



UNIVERSITY OF TARTU

INSTITUTE OF COMPUTER SCIENCE



Basics of Cloud Computing – Lecture 2

Cloud Providers

Satish Srirama



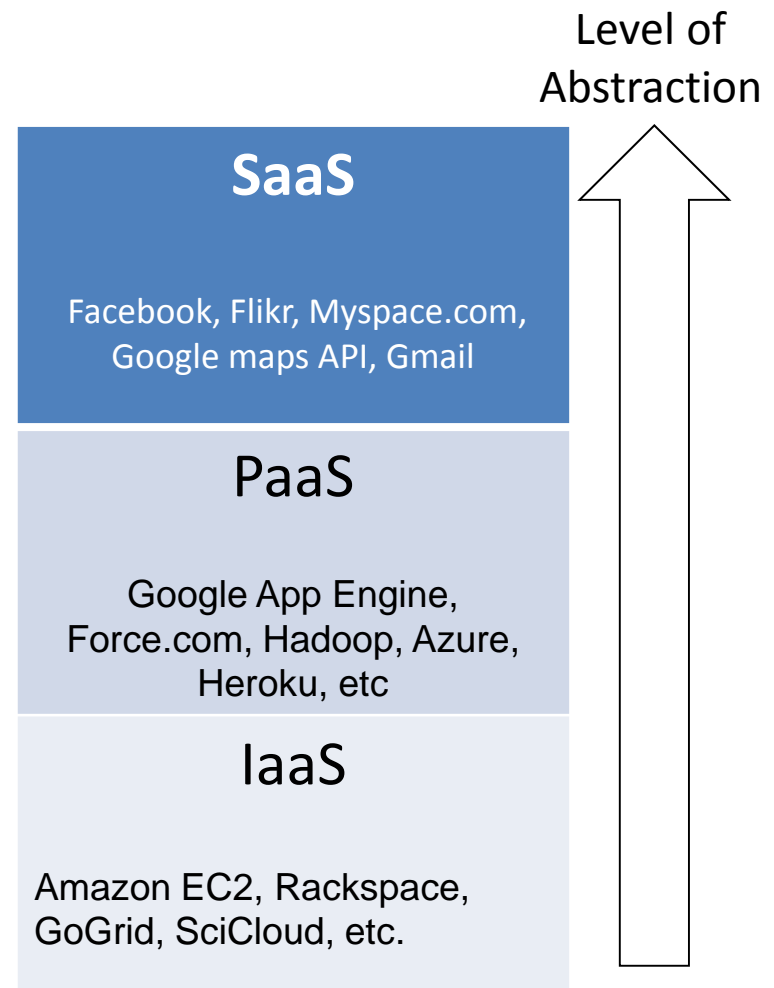
Mobile & Cloud Lab

Outline

- Cloud computing services – recap
- Amazon cloud services
 - Elastic Compute Cloud (EC2)
 - Storage services - Amazon S3 and EBS
- Cloud managers
- Eucalyptus

Cloud Computing - Services

- Software as a Service – SaaS
 - A way to access applications hosted on the web through your web browser
- Platform as a Service – PaaS
 - Provides a computing platform and a solution stack (e.g. LAMP) as a service
- Infrastructure as a Service – IaaS
 - Use of commodity computers, distributed across Internet, to perform parallel processing, distributed storage, indexing and mining of data
 - Virtualization



Cloud Infrastructure

- Provisioning of computing resources
 - CPU, Memory, Processing
 - Basically an “Operating System” on demand
- Usually billed on a per-hour usage model
- Players in this space
 - Amazon EC2, Rackspace, GoGrid, Eucalyptus/Openstack based SciCloud
- Management providers: RightScale, HybridFox, ElasticFox, Amazon Management Console

Cloud Storage

- Provisioning of data storage
 - Either file/object based or Database like functionality
- Billed on bandwidth and storage consumed
- Players in the space
 - Amazon S3, Amazon EBS
 - Amazon's SimpleDB, Google's BigTable, Apache Cassandra
- Management Providers: Amazon Management Console, Jungle Disk, Elephant Disk, PutPlace.com

Cloud Platforms

- Provides a complete software stack
 - Provides a computing platform and a solution stack
 - An IDE for the cloud
- Takes care of: Runtimes, Load balancing, Resource provisioning
- Players in the space
 - Google AppEngine - Python (initially, now also Java)
 - Force.com (SalesForce)
 - Heroku

Cloud Applications

- Applications that are completely 'online'
- Operate on data that is stored in the 'cloud' or 'ether'
- No client software generally required
- Billing: Ad. Revenue, Premium Services
- Players in this space
 - Google Apps - Gmail / Google Docs
 - Apple's MobileMe
 - Microsoft's Live - Hotmail, Live Spaces
 - Salesforce.com

Other cloud services

- Provides services, which other applications can utilize
- Usually free for non-commercial use
- Players in the space
 - Google/Yahoo Maps
 - Google/Yahoo Web Services
 - Amazon Merchant Services
 - Amazon Simple Queue Service

Providers we focus at

- Amazon Web Services
 - Amazon EC2
 - Amazon S3
 - Amazon EBS
- Private cloud enabling technologies
 - Eucalyptus
 - OpenStack
 - SciCloud
- Management providers
 - ElasticFox
 - RightScale

Amazon Elastic Compute Cloud (EC2)

- One of the very early pioneers of cloud computing
- In a nutshell:
 - On Demand “Operating System”
- Complete virtual computer with CPU, Memory and disk space
 - Based on the XEN virtual image platform
- Variety of operating systems available
 - Linux (Fedora, Ubuntu, CentOS, etc)
 - Open Solaris
 - Microsoft Windows

EC2 continued...

- Very simple pricing structure
 - CPU hours
 - Machine size
 - Bandwidth in and out of cloud
- Extremely FAST start up
 - 2-3 minutes from start to finish
- Instance snapshotting
- Very large/generous disk space provisioning
 - 160GB minimum for the standard instances
- Flexible API to control everything
- Wide range of virtual machine types

EC2 Instance types – General purpose*

Instance	CPU	Memory	Storage	Platform	API Name	Price (per h.)
Small (default)	1 EC2 computing unit	2.0 GB	EBS Only	32/64 bit	t2.small	\$0.023 (*nix) \$0.032 (win)
Medium	~3 EC2 units	4 GB	EBS Only	32/64 bit	t2.medium	\$0.047 (*nix) \$0.065 (win)
Large	6.5 EC2 units	8 GB	EBS Only	64 bit	m4.large	\$0.108 (*nix) \$0.203 (win)
Extra Large	13 EC2 units	16 GB	EBS Only	64 bit	m4.xlarge	\$0.215 (*nix) \$0.404 (win)
Micro	1 EC2 unit	1 GB	EBS storage only	32/64 bit	t2.micro	\$0.012 (*nix) \$0.017 (win)

*Data taken on 13.02.2017

<http://aws.amazon.com/ec2/instance-types/>

EC2 advanced Instance types*

Instance	CPU	Memory	Storage	Platform	API Name	Price (per h.)
High-Memory Instances	4 - 128 EC2 units	26 – 195 GB	Upto 2 X 320 GB SSD	64 bit	r3.xlarge r3.2xlarge r3.4xlarge	Available at: http://aws.amazon.com/ec2/pricing/
High-CPU Instances	16 - 132 EC2 units	3.75 -60 GB	EBS Only	64 bit	c4.large c4.xlarge	Same as above

- Other types

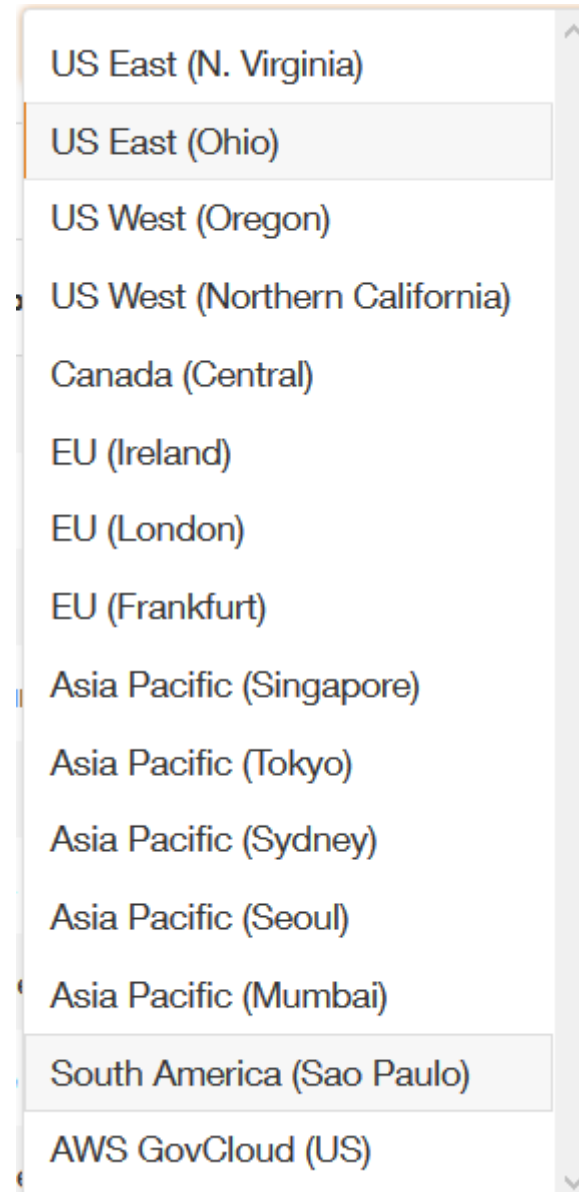
- High I/O Instances
- High Storage Instances
- Cluster Compute Instances

*Data taken on 13.02.2017

<http://aws.amazon.com/ec2/instance-types/>

EC2 Terminology

- Having an account
 - Access Key, Secret Key, Security group
- Availability Zone (15)
- Amazon Machine Image (AMI)
 - A Virtual Machine File
 - ami-XXXX
 - Stored in a special bucket in Amazon's S3
- Public and Private instances available
 - Private instances incur only S3 storage costs
- A Running Machine
 - Amazon Instance (i-XXXX)
 - Booted/Created from an Amazon Image
- Elastic IP addresses



Troubles with EC2

- On power-off all hard disk data is lost
- IP addresses are assigned at random
- Can't turn off public IP address
- Do not forget to terminate the instances

Simple Storage Service (S3)

- Enables you to upload, download, and store data across the Internet
- Buckets store data
 - Buckets are the fundamental container in Amazon S3 for data storage
 - 100 buckets for account
 - No limit on no of objects that can be stored in a bucket
 - Can store up to 5 TB of data in one object
 - Object stores Data and Metadata
 - Objects stored in a Region never leave the Region
 - You cannot modify or append data to an existing object

Simple Storage Service (S3) - continued

- Objects are retrieved via a unique, developer-assigned key
- Keys
 - Example: `http://doc.s3.amazonaws.com/2006-03-01/AmazonS3.wsdl`
"doc" is the bucket name and "2006-03-01/AmazonS3.wsdl" is the key.
- Prices
 - Storage: \$0.023 per GB/Month up to first 50 TB
 - Next 450 TB / month \$0.023 per GB/Month
 - Data Transfer Out from S3 to AWS different zone: \$0.02/GB
 - PUT, COPY, POST, or LIST Requests: \$0.005 per 1,000 Requests
 - GET requests: \$0.004 per 10,000 Requests
- Download Data
 - Downloading is possible from any where
 - You can enable others to download and can charge them (Using Amazon DevPay)

<http://aws.amazon.com/s3/>

Elastic Block Storage (EBS)

- The answer to the persistence problem
- Raw unformatted external block devices
- Allocate 1GB to 1TB volumes
- 20 volumes per account
- Format with your own choice of file system
- Attach to running instances in the same availability zone
- Automatically replicated to prevent data loss
- Create snapshots for backup, or to create new volumes from

Basic ec2-api-tools

- ec2-describe-images
- ec2-describe-instances
- ec2-run-instances
- ec2-create-volume
- ec2-attach-volume
- ec2-allocate-address
- ec2-associate-address
- ec2-terminate-instances

Other cloud services from Amazon

- AWS management console
- Amazon Simple Queue Service (SQS)
- Amazon DynamoDB
- Amazon Relational Database Service
- Amazon CloudWatch
- Amazon Elastic MapReduce
- Elastic Load Balancing
- etc.

Getting started

- Signup for an AWS account
 - Use your existing Amazon.com account if you want
- Register a credit card
 - Billed on the 1st of every month for previous month usage
- <http://aws.amazon.com/ec2>

Account Activity

[View Previous Statement](#)

Summary of This Month's Activity as of October 18, 2009

Billing Cycle for this Report: October 1 - October 31, 2009

AWS service usage charges on this page currently show activity through approximately 10/18/2009 14:59 GMT.

[Expand All](#) | [Collapse All](#)

Rate	Usage	Totals
Amazon Elastic Compute Cloud		
View/Edit Service		
United States		
Elastic IP Addresses		
\$0.00 per Elastic IP address remap - first 100 remaps / month	13 Count	0.00
		0.00
European Union		
Amazon EC2 running Linux/UNIX		
\$0.44 per Large Instance (m1.large) instance-hour (or partial hour)	154 Hrs	67.76
Amazon EC2 Bandwidth		
\$0.100 per GB Internet Data Transfer - all data transfer into Amazon EC2	1,338 GB	0.13
\$0.170 per GB Internet Data Transfer - first 10 TB / month data transfer out of Amazon EC2	0.351 GB	0.06
Amazon EC2 EBS		
\$0.11 per GB-month of provisioned storage	22.742 GB-Mo	2.50
\$0.11 per 1 million I/O requests	778,528 IOs	0.09
Elastic IP Addresses		
\$0.01 per non-attached Elastic IP address per complete hour	269 Hrs	2.69
		73.23
View Usage Report		73.23
Amazon Simple Storage Service		
View/Edit Service		
United States		
\$0.170 per GB - first 10 TB / month data transfer out	0.000053 GB	0.01
\$0.01 per 1,000 PUT, COPY, POST, or LIST requests	10 Requests	0.01
\$0.01 per 10,000 GET and all other requests	4 Requests	0.01
		0.03
European Union		
\$0.18 per GB - first 50 TB / month of storage used	0.451 GB-Mo	0.08
\$0.170 per GB - first 10 TB / month data transfer out	0.000739 GB	0.01
\$0.012 per 1,000 PUT, COPY, POST, or LIST requests	78 Requests	0.01
\$0.012 per 10,000 GET and all other requests	90 Requests	0.01
		0.11
View Usage Report		0.14
Taxes		
Estimated Taxes		
(Due November 1, 2009)		
		0.00
Charges due on November 1, 2009†		73.37

14/02/2017

Satish Srirama

22/40

Getting started – for free

- <http://aws.amazon.com/free/> - free tier
- Valid for 12 months
- 750 hours / month of EC2 Linux Micro Instance usage
- 5 GB of S3 standard storage, 20,000 Get Requests, and 2,000 Put Requests
- 30 GB of Amazon EBS
- 1 GB Regional Data transfer
- Be careful
 - You cross the limits, you are charged without any notice

Note: All values get updated regularly

ElasticFox

- Mozilla Firefox extension for managing your Amazon EC2 account
- HybridFox is a Fork of ElasticFox
- Supported services
 - List available AMIs
 - List your running instances
 - Launch new instances of an AMI
 - Manage security groups and launch permissions associated with your instances
 - Manage Amazon EBS volumes


ElasticFox - GUI

The screenshot displays the ElasticFox web interface within a Mozilla Firefox browser. The browser's address bar shows the URL `chrome://ec2ui/content/ec2ui_main_window.xul`. The ElasticFox interface includes a navigation menu at the top with options like 'Instances', 'Images', 'KeyPairs', 'Security Groups', etc. The 'Images' section is active, showing a table of available images. The table has columns for ID, Manifest, State, Owner, Visibility, Architecture, Platform, Root Device Type, Name, Description, and Tag. The 'State' column for all images is 'available'. The 'Owner' column lists various owners including 'amazon', 'public', and '432018295444'. The 'Architecture' column is mostly 'i386', with some 'x86_64' entries. The 'Platform' column is mostly 'instance-store'. The 'Root Device Type' column is mostly 'instance-store', with some 'instance-store' entries. The 'Name' column is mostly empty, with some entries like 'ubuntu-kernels-...'. The 'Description' column is mostly empty, with some entries like 'ubuntu-kernels-...'. The 'Tag' column is mostly empty, with some entries like 'x86_64'. On the right side of the interface, there is a 'Launch Permissions' section with a list of permissions and a search bar.

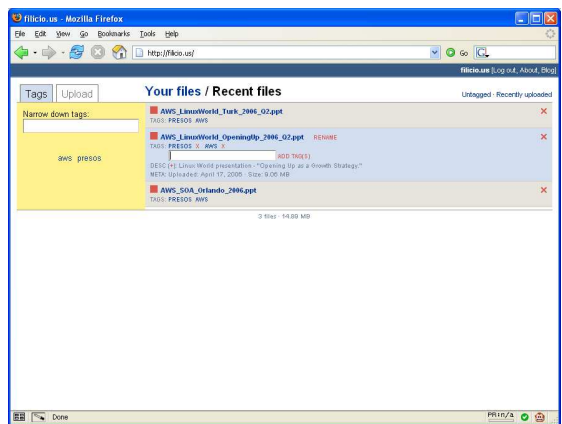
ID	Manifest	State	Owner	Visibility	Architecture	Platform	Root Device Type	Name	Description	Tag
aki-00806369	karmic-kernel-zu/ubuntu-kernel-2.6.31-3...	available	099720109477	public	i386		instance-store			
aki-00896a69	karmic-kernel-zu/ubuntu-kernel-2.6.31-3...	available	099720109477	public	i386		instance-store			
aki-008b6869	redhat-cloud/RHEL-5-Server/5.4/x86_64...	available	432018295444	public	x86_64		instance-store			
aki-00f41769	karmic-kernel-zu/ubuntu-kernel-2.6.31-3...	available	099720109477	public	i386		instance-store			
aki-010be668	ubuntu-kernels-milestone-us/ubuntu-luc...	available	099720109477	public	i386		instance-store			
aki-0251b36b	ec2-paid-ibm-images-ids/vmlinux-2.6.16.60...	available	amazon	public	x86_64		instance-store			
aki-03e9086a	canonical-cloud-us/vmlinux-2.6.28-12-xe...	available	099720109477	public	x86_64		instance-store			
aki-07c5266e	ubuntu-kernels-sandbox-us/ubuntu-karm...	available	099720109477	public	x86_64		instance-store			
aki-084dad61	canonical-cloud-test/vmlinux-chuck-2009...	available	099720109477	public	i386		instance-store			
aki-0942a160	ubuntu-kernels-testing-us/ubuntu-hardy...	available	099720109477	public	i386		instance-store			
aki-0a4aa863	ec2-paid-ibm-images/vmlinux-2.6.16.60...	available	amazon	public	x86_64		instance-store			
aki-0c8b6865	redhat-cloud/RHEL-5-Server/5.3/x86_64...	available	432018295444	public	x86_64		instance-store			
aki-0d42a164	ubuntu-kernels-testing-us/ubuntu-hardy...	available	099720109477	public	x86_64		instance-store			
aki-0d9f7b64	oracle_linux_kernels/2.6.18-53.1.13.9.1...	available	725966715235	public	x86_64		instance-store			
aki-0dc52664	ubuntu-kernels-sandbox-us/ubuntu-karm...	available	099720109477	public	i386		instance-store			
aki-0fbb5766	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	x86_64		instance-store			
aki-11c72878	ubuntu-kernels-testing-us/ubuntu-lucid-a...	available	099720109477	public	x86_64		instance-store	ubuntu-kernels-...		
aki-12f0127b	ec2-paid-ibm-images/vmlinux-2.6.16.60...	available	amazon	public	i386		instance-store			
aki-131df27a	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	x86_64		instance-store			
aki-13c2227a	karmic-kernel-zu/test-kernel-01.manifest...	available	099720109477	public	i386		instance-store			
aki-13dc3f7a	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	x86_64		instance-store			
aki-14b4597d	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	x86_64		instance-store			
aki-1521c17c	canonical-cloud-test/vmlinux-2.6.29-test...	available	099720109477	public	i386		instance-store			
aki-1568877c	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	x86_64		instance-store			
aki-162ac87f	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	i386		instance-store			
aki-174aae7e	oracle-corporation/oracle_linux_kernels/...	available	725966715235	public	i386		instance-store			
aki-1783627e	sun-opensolaris-2009-06/unix_32_7.0.m...	available	327216928991	public	i386		instance-store			
aki-1920c070	karmic-kernel-zu/test-kernel-01.manifest...	available	099720109477	public	i386		instance-store			
aki-1b3ada72	karmic-kernel-zu/test-kernel-01.manifest...	available	099720109477	public	i386		instance-store			
aki-1ddc3f74	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	i386		instance-store			
aki-1e638377	karmic-kernel-zu/ubuntu-kernel-2.6.31-3...	available	099720109477	public	i386		instance-store			
aki-1fcd2e76	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	x86_64		instance-store			
aki-204cac49	canonical-cloud-test/vmlinux-chuck-test...	available	099720109477	public	i386		instance-store			
aki-20c12649	canonical-beta-us/vmlinux-2.6.27-22-xe...	available	099720109477	public	i386		instance-store			
aki-21f01148	canonical-cloud-us/vmlinux-2.6.28-12-xe...	available	099720109477	public	i386		instance-store			
aki-24c3204d	karmic-kernel-zu/ubuntu-kernel-2.6.31-3...	available	099720109477	public	i386		instance-store			
aki-25de3b4c	redhat-cloud/RHEL-5-Server/5.1/i386/ke...	available	432018295444	public	i386		instance-store			
aki-266b894f	sles-beta-ibm/vmlinux-2.6.16.60-0.29-xe...	available	140642109189	public	i386		instance-store			
aki-27d9364e	ubuntu-images-sandbox/ubuntu-kernel-2...	available	099720109477	public	i386		instance-store			

Independent software vendors (ISV) for S3

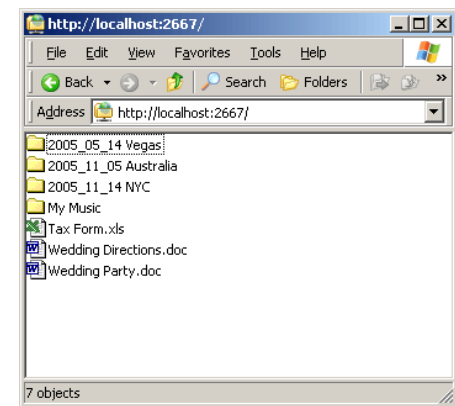
S3 Explorer



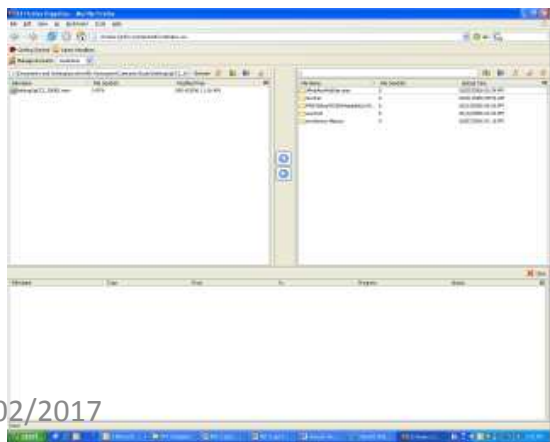
filicio.us



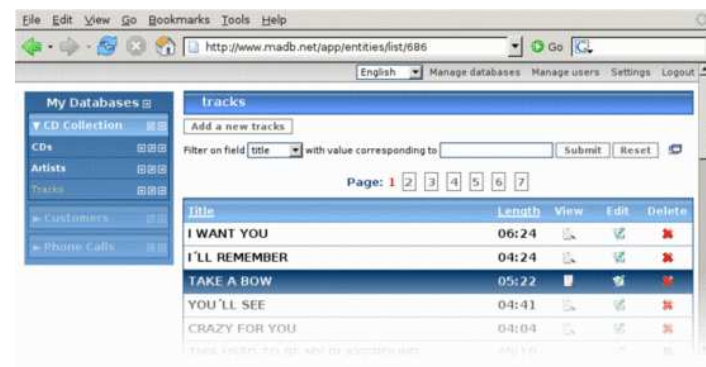
Jungle Disk



S3 Firefox Organizer



MyOwnDB



RightScale

- RightScale cloud management platforms
- Automated configuration
 - From boot to production on auto-pilot
 - Server templates & Right scripts
 - Cloning/ replication of set-ups
- Automated management
 - Monitoring, alerts, SLAs -> actions
- Have several examples
- Can login as a developer
<https://my.rightscale.com>

PRIVATE CLOUD ENABLING TECHNOLOGIES

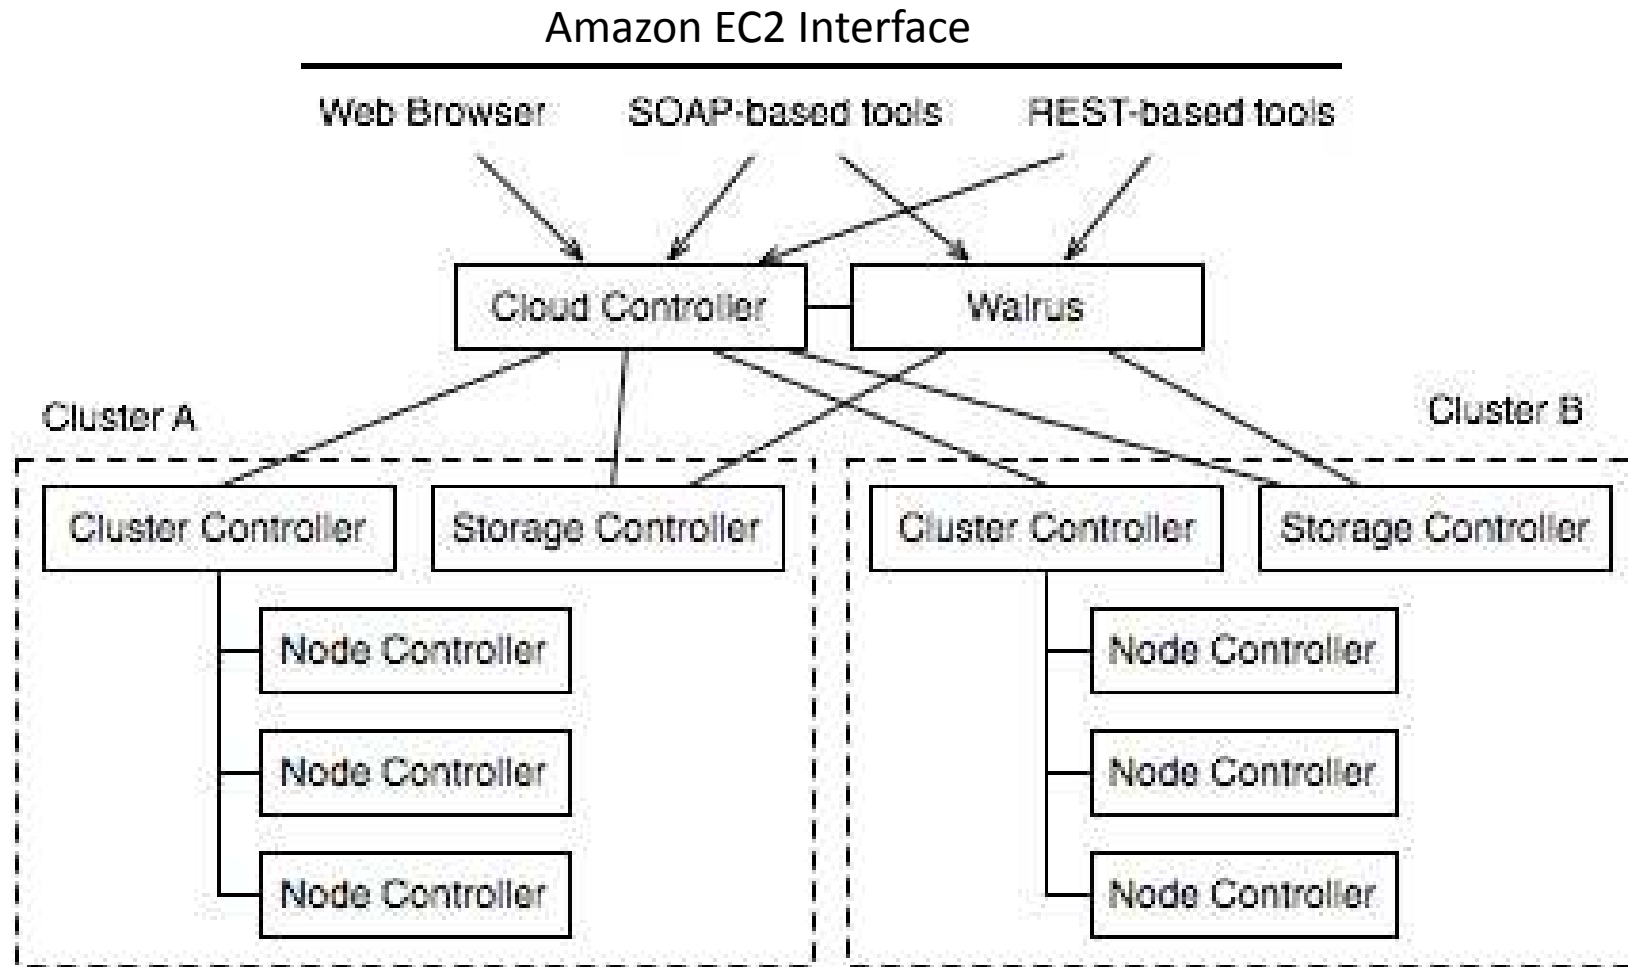
Eucalyptus

- Are you OK with using your credit card?
- Open source project
- Elastic Utility Computing Architecture Linking Your Programs To Useful Systems
- Web services based implementation of EC2 infrastructure
- For establishing private clouds
- Functions as a software overlay
 - Existing installation should not be violated (too much)
- Focus on installation and maintenance

Eucalyptus Usage

- Experimentation vehicle prior to buying commercial services
 - Provide development, debugging, and “tech preview” platform for Public Clouds
- Homogenize local IT environment with Public Clouds
 - AWS functionality locally makes moving & using Amazon AWS easier, cheaper, and more sustainable
- Provide a basic software development platform for the open source community
 - E.g. the “Linux Experience”
- **Not** designed as a replacement technology for AWS or any other Public Cloud service

Eucalyptus - architecture



Basic euca2ools

- euca-describe-images
- euca-describe-instances
- euca-run-instances
- euca-create-volume
- euca-attach-volume
- euca-terminate-instances
- euca-describe-availability-zones

OpenStack

- Founded by **NASA** and **Rackspace**
- The open source cloud computing platform
- Feature-rich and massively scalable
- Powers cloud storage, compute, and networking
- A world-wide open source collaboration

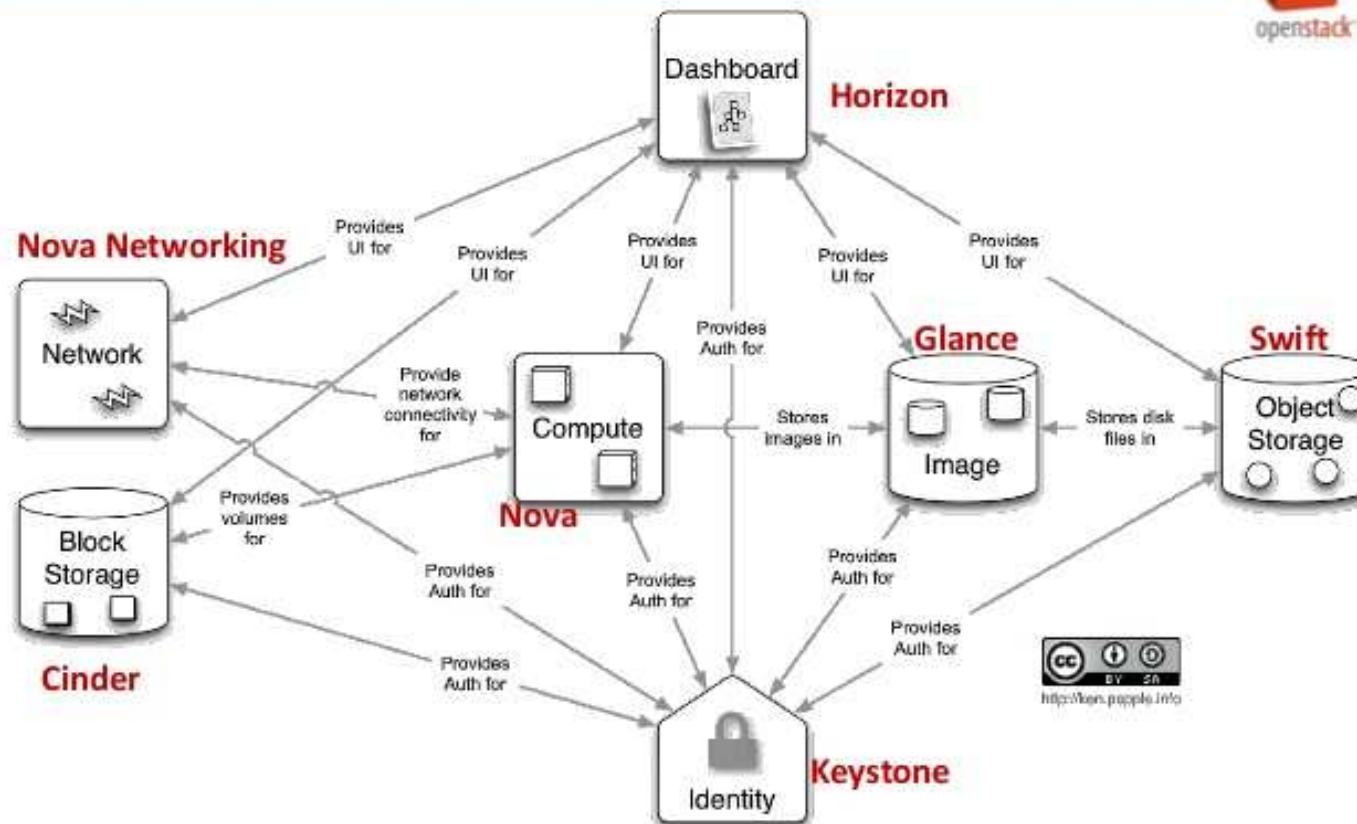
OpenStack has 105+ participating companies



...and over 1,350 individual contributors. [source: Piston]

OpenStack - High level overview

High Level Architecture - Grizzly



Adapted from : <http://25a0f8eada321387d14b7111c577a00b96c94eca71eac0173e073e077.com/openstack-conceptual-arch-tutorial.pdf>

10

© 2013 Cloud Technology Partners, Inc.

3 Major OpenStack Components

- Compute (“**Nova**”)
 - Orchestrates large networks of Virtual Machines.
 - Responsible for VM instance lifecycle, network management, and user access control.
- Object Storage (“**Swift**”)
 - Provides scalable, redundant, long-term storage for things like VM images, data archives, and multimedia.
- Image Service (“**Glance**”)
 - Manages VM disk images.
 - Can be a stand-alone service.
 - Supports private/public permissions, and can handle a variety of disk image formats.
- Other components: Dashboard, Load Balancing, Authentication...

Scientific Computing Cloud (SciCloud)

[Srirama et al, CCGrid 2010]

- Mobile & Cloud Lab owned private cloud infrastructure
- Collection of smaller clouds built using Eucalyptus/OpenStack platforms
- Goal of the project
 - To efficiently use the already existing resources of universities
 - To address computationally intensive scientific, mathematical, and academic problems

<http://mc.cs.ut.ee/mcsite/projects/scicloud>

Working with OpenStack

- Last week keys are created for you all
- You work with the OpenStack private cloud
- Just like Eucalyptus euca2ools are sufficient to invoke the OpenStack services

This week in lab

- You work with SciCloud
 - Creating instances
 - Creating security groups
 - Connecting to instances through ssh
 - Installing software
 - Preparing images and snapshots

Next Lecture

- Scale on cloud

References

- Amazon Web (Cloud) Services – documentation
<http://aws.amazon.com/documentation/>
- SciCloud – homepage
<http://mc.cs.ut.ee/mcsite/projects/scicloud>
- Eucalyptus User's Guide
<http://www.eucalyptus.com/docs/3.1/ug-3.1.0.pdf>
- OpenStack <https://www.openstack.org/>
- S. N. Srirama, O. Batrashev, E. Vainikko: [SciCloud: Scientific Computing on the Cloud](#), 10th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2010), May 17-20, 2010, pp. 579. IEEE CS.