

CS 683 Emerging Technologies
Fall Semester, 2008
Doc 12 EC2
Oct 14 2008

Copyright ©, All rights reserved. 2008 SDSU & Roger Whitney, 5500
Campanile Drive, San Diego, CA 92182-7700 USA. OpenContent ([http://
www.opencontent.org/openpub/](http://www.opencontent.org/openpub/)) license defines the copyright on this
document.

References

Amazon Elastic Compute Cloud Developer Guide API Version 2008-05-05, <http://developer.amazonwebservices.com/connect/entry.jspa?externalID=1665&categoryID=87>

Comparison of Virtual Machines, http://en.wikipedia.org/wiki/Comparison_of_virtual_machines

Virtualization

Virtual machine emulates hardware

Allows running any software that runs on emulated hardware

VMs supporting Languages

- Java VM

- Erlang VM

VMs supporting operating systems

- Parallels Desktop for Mac

- Microsoft Virtual PC

- VMware

- Wine

- Mac-on-Linux

Paravirtualization

Operating system is modified to be aware of the fact it is virtualized.

What is Elastic Compute Cloud (EC2)?

Amazon Machine Image (AMI)

Linux virtual machine image

Copy of the OS, applications & file system

Soon to support Windows

Instances

Running machine images

Can be run on different machines

What is Elastic Compute Cloud (EC2)?

Amazon has server farms running VMs

You can run AMIs on those machines

Charged per usage

Instances have access to

local files

Network

S3

Simple Queue Service (SQS)

Sending messages between machines

EC2 Instance Types

Type	Memory	Cores	Storage	Platform	Cost/hour
Small	1.7 GB	1	160 GB	32-bit	\$0.10
Large	7.5 GB	2	850 GB	64-bit	\$0.40
Extra Large	15 GB	4	1,690 GB	64-bit	\$0.80
High-CPU Medium	1.7 GB	2	350 GB	32-bit	\$0.20
High-CPU Extra Large	7 GB	8	1,690 GB	64-bit	\$0.80

Creating AMI

Start with Existing AMI and modify it

<http://docs.amazonwebservices.com/AWSEC2/2008-05-05/GettingStartedGuide/>

Full install of the OS

See Amazon Elastic Compute Cloud Developer Guide

Detailed operations

Need certificates for security

Lots of steps

Instance Network Access

Each instance on start up gets

- Private network address & DNS name

 - Valid only on EC2 network

 - High bandwidth, low latency

 - For communication between instances

- Public network address & DNS name

 - Publicly accessible

 - DNS names take up to 24 hours to propagate

Both types are valid only when instance is running

Address are not persistent across restarting an instance

Elastic IP Addresses

Elastic IP addresses are assigned to EC2 accounts

Can dynamically map it to an EC2 instance

Allows you to upgrade/downgrade the instance at the address

Elastic Block Store (EBS)

Volume of block storage

Stored in S3

Can be mounted/unmounted to an EC2 instance

Up to one TB of storage

Can make snapshots

Elastic Block Store (EBS) usage

Store data

Back up data in case of machine failure

Start multiple instances that have copies of same data

Store data while instance is not running

How does this all work?

AMI instance is just a Linux server running

- Apache

- Database

- Any program that can run on Linux

Use EBS for data storage

- Backups

- Copying data between instances

Communicate between AMI instances via

- Internet protocols (tcp/ip etc)

- Amazon Simple Queue Service

What about Scaling?

Starting point - one AMI instance on one core

Apache, application & database

Instance attached to elastic IP address

Data stored on EBS

When need more processing power

Run copy of AMI instance on up to 8 cores

Snapshot EBS and attach to new instance

Point elastic IP address to new instance

More Scaling

Separate parts in to different AMI instances

Apache

- one instance

- Front end

- Load balances requests

Application

- Run application instance on N machines

Database

- Run on separate instance