

Introduction to Cloud Computing on Amazon Web Services (AWS) with focus on EC2 and S3

Horst Lueck

cloud@horstlueck.com

2011-05-17


IT Pro Forum

http://itproforum.org

Thanks to Open Office Impress

The Cloud – the Name

The 90s



Today



The Cloud – the Name (cont.)

- Computing as a Utility Service
 - “Pay by the drink”
- Grid Computing – more descriptive but sounds less sexy
 - However, serious “grids” in science: globus.org [4] [5]
- Cloud sounds more sexy, so it wins
 - Term is more marketable – and “hype-able”
 - Fuzziness helps with acceptance by non-technical people

The Cloud – the Hype

- “The clouded world of naming cloud startups” [9]
- Google filter “inurl:cloud” 19 Mill. Hits – many real clouds, but top 5:
 - www.ubuntu.com/business/cloud/overview (Ubuntu on AWS)



Concerns

- Lack of control and ownership of physical data.
 - Richard Stallman "Careless Computing" [6] [7]
- Third party/government monitoring, all the way to loss of data and services, think WikiLeaks
- Service Provider disappears, e.g. Coghead [8]
- Σ : Keep important stuff backed up on physical media under your control

Why Choosing Amazon Web Services ?

(Author's personal reasons, Dec 2010)

- Oldest (2006) and best established
- Free Tier for 1 year (1 micro instance, 750 hr/month) [2]
- Hibernate instance w/o pay (except min. storage costs)
- Pay only for what you use
 - service/hr, storage/month, data transfer \$/GB
- Support
 - REST (HTTP/1.1) and SOAP (XML) interfaces
 - Tools and libraries in Java, PHP, Python...
 - Working examples to get started with most services
 - Good documentations, html and .pdf
 - Forums

AWS Virtual Machines, “Instances”

- Micro Instance "t1.micro" (free for Free Tier)
 - 633 MB memory, EBS storage only (8 GB AMI),
 - 32/64-bit platform, I/O Performance: Low
- Small Instance "m1.small" (default)
 - 1.7 GB memory, 160 GB instance storage,
 - 32-bit platform, I/O Performance: Moderate
- ... skipping...
- High-Memory Quadruple Extra Large Instance "m2.4xlarge"
 - 68.4 GB of memory, 1690 GB of instance storage, 64-bit platform, I/O Performance: High
 - 26 EC2 Compute Units (8 virtual cores with 3.25 EC2 Compute Units each)

Sign up for Free Tier

- To get started see ref. [2]
- Required:
 - Credit card
 - Address, (phone?)
 - Email, used for
 - Login (instead of AWS ID)
 - Contact
 - Newsletter (good tips)
 - Billing info

AWS Management Console

(looking at live console at presentation)

- Elastic Beanstalk -- quick deployment of applications
- Simple Storage Service (S3)
- Elastic Compute Cloud (EC2)
- Virtual Private Cloud (VPC)
- CloudWatch – Monitoring, Alarms, Auto Scaling
- Elastic MapReduce – manage work flow (S3, EC2)
- CloudFront – content delivery network / S3 => www
- CloudFormation – create server stack / JSON templates
- Relational Database Service (RDS), currently MySQL
- Simple Notification Service (SNS), beta
- Identity and Access Management (IAM), mng groups, users

- AWS Elastic Beanstalk
- Amazon S3
- Amazon EC2**
- Amazon VPC
- Amazon CloudWatch
- Amazon Elastic MapReduce
- Amazon CloudFront
- AWS CloudFormation
- Amazon RDS
- Amazon SNS
- AWS IAM

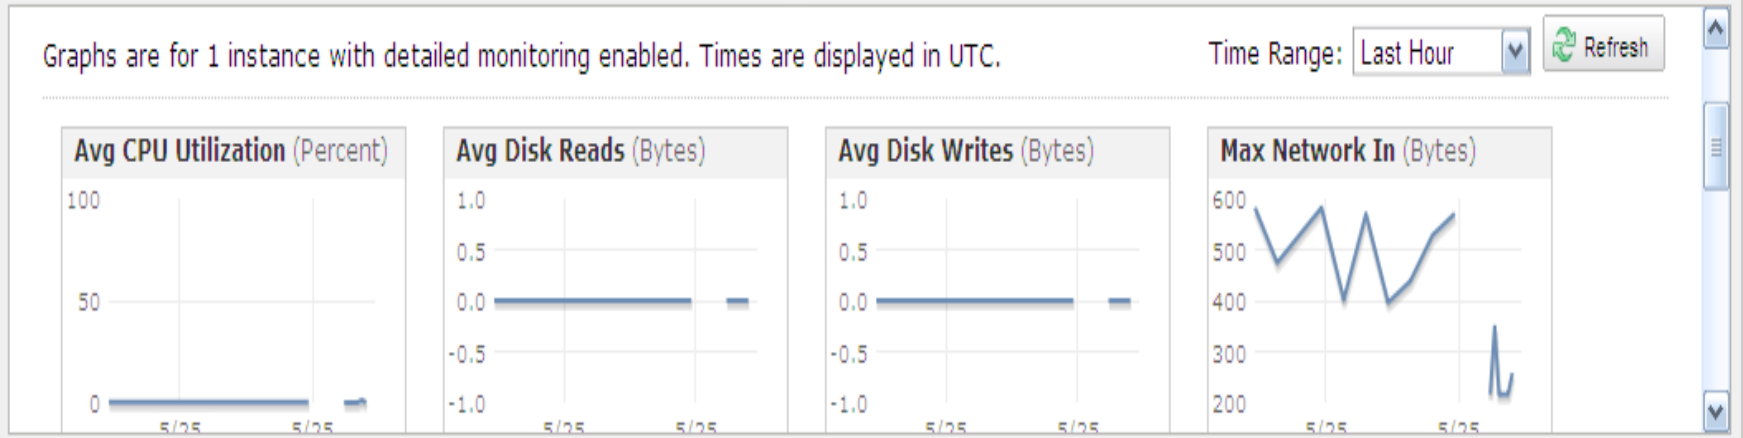
- Navigation**
- Region:**
US East (Virginia)
- > EC2 Dashboard
 - INSTANCES
 - > Instances
 - > Spot Requests
 - > Reserved Instances
 - IMAGES
 - > AMIs
 - > Bundle Tasks
 - ELASTIC BLOCK STORE
 - > Volumes
 - > Snapshots
 - NETWORKING & SECURITY
 - > Security Groups
 - > Elastic IPs
 - > Placement Groups
 - > Load Balancers
 - > Key Pairs

My Instances

Viewing:

1 to 9 of 9 Instances

	Name	Instance	AMI ID	Root Device	Type	Status	Security Groups	Key Pair Name	Monitoring	Vir
<input type="checkbox"/>	ubu_NIC	i-444e4629	ami-00e91f69	ebs	t1.micro	running	quick-start-1	HLaws1	basic	
<input type="checkbox"/>	HL-wiki	i-c9859ca5	ami-2613e04f	ebs	m1.small	running	HL-wiki-EC2Securit	HL_demo	basic	
<input checked="" type="checkbox"/>	Amaz_RH_itpro-1	i-c1ce3caf	ami-80ec12e9	ebs	t1.micro	running	itpro	HL_demo	<input checked="" type="checkbox"/> detailed	
<input type="checkbox"/>	ubu_2	i-31e7165d	ami-5cae5f35	ebs	t1.micro	stopped	quick-start-1	HLaws1	basic	
<input type="checkbox"/>	Fedora13 LAMP	i-155b1579	ami-08728261	ebs	t1.micro	stopped	quick-start-1	HLaws1	basic	
<input type="checkbox"/>	ubu_LAMP_itpro	i-b85331d7	ami-028f7b6b	ebs	t1.micro	stopped	itpro	itpro	basic	
<input type="checkbox"/>	Amaz_RH_20apr	i-72f1be1d	ami-8c1feca5	ebs	t1.micro	stopped	quick-start-1	HLaws1	basic	
<input type="checkbox"/>	Amaz_RH_itpro-3	i-ff45bb91	ami-80ec12e9	ebs	t1.micro	terminated	itpro	itpro	basic	



Create a new server instance on EC2


(preview of hands-on part)

- EC2 Dashboard / Launch Instance: "All" ... "Owned by me"
- Select an AMI (Amazon Machine Image)
 - public, private (=yours), re\$erved, stored on S3, etc...
 - Amazon Linux, LNX distros like Ubuntu, Suse, RHEL, Windows...
 - Must have valid AKI (Amazon Kernel Image)
- Instance Details
 - Type (t1.micro), Availability Zone, Number of Instances
 - Naming (matching Names of EBS, AMI... helps w/ housekeeping !)
- Chose or create Key Pair: 'itpro'
 - At Create, download opens for 'itpro.pem' SAFE IT !!!
- Chose or create "Security Group": 'itpro'
 - "Inbound" Select from drop-down: SSH HTTP HTTPS

Request Instances Wizard



Please review the information below, then click **Launch**.

AMI:  Other Linux AMI ID ami-d4fc02bd (i386) [Edit AMI](#)

Number of Instances: 1
Availability Zone: No Preference
Instance Type: Micro (t1.micro)
Instance Class: On Demand [Edit Instance Details](#)

Monitoring: Disabled **Termination Protection:** Disabled
Tenancy: Default
Kernel ID: Use Default **Shutdown Behavior:** Stop
RAM Disk ID: Use Default
User Data: [Edit Advanced Details](#)

Key Pair Name: itpro [Edit Key Pair](#)

Security Group(s): sg-48fd9d21 [Edit Firewall](#)

Elements of an Instance – Summary

- AMI (Amazon Machine Image)
 - Type(size, etc.), Availability Zone,
 - AKI (Amazon Kernel Image)
- Key Pair -- for 1st login (default: password login is disabled)
- Security Group -- listening ports for inbound traffic
- EBS (Elastic Block Store) -- e.g. instance-ID:/dev/sda1
- Snapshot -- for speedy start-up on EBS
 - via `CreateImage(instance-ID)` for ami-ID from vol-ID_EBS
- Optional:
 - Elastic IP, DNS
 - Monitoring details, etc.

S3 – Simple Storage Service

- Container based – using bucket-, folder-, and file objects
 - NOT a traditional file system
- Complex, customizable access privileges, using
 - ARN (Amazon Resource Name) – may include time zones
 - Simple example: "Resource":["arn:aws:s3:::bucket/itpro"]
 - Rules and templates in JSON format
- Each object can be made public through unique URL
 - either as a web object (normal browser view)
 - or as XML (bucket/folders, or error message)
- Authentication with Access Keys (key-ID:'Secret' value)
 - used in config files (e.g. s3cmd) or as Unix environment variables

Working with S3

- Command line tool 's3cmd', part of S3tools project [10]
 - `s3cmd ls s3://itpro-eug`
 - `s3cmd sync -P itproforum.org s3://itpro-eug`
- View as bucket
 - `https://s3.amazonaws.com/itpro-eug/itproforum.org/index.html`
 - Format: `https, s3-domain/bucket/folder/file`
- View as static website (provide index and error file `.html .txt / not .cgi`)
 - `http://itpro-eug.s3-website-us-east-1.amazonaws.com/itproforum.org/index.html`
 - Format: `http, bucket.s3-zone-domain/folder/file`

Begin of Hands-on Part – EC2

- Presenter and audience communicate through a wiki:
 - <http://50.19.95.82/itpro>
 - (Gollum wiki, started through CloudFormation server stack)
- ElasticIPs were allocated just for this presentation
 - The IPs will be released and won't have any meaning in the future
- Instance “Amaz_RH_itpro(#)” is running – (#) TBA at presentation
 - Serving <http://50.19.103.131/itpro/>
 - It's a modified mirror image of <http://itproforum.org/>
 - Next we'll update our “mirror” server ...

itpro



"Group therapy for people doing real work with computers."

FIRST HOST!!! <<< (some user's input at the presentation:-)

<http://50.19.95.82/itpro> -- This page: Gollum wiki server

For communication between presenter and audience.

<http://50.19.103.131/itpro/> -- 'mirror' website of itproforum.org (demo)

This is where " wget -m -nv http://itproforum.org/" and "s3cmd sync itproforum.org s3://itpro-eug" will be run at the presentation.

"Mirror" web site is also sync'ed to S3:

"s3cmd sync itproforum.org s3://itpro-eug"

a) View as bucket <https://s3.amazonaws.com/itpro-eug/itproforum.org/index.html>

Both 'good links' (like existing folders), and 'bad links' (like invalid URLs) generate XML response

b) View as static page <http://itpro-eug.s3-website-us-east-1.amazonaws.com/itproforum.org/index.html>

This view can serve a baseURL/index.* (.html, .txt etc. but nothing dynamic, like .cgi) ,

plus a customized error page to users entering a bad URL

We'll finish with a video showing Elvis (and the author) at Addi's in Springfield

<http://50.19.103.131/Elvis/> -- credits to local TV station KVAL

Very cool -- but why is this relevant to AWS ?

Well, look at the embedded video feed:

file=http%3A%2F%2Fkidkbim.s3.amazonaws.com%2FElvis-1301084730.she.mp4

which translates to:

<http://kidkbim.s3.amazonaws.com/elvis-1301084730.she.mp4>

Note to Kay: Even as an app/web developer you want to understand the basics of AWS/S3 :-)

Hands-on Part – EC2/S3 (cont.)

- Updating EC2 webserver
 - `ssh -i ~/path/to/itpro.pem ec2-user@50.19.103.131`
 - `cd /var/www/html/`
 - `wget -m -nv http://itproforum.org/` // visitors watching live changes
- Propagate changes to static mirror on S3, using s3cmd tools:
 - `s3cmd sync itproforum.org s3://itpro-eug`
 - visitors watching live changes:
 - `https://s3.amazonaws.com/itpro-eug/itproforum.org/index.html`
 - `http://itpro-eug.s3-website-us-east-1.amazonaws.com/itproforum.org/index.html`
- Command line tools <http://awsdocs.s3.amazonaws.com/EC2/latest/ec2-clt.pdf>
 - Over 100 tools in `/opt/aws/bin/ec2-*`

References (sequence may be different from contents, but numbers match)

[1] http://en.wikipedia.org/wiki/Cloud_computing

[1b] http://en.wikipedia.org/wiki/Amazon_Web_Services – Overview, many links

[1c] http://en.wikipedia.org/wiki/Amazon_Elastic_Compute_Cloud – EC2

[2] <http://aws.amazon.com/free/> – AWS Free Usage Tier

[3] <http://cloud.ubuntu.com/ami/>

[4] <http://globus.org/>

[5] <http://www.globus.org/demogrid/>

[6] <http://www.guardian.co.uk/technology/2008/sep/29/cloud.computing.richard.stallman>

[7] <http://techcrunch.com/2010/12/14/stallman-cloud-computing-careless-computing/>

[8] <http://www.infoworld.com/d/cloud-computing/what-do-if-your-cloud-provider-disappears-508?page=0,3>

[9] <http://www.jackofallclouds.com/2011/02/naming-cloud-startups/>

[10] <http://s3tools.org/s3cmd>

Credits for 3rd party images:

wikipedia and images.google.com