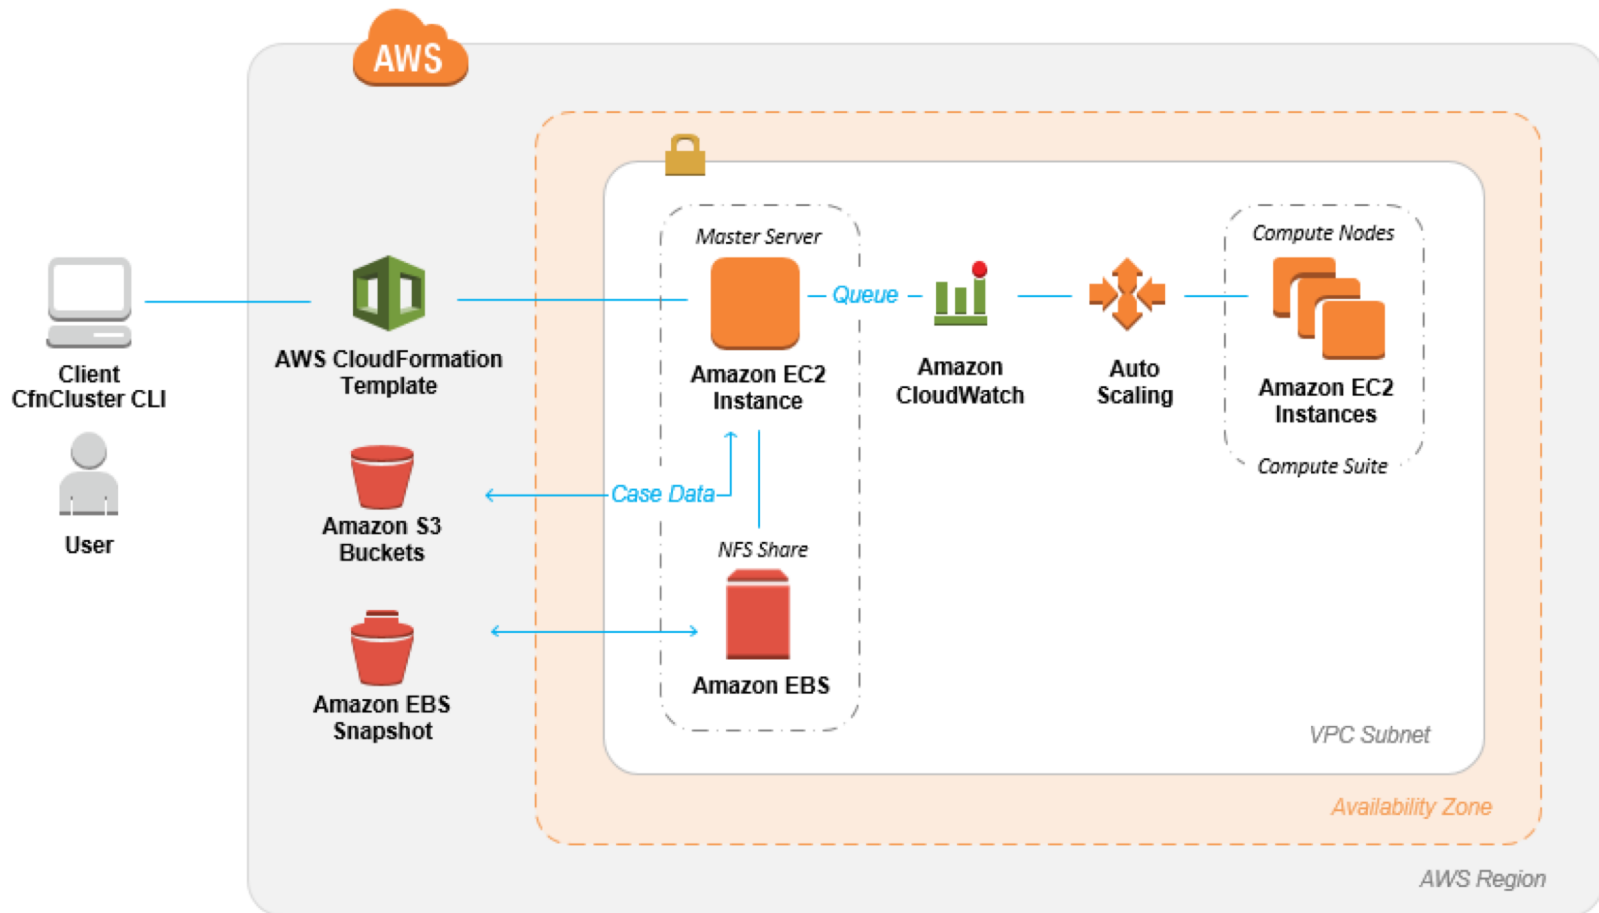
### EMR Releases



# Cloud Services for Parallel Processing

## Cloud as Infrastructure Tool: AWS – HPC

- Enhanced networking (20 Gigabit)
- Different high performance storage solution to build the cluster
- Managed by Engine Frame (CfnCluster)

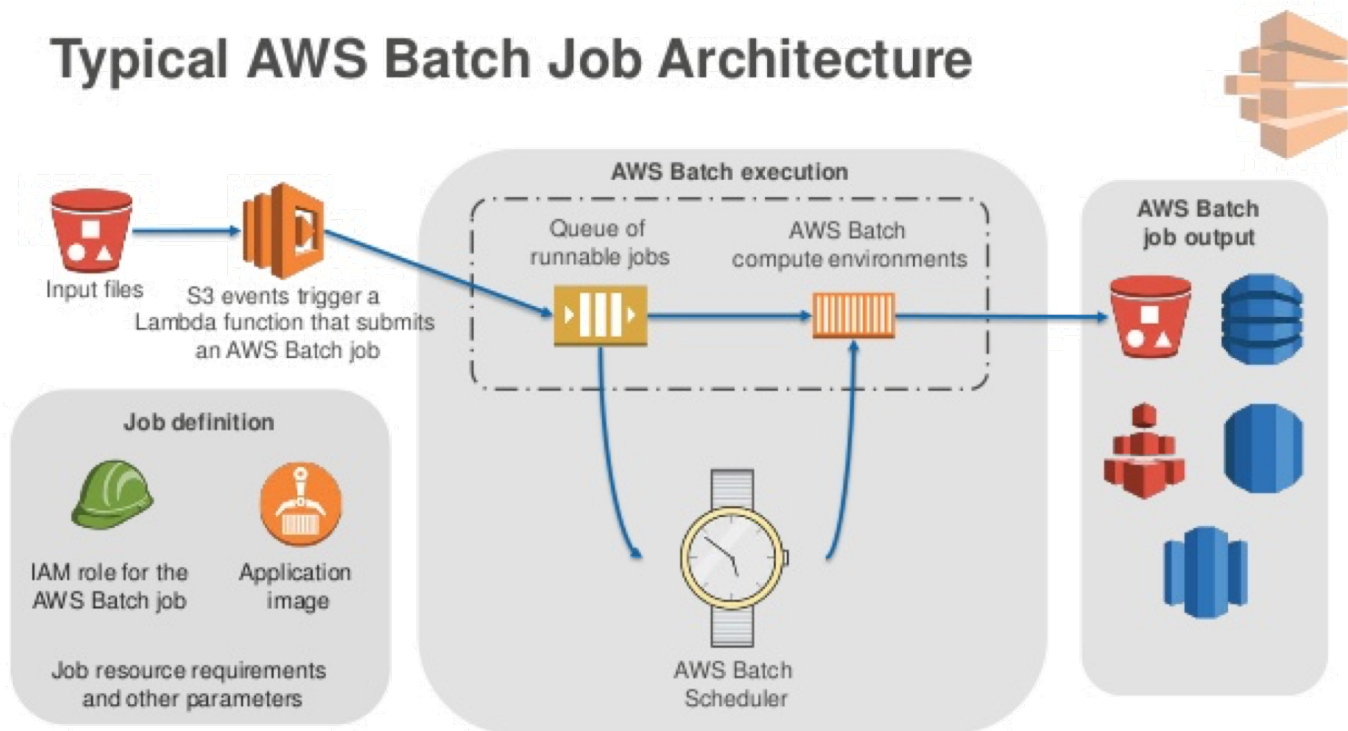


# Cloud Services for Parallel Processing

## Cloud as Infrastructure Tool: AWS - Batch

- Fully-managed batch system for HTC (including DAG-based workflows)
- “Containerized” jobs (unit of execution), queues and environments
- Job definition includes container properties, capacity req., variables...
- Jobs reside in queues until they are executed (priorities)
- Jobs are mapped into computing environments (EC2 instances)
- Basic scheduler and simple job life-cycle

### Typical AWS Batch Job Architecture



# Cloud Services for Parallel Processing

## Performance Penalty as a Small Tax You Have to Pay

### Overhead in Computing

- Single node has processor performance penalty between 1% and 5%
- NASA reported an overhead between 9% and 25% (HPCC and NPB)<sup>1</sup>
- Growing number of clouds using containers (OpenVZ and LXC)

### Overhead in Input/Output

- Growing number of Big Data apps
- Support for multiple system datastores including automatic scheduling

### Need for Low-Latency High-Bandwidth Interconnection

- Lower performance, 10 GigE typically, used in clouds has a significant negative (x2-x10, especially latency) impact on HPC applications<sup>1</sup>
- PCI passthrough available for VMs that need consumption of raw GPU devices and Infiniband access
- FermiCloud has reported MPI performance (HPL benchmark) on VMs and SR-IOV/Infiniband with only a 4% overhead<sup>2</sup>

- (1) An Application-Based Performance Evaluation of Cloud Computing, NASA Ames, 2013
- (2) FermiCloud Update, Keith Chadwick, Fermilab, HePIX Spring Workshop 2013



# Reading Assignments / Open Discussion

## Potential of Cloud for Scientific Applications

I. Sadooghi et al, *“Understanding the Performance and Potential of Cloud Computing for Scientific Applications”*,  
IEEE TCC, Issue No. 02 - April-June (2017 vol. 5)

What is the aim of the study?

Which infrastructures are compared?

What is the methodology?

What is the overhead in the memory, network, compute and storage components?

# Next Steps

- **Lab session** tomorrow:  
Help with HWA or Infrastructure Guide  
“First Access to AWS”?
- Get ready for **next lecture**:  
A.3. Practical Aspects of Cloud Computing
- **Reading Assignment**:

J. Riley, J. Noss, W. Dillingham, J. Cuff, I. M. Llorente, “*A High-Availability Cloud for Research Computing*”, IEEE Computer, Volume: 50, Issue: 6, 2017

# Questions

Large-Scale Parallel Processing on the Cloud

