



Queen's
UNIVERSITY



Shadi Khalifa
Database Systems Laboratory (DSL)
khalifa@cs.queensu.ca

what is Amazon !!

- American international multibillion dollar electronic commerce company with headquarters in Seattle, Washington, USA.
 - started in 1995 by Jeff Bezos as an online bookstore.
 - but soon diversified, selling DVDs, VHSs, CDs, video and MP3 downloads/streaming, software, video games, electronics, apparel, furniture, food, toys, and jewelry.
 - The company also produces consumer electronics: Kindle e-book reader and the Kindle Fire tablet computer.
 - In 2006, Amazon officially launched the Amazon web Services (AWS) to become a major provider of cloud computing services.

What is Amazon Web Services?

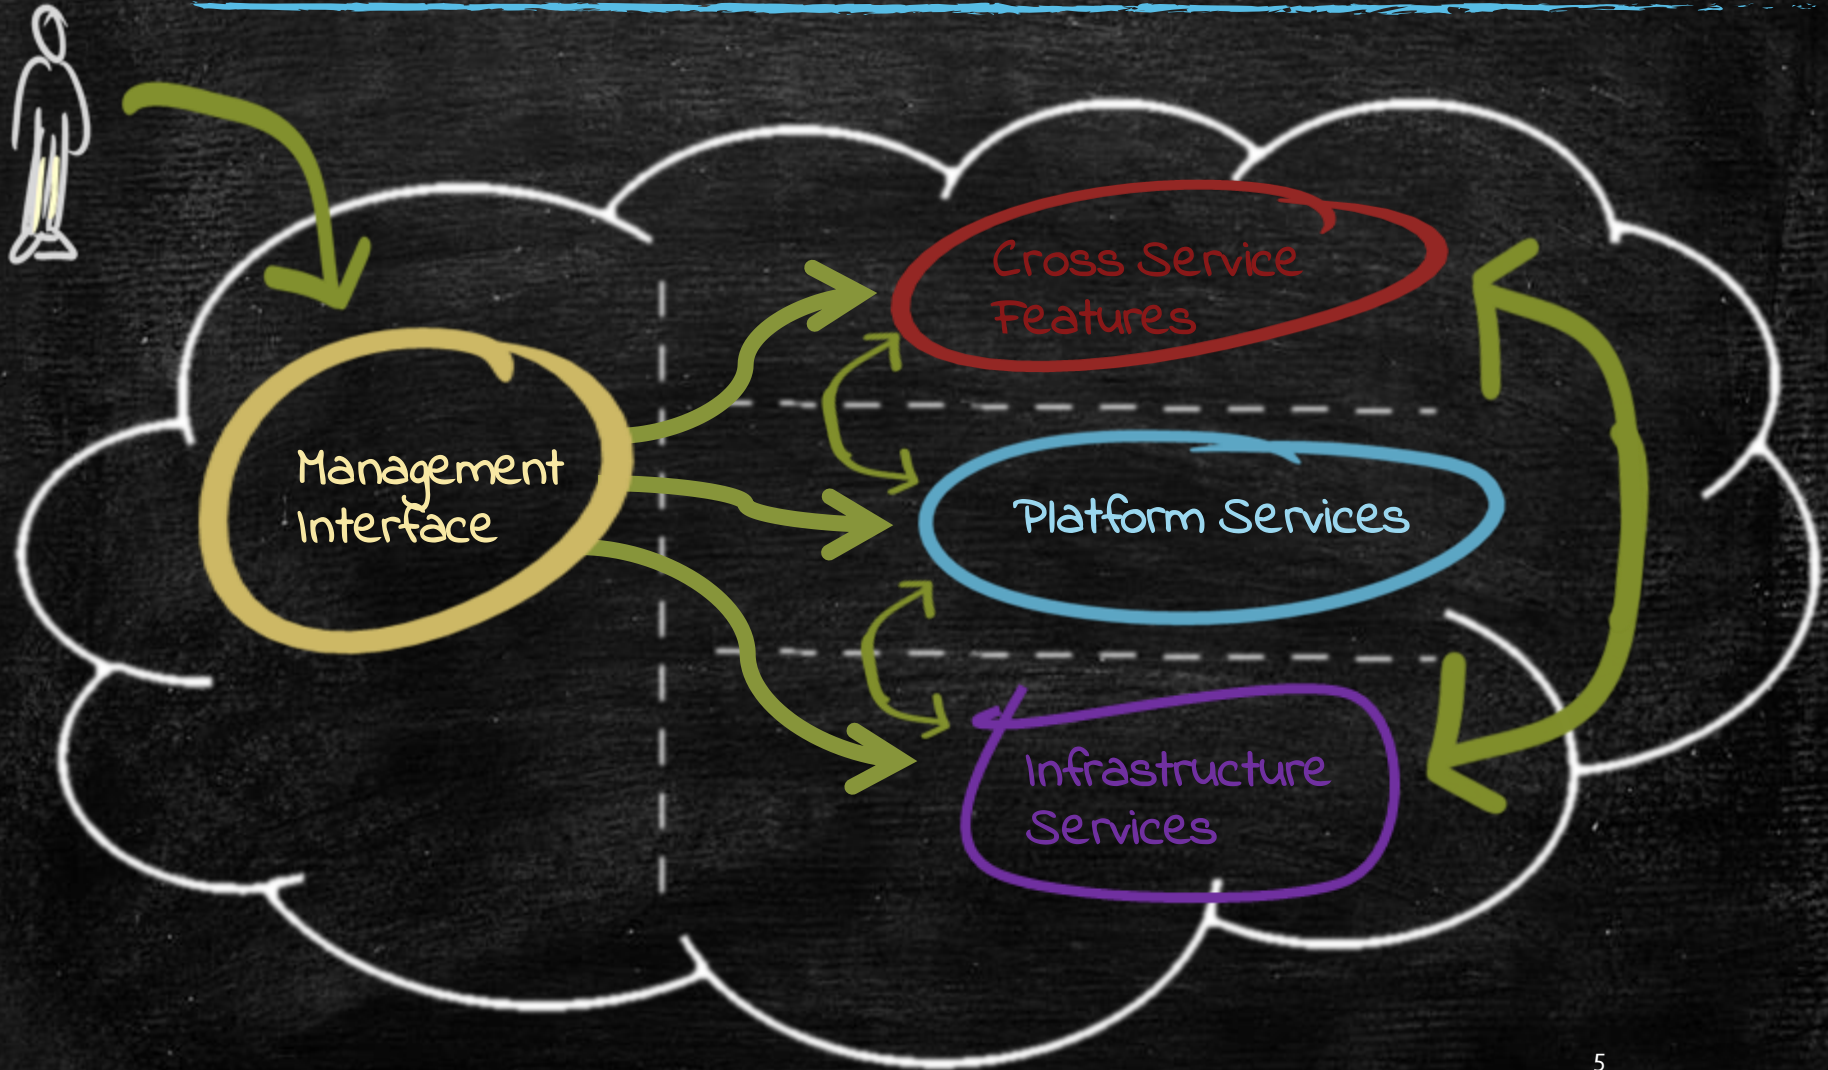
- Amazon Web Services (AWS) is a collection of remote computing services (web services) that together make up a cloud computing platform, offered over the Internet by Amazon.com.
- website: <http://aws.amazon.com>
- AWS is located in 9 geographical 'Regions'. Each Region is wholly contained within a single country and all of its data and services stay within the designated Region.
- Each Region has multiple 'Availability Zones', which are distinct data centers providing AWS services. Availability Zones are isolated from each other to prevent outages from spreading between zones. However, several services operate across Availability Zones (e.g. S3, DynamoDB).



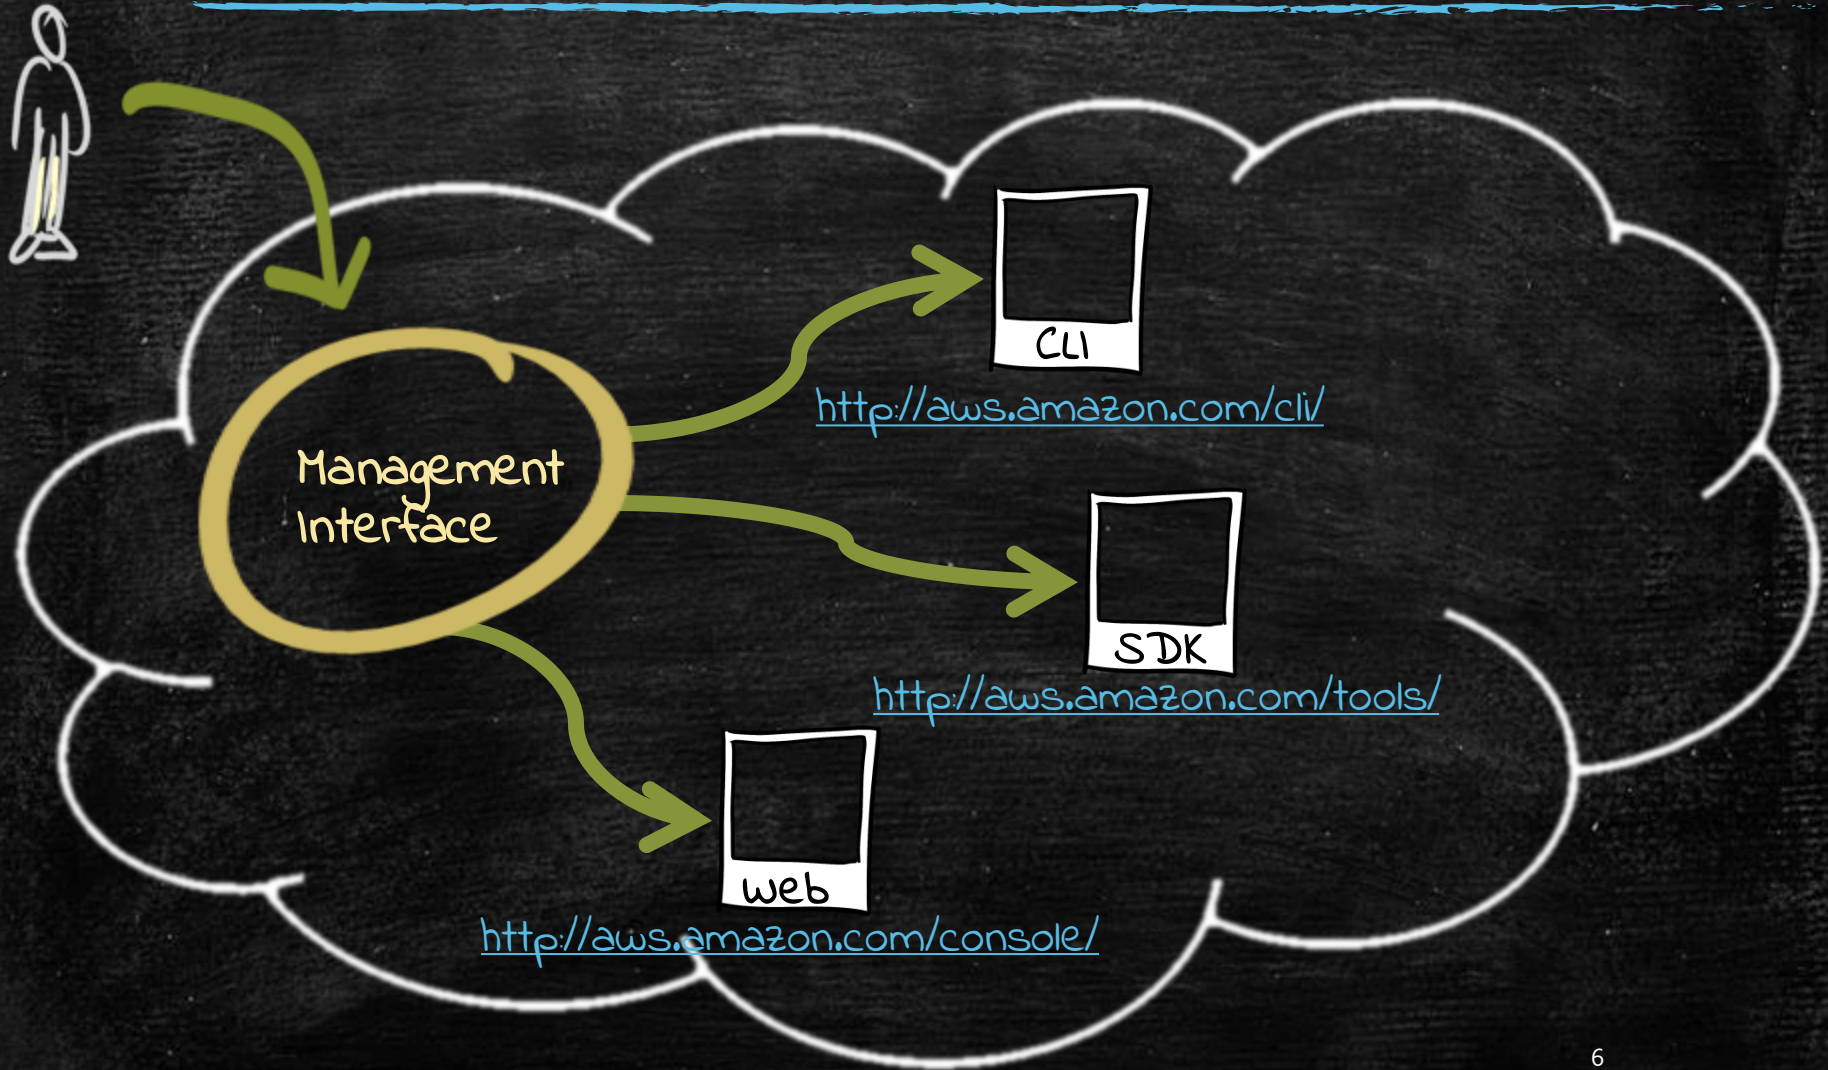
What is AWS offering?

- **Low Ongoing Cost;** pay-as-you-go pricing with no up-front expenses or long-term commitments.
- **Instant Elasticity & Flexible Capacity: (scaling up and down)** Eliminate guessing on your infrastructure capacity needs.
- **Speed & Agility:** Develop and deploy applications faster instead of waiting weeks or months for hardware to arrive and get installed.
- **Apps not Ops:** Focus on projects. Lets you shift resources away from data center investments and operations and move them to innovative new projects.
- **Global Reach:** Take your apps global in minutes.
- **Open and Flexible:** You choose the development platform or programming model that makes the most sense for your business.
- **Secure:** Allows your application to take advantage of the multiple layers of operational and physical security in the AWS data centers to ensure the integrity and safety of your data.

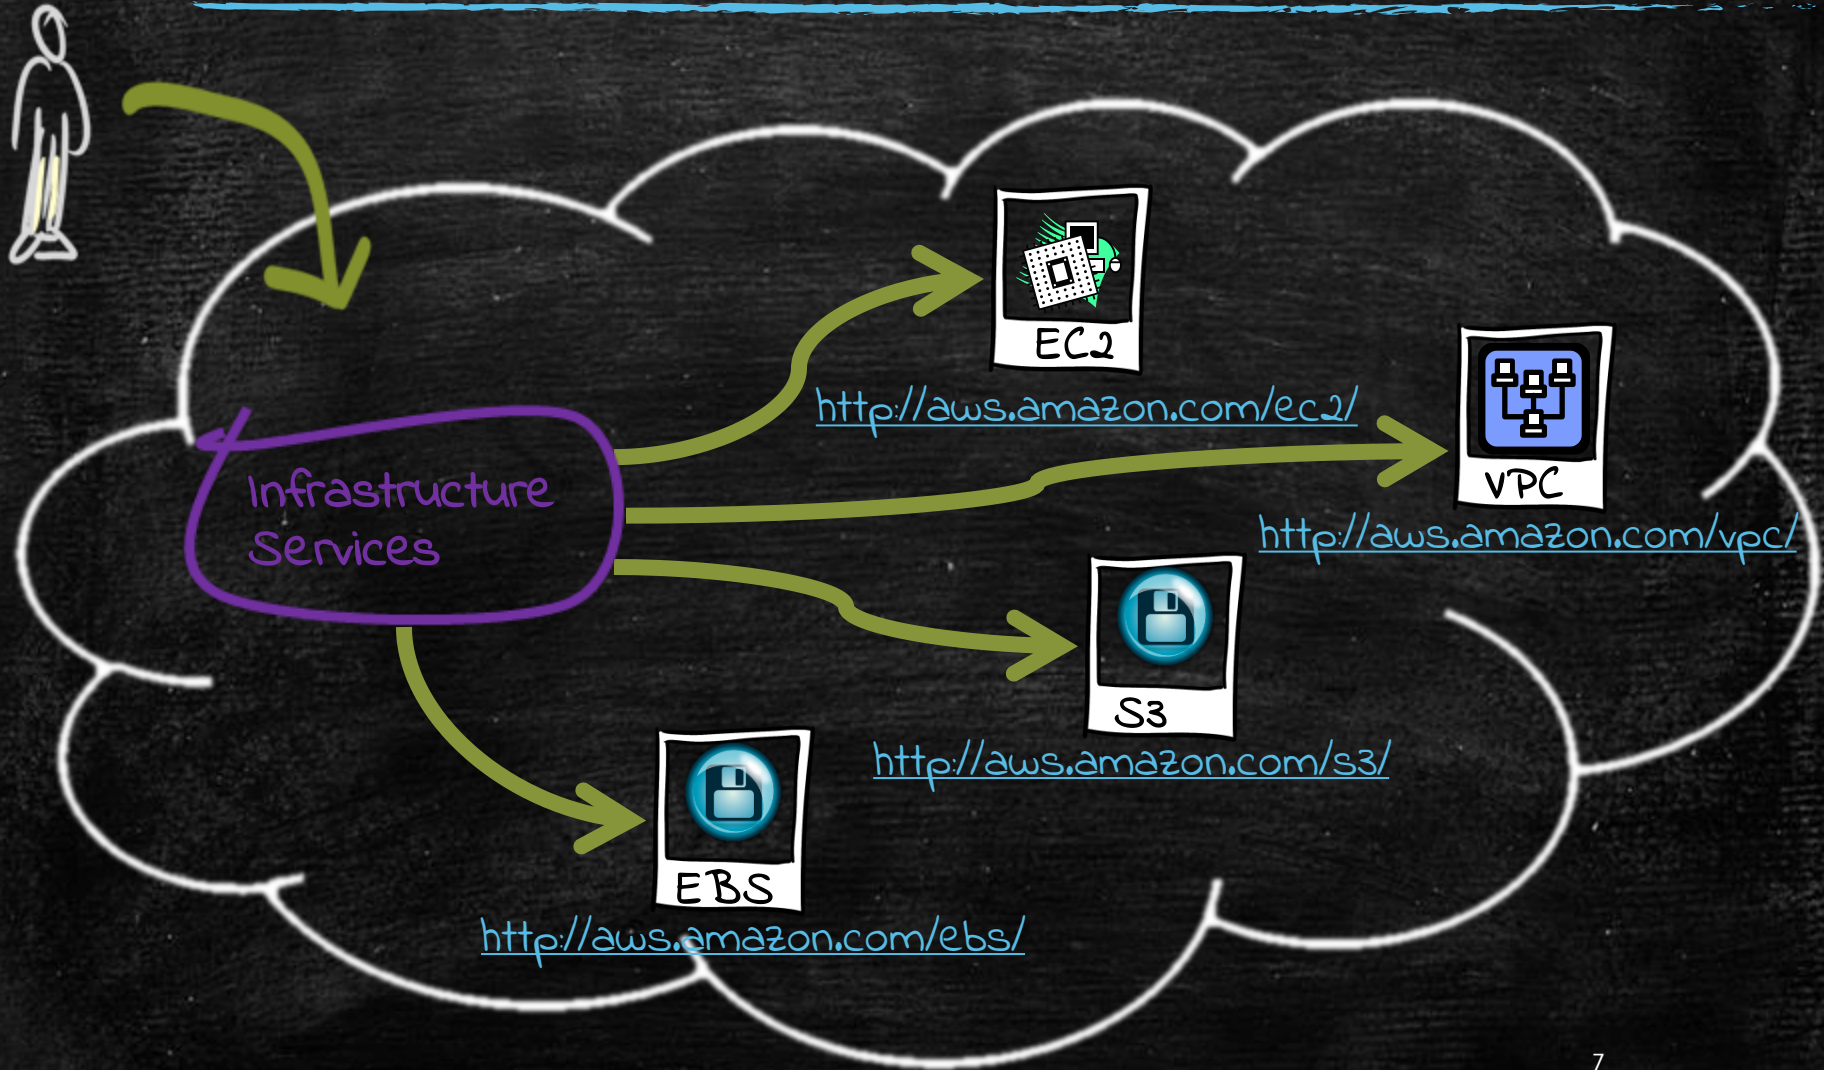
The Amazon web Services Universe



Management Interface



Infrastructure Services





Amazon Elastic Compute Cloud (EC2)

- A web service that provides **resizable compute capacity** in the cloud.
- EC2 allows **creating virtual Machines (VM) on-demand**. Pre-configured **templated Amazon Machine Image (AMI)** can be used get running immediately. Creating and sharing your own AMI is also possible via the **AWS Marketplace**.
- **Auto Scaling** allows **automatically scale of the capacity up** seamlessly during **demand spikes** to maintain performance, and **scales down** during **demand lulls** to minimize costs.
- **Elastic Load Balancing** automatically distributes incoming application traffic across multiple Amazon EC2 instances.
- Provide tools to build **failure resilient applications** by launching application instances in **separate Availability Zones**.
- Pay only for resources actually consume, **instance-hours**.
- **VM Import/Export** enables you to easily import virtual machine images from your existing environment to Amazon EC2 instances and export them back at any time.



EC2 Instances

■ Micro instances (t1.micro):

- Micro Instance 613 MiB of memory, up to 2 ECUs (for short periodic bursts), EBS storage only, 32-bit or 64-bit platform.

■ Standard Instances provide customers with a balanced set of resources and a low cost platform.

- **M1 Small Instance (Default)** 1.7 GiB of memory, 1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit), 160 GiB of local instance storage, 32-bit or 64-bit platform
- **M1 Medium Instance** 3.75 GiB of memory, 2 EC2 Compute Units (1 virtual core with 2 EC2 Compute Units each), 410 GiB of local instance storage, 32-bit or 64-bit platform
- **M1 Large Instance** 7.5 GiB of memory, 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each), 850 GiB of local instance storage, 64-bit platform
- **M1 Extra Large Instance** 15 GiB of memory, 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each), 1690 GiB of local instance storage, 64-bit platform
- **M3 Extra Large Instance** 15 GiB of memory, 13 EC2 Compute Units (4 virtual cores with 3.25 EC2 Compute Units each), EBS storage only, 64-bit platform
- **M3 Double Extra Large Instance** 30 GiB of memory, 26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units each), EBS storage only, 64-bit platform

one EC2 Compute Unit (ECU) provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor.



EC2 High Performance Instances

High-Memory Instances:

- **High-Memory Extra Large Instance** 17.1 GiB memory, 6.5 ECU (2 virtual cores with 3.25 EC2 Compute Units each), 420 GiB of local instance storage, 64-bit platform
- **High-Memory Double Extra Large Instance** 34.2 GiB of memory, 13 EC2 Compute Units (4 virtual cores with 3.25 EC2 Compute Units each), 850 GiB of local instance storage, 64-bit platform
- **High-Memory Quadruple Extra Large Instance** 68.4 GiB of memory, 26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units each), 1690 GiB of local instance storage, 64-bit platform

High-CPU Instances

- **High-CPU Medium Instance** 1.7 GiB of memory, 5 EC2 Compute Units (2 virtual cores with 2.5 EC2 Compute Units each), 350 GiB of local instance storage, 32-bit or 64-bit platform
- **High-CPU Extra Large Instance** 7 GiB of memory, 20 EC2 Compute Units (8 virtual cores with 2.5 EC2 Compute Units each), 1690 GiB of local instance storage, 64-bit platform

High Storage Instances

- **High Storage Eight Extra Large** 117 GiB memory, 35 EC2 Compute Units, 24 * 2 TB of hard disk drive local instance storage, 64-bit platform, 10 Gigabit Ethernet

High I/O Instances

- **High I/O Quadruple Extra Large** 60.5 GiB memory, 35 EC2 Compute Units, 2 * 1024₁₆ GiB of SSD-based local instance storage, 64-bit platform, 10 Gigabit Ethernet



EC2 Cluster Instances

- **Cluster Compute Instances** provide proportionally high CPU resources with increased network performance and are well suited for High Performance Compute (HPC) applications and other demanding network-bound applications.
 - **Cluster Compute Eight Extra Large** 60.5 GiB memory, 88 EC2 Compute Units, 3370 GiB of local instance storage, 64-bit platform, 10 Gigabit Ethernet
- **High Memory Cluster Instances** provide proportionally high CPU and memory resources with increased network performance, and are well suited for memory-intensive applications including in-memory analytics, graph analysis, and scientific computing.
 - **High Memory Cluster Eight Extra Large** 244 GiB memory, 88 EC2 Compute Units, 240 GiB of local instance storage, 64-bit platform, 10 Gigabit Ethernet
- **Cluster GPU Instances** provide general-purpose graphics processing units (GPUs) with proportionally high CPU and increased network performance for applications benefitting from highly parallelized processing, including HPC, rendering and media processing applications.
 - **Cluster GPU Quadruple Extra Large** 22 GiB memory, 33.5 EC2 Compute Units, 2 x NVIDIA Tesla "Fermi" M2050 GPUs, 1690 GiB of local instance storage, 64-bit platform, 10 Gigabit Ethernet.



EC2 Payment methods

- **On-Demand Instances** let you pay for compute capacity by the hour with no long-term commitments.
- **Reserved Instances** give you the option to make a low, one-time payment for each instance you want to reserve and in turn receive a significant discount on the hourly charge for that instance.
- **Spot Instances** allow customers to bid on unused Amazon EC2 capacity and run those instances for as long as their bid exceeds the current Spot Price.

Infrastructure
Services



EC2

Virtual Servers in the Cloud

Amazon Elastic Compute Cloud (EC2)

DEMO!

Amazon Elastic Block Store (EBS)

- Provides **block level storage** volumes (1 GB to 1 TB) for use with Amazon EC2 instances.
 - **Multiple volumes** can be mounted to the same instance.
 - EBS volumes are **network-attached**, and **persist independently** from the life of an instance.
 - Storage volumes behave like **raw, unformatted block devices**, allowing users to **create a file system** on top of Amazon EBS volumes, or use them in any other way you would use a block device (**like a hard drive**).
- EBS volumes are **placed in a specific Availability Zone**, and can then be **attached to instances also in that same Availability Zone**.
- Each storage volume is **automatically replicated within the same Availability Zone**.
- EBS provides the ability to **create point-in-time snapshots of volumes**, which are **persisted to Amazon S3**.
 - These snapshots can be **used as the starting point for new Amazon EBS volumes**, and protect data for long-term durability.
 - The **same snapshot can be used to instantiate as many volumes** as you wish.
 - These snapshots **can be copied across AWS regions**.

EBS volumes


- **Standard volumes** offer storage for applications with **moderate or bursty I/O** requirements.
 - Standard volumes deliver approximately **100 IOPS** on average.
 - well suited for use as **boot volumes**, where the burst capability provides fast instance start-up times.
- **Provisioned IOPS volumes** are designed to deliver **predictable, high performance for I/O intensive workloads such as databases**.
 - You specify an **IOPS rate** when creating a volume, and EBS provisions that rate for the lifetime of the volume.
 - Amazon EBS currently supports **up to 4000 IOPS per Provisioned IOPS volume**.
 - You can **stripe multiple volumes together** to deliver thousands of IOPS per EC2 instance.
- To enable your EC2 instances to fully utilize the IOPS provisioned on an EBS volume;
 - Launch selected Amazon EC2 instance types as **"EBS-optimized"** instances.
 - **EBS-optimized instances deliver dedicated throughput between Amazon EC2 and Amazon EBS**, with options between **500 Mbps and 1000 Mbps** depending on the instance type used.
- EBS charges based on **per GB-month AND per 1 million I/O requests**



Amazon Simple Storage Service (S3)

- Amazon S3 provides a simple web services interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web.
- write, read, and delete objects containing from 1 byte to 5 terabytes of data each. The number of objects you can store is unlimited.
- Each object is stored in a bucket and retrieved via a unique, developer-assigned key.
 - A bucket can be stored in one of several Regions.
 - You can choose a Region to optimize for latency, minimize costs, or address regulatory requirements.
 - objects stored in a Region never leave the Region unless you transfer them out.
- Authentication mechanisms are provided to ensure that data is kept secure from unauthorized access.
 - objects can be made private or public, and rights can be granted to specific users.
- S3 charges based on per GB-month AND per I/O requests AND per data modification requests.

Infrastructure Services



Amazon Simple Storage Service (S3)

Create Bucket Actions ▾ None Properties Transfers ↻ ?

All Buckets

Name
aws-test-grep
aws-test-HE
cloud-final
cloud-final-results
cloudfinal-out1
dasbucket1
data432
eh-test
EH_test
final-report1
final-reprot1
finalreport
hanaEsi_final
mmyawsbucket
mmyawsbucket1
ms499-mapredcapri-test
ms499-wordcount-test

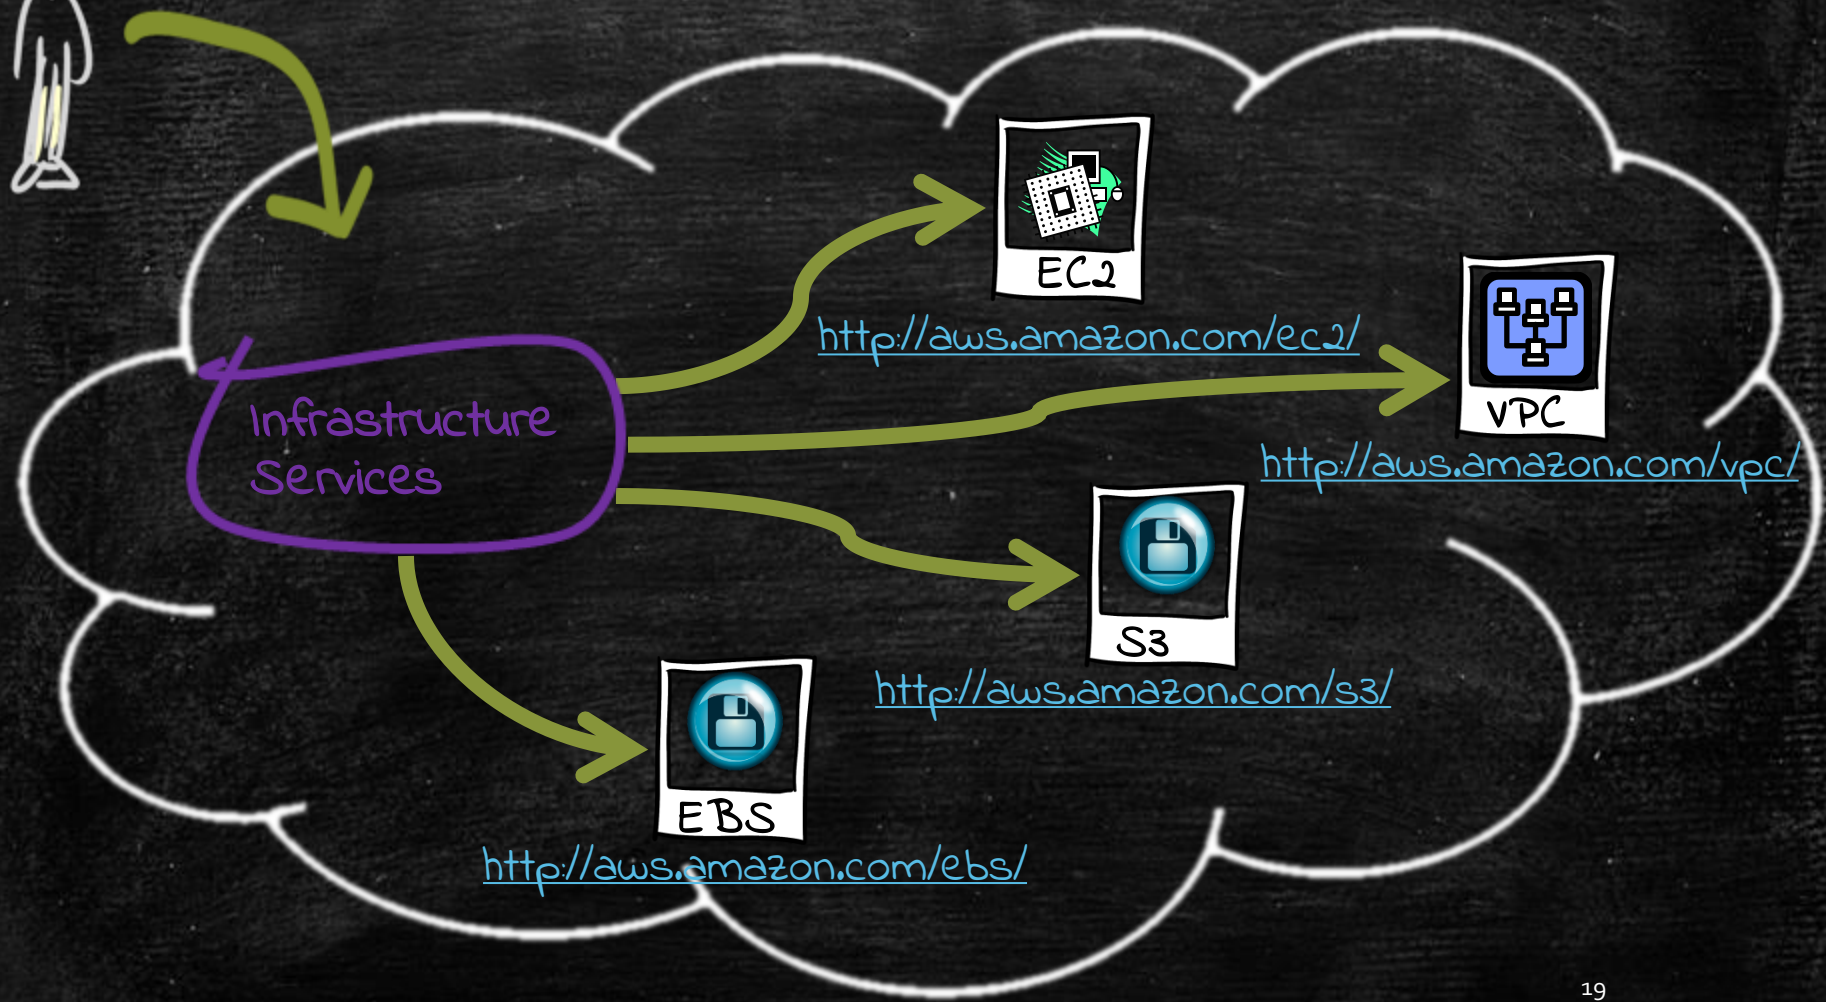
S3 Console
Let us know what you think, with the Feedback button at the bottom of the page.



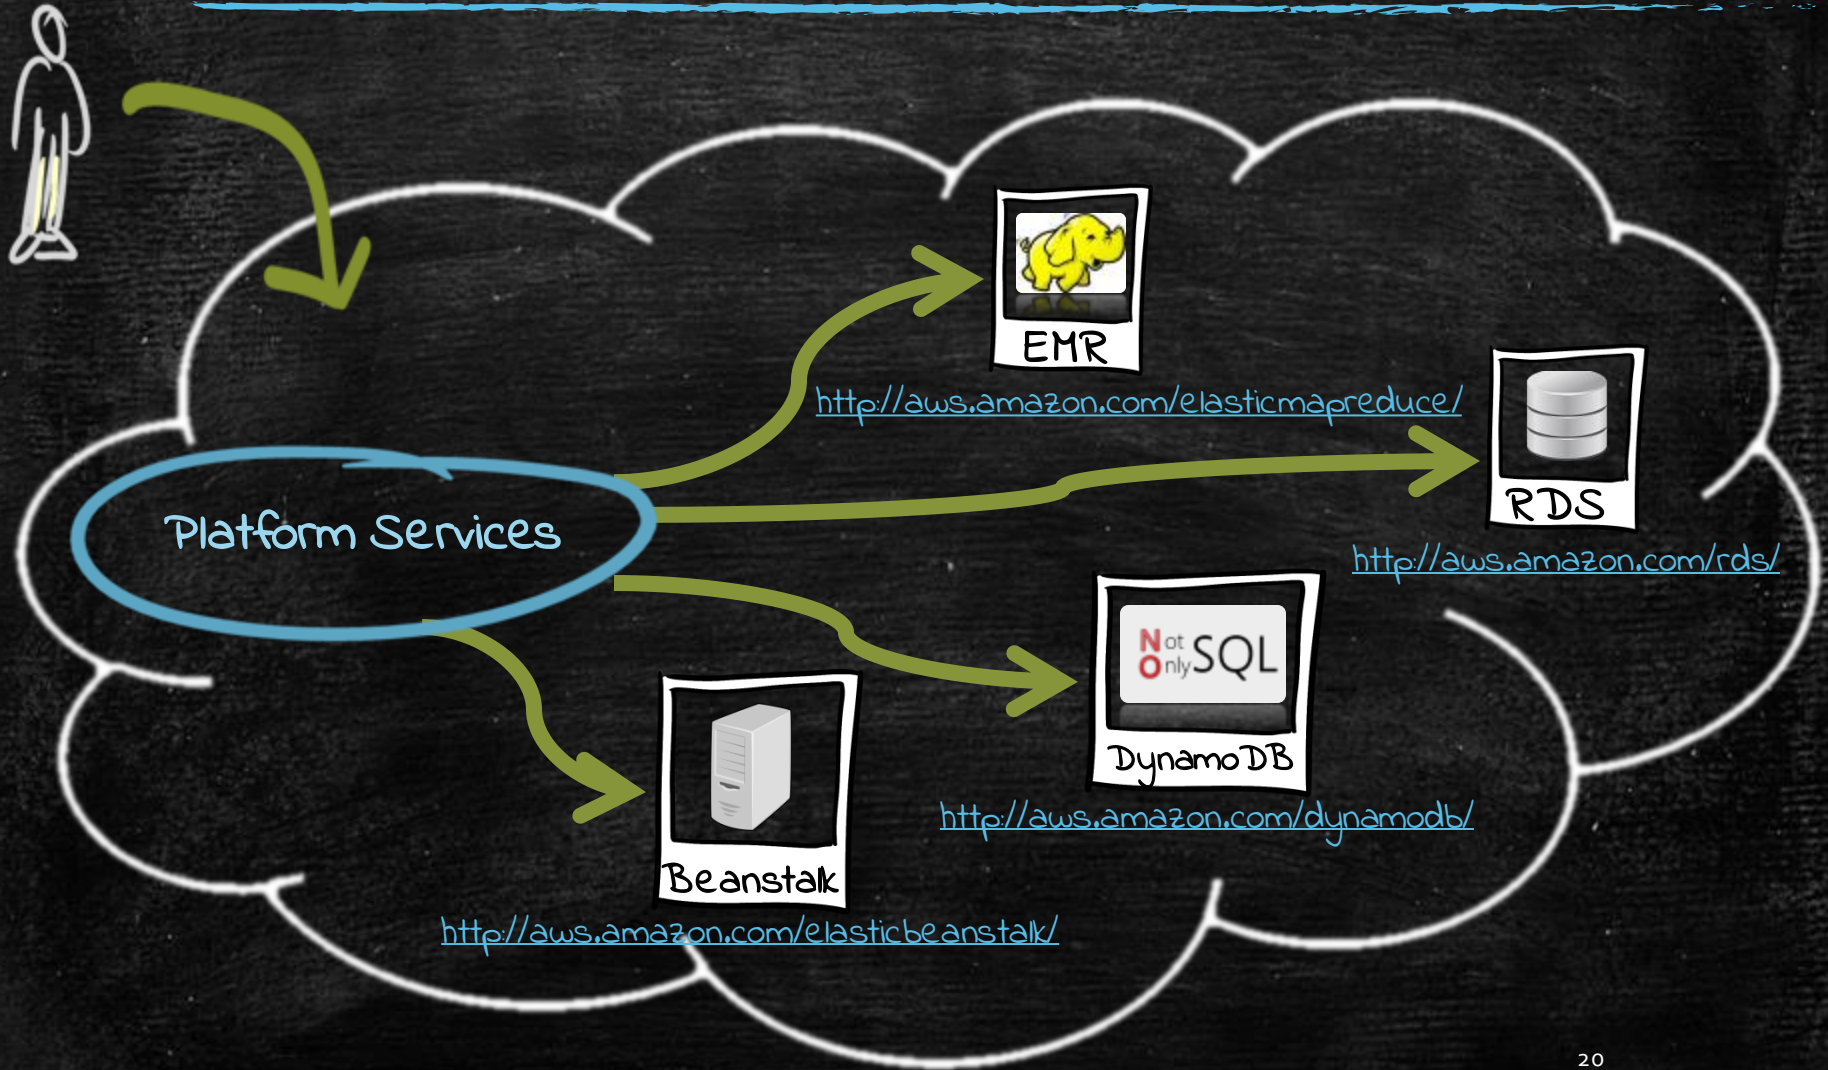
Amazon Virtual Private Cloud (VPC)

- Amazon VPC lets you provision a logically isolated section of the Amazon web Services (AWS) Cloud.
- You have complete control over your virtual networking environment, including:
 - selection of your own IP address range,
 - creation of subnets, and
 - configuration of route tables and network gateways.
- VPC allows bridging with an onsite IT infrastructure with an encrypted VPN connection with an extra charge per VPN Connection-hour.
- There is no additional charge for using Amazon Virtual Private Cloud, aside from the normal Amazon EC2 usage charges.

Demo & Questions



Platform Services





Amazon Elastic MapReduce (EMR)

- Amazon EMR is a web service that makes it easy to quickly and cost-effectively process vast amounts of data using Hadoop.
- Amazon EMR distribute the data and processing across a resizable cluster of Amazon EC2 instances.
- with Amazon EMR you can launch a persistent cluster that stays up indefinitely or a temporary cluster that terminates after the analysis is complete.
- Amazon EMR supports a variety of Amazon EC2 instance types and Amazon EC2 pricing options (On-Demand, Reserved, and Spot).
- when launching an Amazon EMR cluster (also called a "job flow"), you choose how many and what type of Amazon EC2 instances to provision.
- The Amazon EMR price is in addition to the Amazon EC2 price.
- Amazon EMR is used in a variety of applications, including log analysis, web indexing, data warehousing, machine learning, financial analysis, scientific simulation, and bioinformatics.

Platform Services



Elastic MapReduce
Managed Hadoop Framework

Amazon Elastic MapReduce (EMR)

Your Elastic MapReduce Job Flows

Create New Job Flow Terminate Debug

Show/Hide Refresh Help

Viewing: All

1 to 3 of 3 Job Flows

Name	State	Creation Date	Elapsed Time	Normalized Instance Hours
My Job Flow	COMPLETED	2013-10-17 17:07 EDT	0 hours 6 minutes	3
My Job Flow	COMPLETED	2013-10-17 16:42 EDT	0 hours 20 minutes	3
My Job Flow	COMPLETED	2013-10-17 16:40 EDT	0 hours 7 minutes	3

1 Job Flow selected

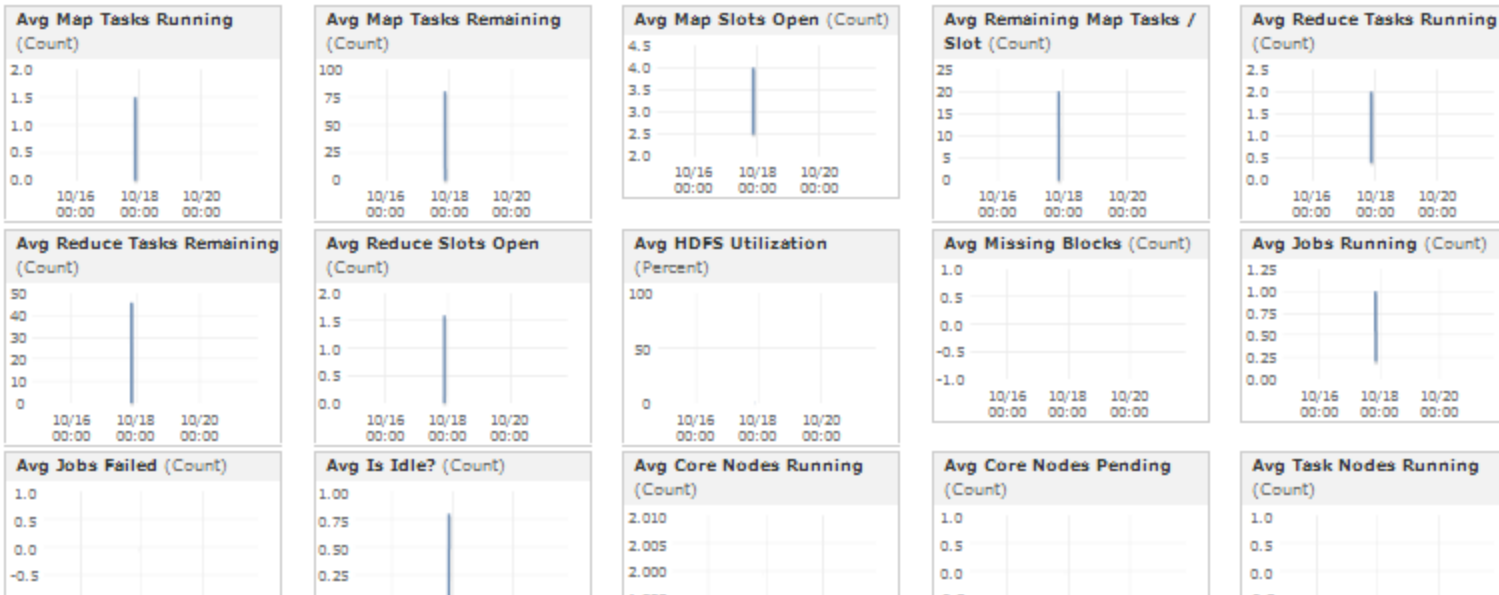
Job Flow: j-7D8QY3YLLCN6

Last State Change: Steps completed

Description Steps Bootstrap Actions Instance Groups **Monitoring**

Times are displayed in UTC.

Time Range: Last Week Refresh





Amazon Relational Database Service (RDS)

- Amazon RDS is a web service that makes it easy to set up, operate, and scale a relational database in the cloud.
- Amazon RDS gives access to the capabilities of a familiar MySQL, Oracle or Microsoft SQL Server database engine.
 - Code, applications, and tools already used with existing databases can be used with RDS.
- Amazon RDS automatically patches the database software and backs up the database, storing the backups for a user-defined retention period and enabling point-in-time recovery.
- Amazon RDS provides scaling the compute resources or storage capacity associated with the Database Instance.
- Pay only for the resources actually consumed, based on the DB Instance hours consumed, database storage, backup storage, and data transfer.
 - On-Demand DB Instances let you pay for compute capacity by the hour with no long-term commitments.
 - Reserved DB Instances give the option to make a low, one-time payment for each DB Instance and in turn receive a significant discount on the hourly usage charge for that DB Instance.

SQL Databases

- In relational databases (SQL Databases), ACID (Atomicity, Consistency, Isolation, Durability) is a set of properties that guarantee that database transactions are processed reliably.
 - Atomicity requires that each transaction is "all or nothing": if one part of the transaction fails, the entire transaction fails, and the database state is left unchanged.
 - The consistency property ensures that any transaction will bring the database from one valid state to another.
 - The isolation property ensures that the concurrent execution of transactions results in a system state that would be obtained if transactions were executed serially.
 - Durability means that once a transaction has been committed, it will remain so, even in the event of power loss, crashes, or errors.



Amazon DynamoDB

- **DynamoDB** is a fast, fully managed **NoSQL database service** that makes it simple and cost-effective to store and retrieve any amount of data, and serve any level of request traffic.
- All data items are stored on **Solid State Drives (SSDs)**, and are **replicated** across **3 Availability Zones** for high availability and durability.
- DynamoDB **tables do not have fixed schemas**, and each item may have a **different number of attributes**.
- DynamoDB has **no upfront costs** and implements a **pay as you go** plan as a **flat hourly rate** based on the **capacity reserved**.

NoSQL Databases

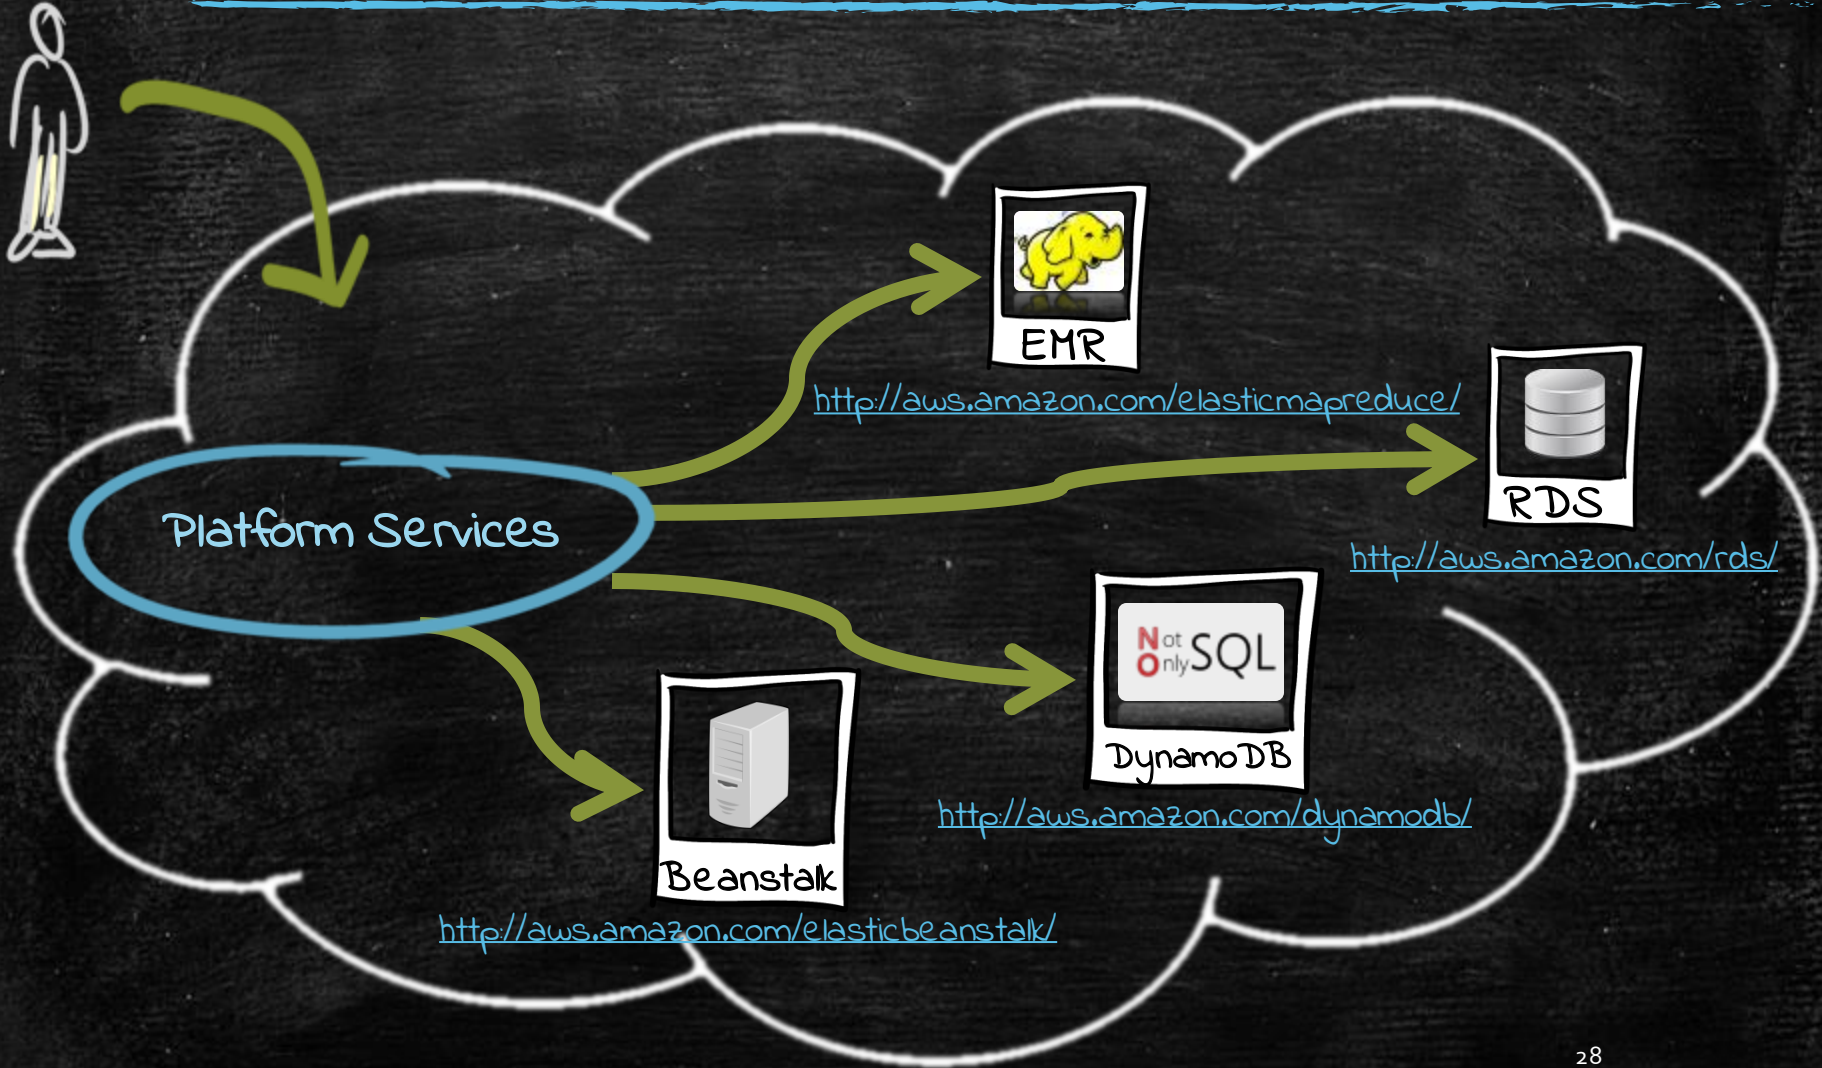
- A **NoSQL database** provides a mechanism for storage and retrieval of data that **employs less constrained consistency models** than traditional relational databases.
- NoSQL databases only support **Eventual Consistency** which is a consistency model used in distributed computing that **informally guarantees that, if no new updates are made to a given data item, eventually all accesses to that item will return the last updated value.**
- NoSQL databases are often **highly optimized key-value stores** intended for **simple retrieval and appending operations**, with the goal being significant performance benefits in terms of latency and throughput.
- **Key-value stores** allow the application to **store its data in a schema-less way.**
 - The data could be stored in a **datatype of a programming language** or an object. Because of this, there is **no need for a fixed data model.**



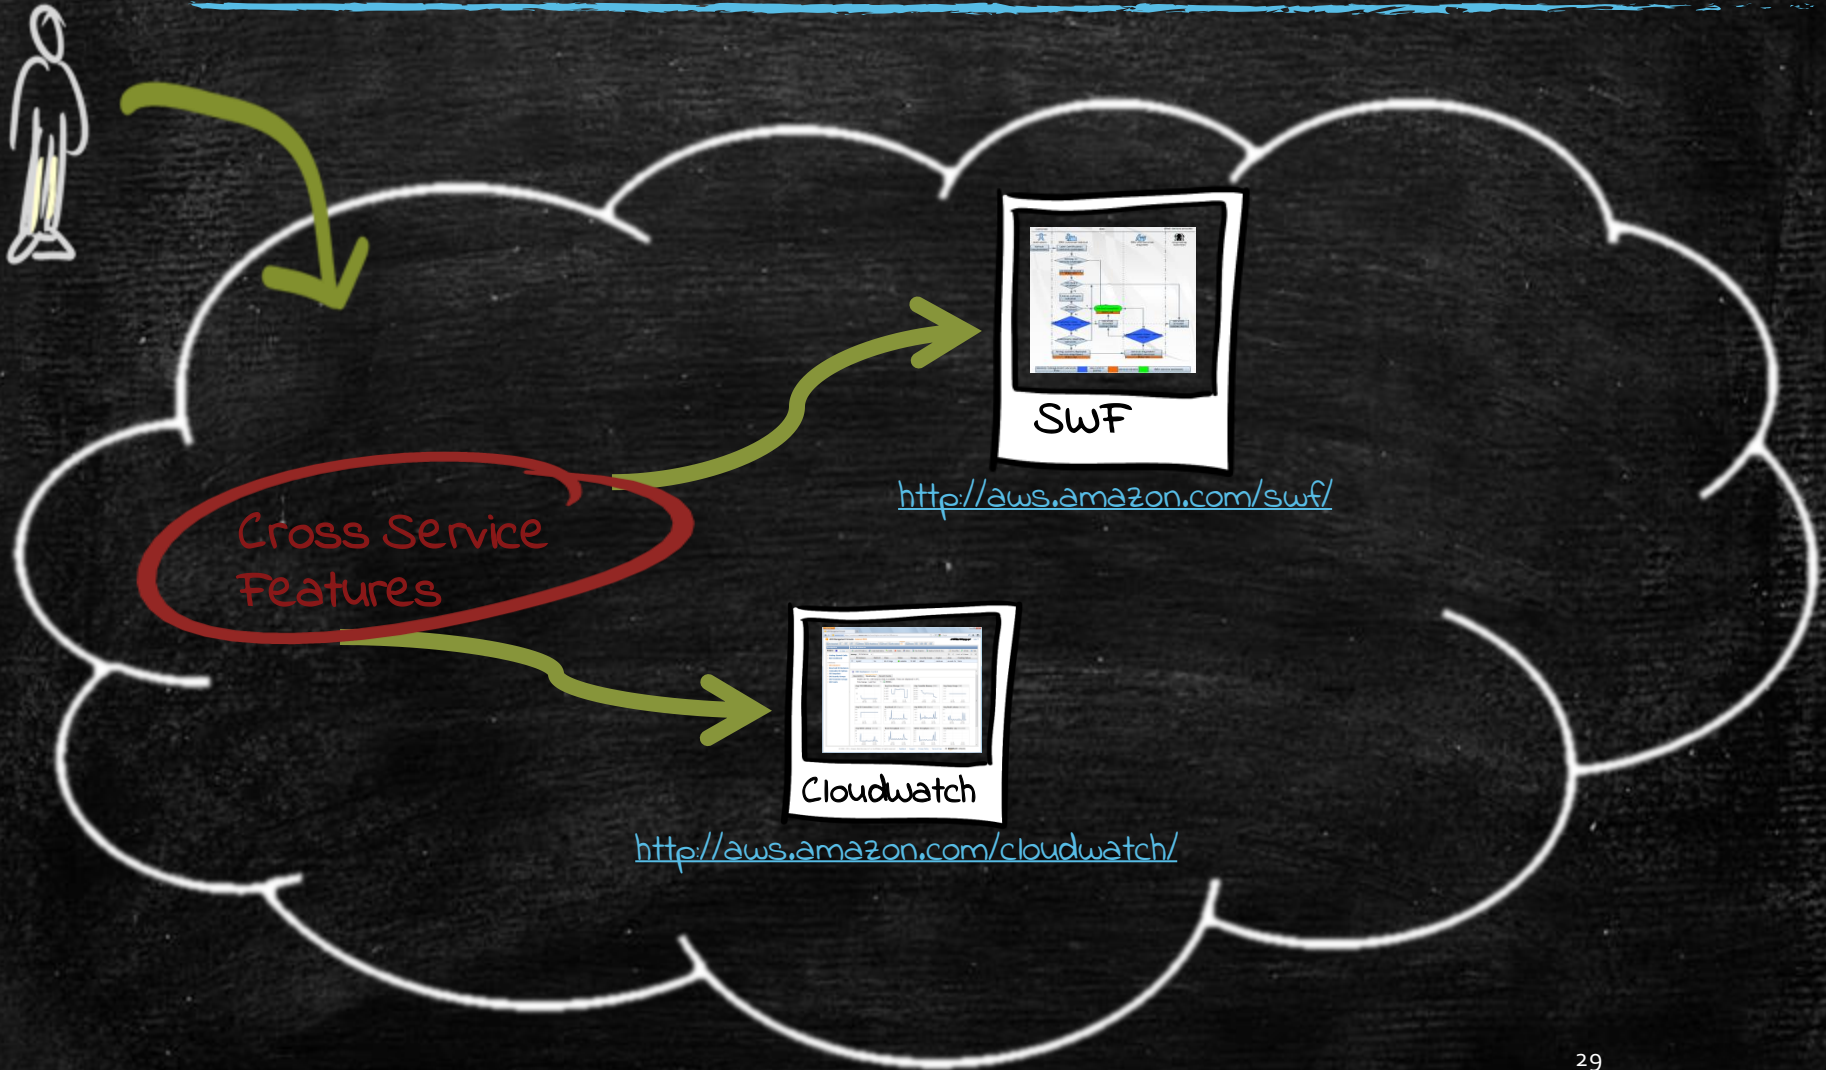
Amazon Elastic Beanstalk

- **AWS Elastic Beanstalk** provides a solution to **quickly deploy** and **manage** applications in the AWS cloud.
- You **simply upload your application**, and Elastic Beanstalk **automatically** handles the deployment details of capacity **provisioning**, **load balancing**, **auto-scaling**, and application health **monitoring**.
- Elastic Beanstalk leverages AWS services such as **Amazon EC2**, **Amazon S3**, ...
- To ensure easy portability of your application, Elastic Beanstalk is built using familiar software stacks such as:
 - Apache HTTP Server for Node.js, PHP and Python
 - Passenger for Ruby
 - IIS 7.5 for .NET
 - Apache Tomcat for Java.
- There is **no additional charge** for Elastic Beanstalk - you **pay only for the AWS resources** needed to store and run your applications.

Questions



Cross Service Features



Cross-Service
Features



CloudWatch

Resource and Application Monitoring

Amazon Cloudwatch

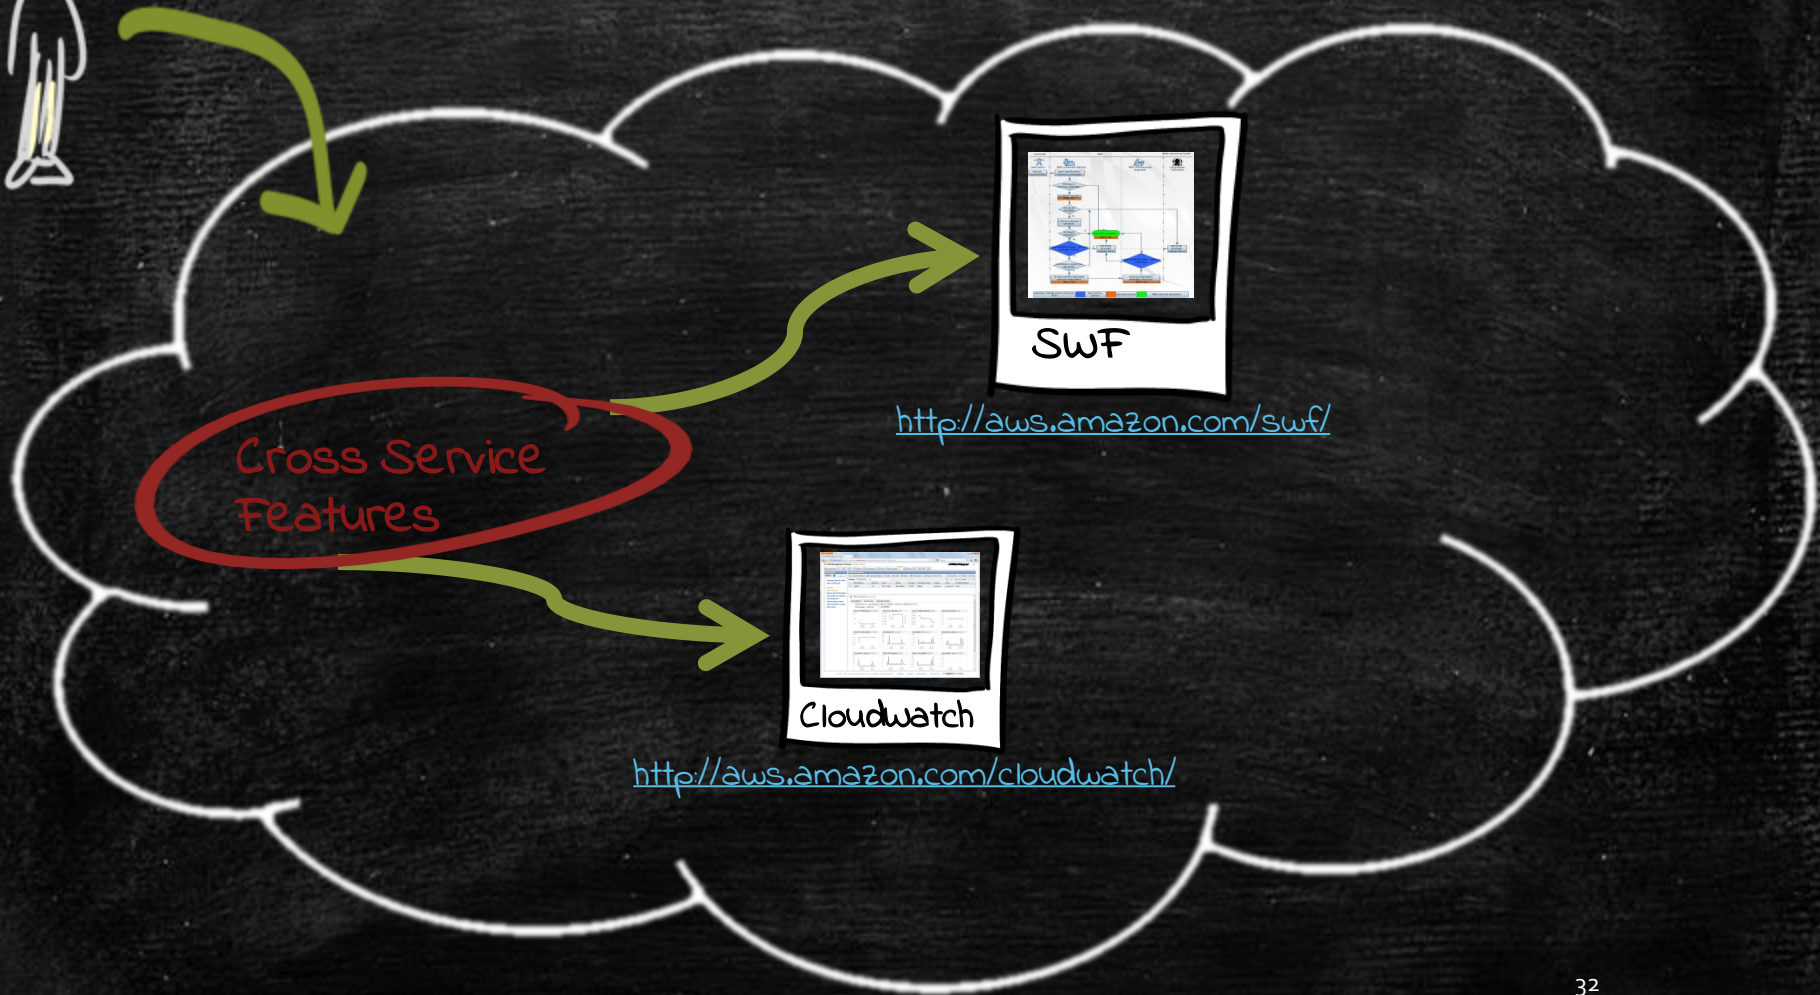
- Amazon **Cloudwatch** provides **monitoring for AWS cloud resources** and the applications customers run on AWS.
- Amazon Cloudwatch lets you **programmatically** retrieve your **monitoring data**, view **graphs**, and **set alarms** to help you troubleshoot, spot trends, and **take automated action** based on the state of your cloud environment.
- Amazon Cloudwatch enables you to monitor your AWS resources **up-to-the-minute in real-time**, including:
 - Amazon EC2 instances,
 - Amazon EBS volumes,
 - Elastic Load Balancers,
 - Amazon RDS DB instances.
- Metrics such as **CPU utilization**, **latency**, and **request counts** are provided automatically for these AWS resources.
- Customers can also supply their **own custom application and system metrics**, such as **memory usage**, **transaction volumes**, or **error rates**.



Amazon Simple workflow Service (SWF)

- Amazon SWF is a task coordination and state management service for cloud applications.
- Using Amazon SWF, you structure the various processing steps in an application that runs across one or more machines as a set of "tasks."
- Amazon SWF manages dependencies between the tasks, schedules the tasks for execution, and runs any logic that needs to be executed in parallel.
- The service also tracks the tasks' progress.
- As the business requirements change, Amazon SWF makes it easy to change application logic without having to worry about the underlying state machinery and flow control.

Questions





watch out for unexpected Costs

- When you finish your work remember to make sure of the following to **avoid unwanted costs**:
 - **Delete** your **S3** objects.
 - Stop or **Shut Down** your **EC2** and **RDS** instances.
- The customer is responsible for the resources he's using. AWS **declines any responsibility** if the customer forgets to shut down resources.

AWS Free Usage Tier

MORE
INFO ↘

- <http://aws.amazon.com/free/>