

Amazon EC2: Cheaper than Tinkertoys™

Amazon Web Services (AWS), especially the Elastic Compute Cloud (EC2), are so inexpensive in the small scale that they lend themselves to easy ad-hoc use and experimentation.

Warning: Amazon may be evil !

Q (no &A)

Rapid Demo

You will need:

- Email account
- Major Credit Card
- Telephone (for ringback account activation)
- Computer with SSH (to connect to Unix) or Remote Desktop (to connect to Windows)

You will create or install:

- Amazon account (existing book acct OK)
- EC2 API Tools

Also helpful:

- Firefox browser with ElasticFox extension
- Perl modules
Net::Amazon
Net::Amazon::EC2

Amazon Web Services

- Compute
 - Elastic Compute Cloud (EC2)
 - Elastic MapReduce (Hadoop)
 - Auto Scaling
- Content Delivery
 - CloudFront
- Database
 - SimpleDB
 - Relational Database Service (RDS)
- E-Commerce
 - Fulfillment Web Service (FWS)
- Messaging
 - Simple Queue Service (SQS)
- Monitoring
 - CloudWatch
- Networking
 - Virtual Private Cloud (VPC)
 - Elastic Load Balancing
 - Elastic IP Addressing
- Payments & Billing
 - Flexible Payments Service (FPS)
 - DevPay
- Storage
 - Simple Storage Service (S3)
 - Elastic Block Storage (EBS)
 - AWS Import/Export
- Workforce
 - Mechanical Turk

AWS use cases

- Application Hosting
- Backup Storage
- Content Delivery
- E-Commerce Applications
- Enterprise IT
- HPC Applications
- Media Hosting
- Search Engines
- Web Hosting
- iPhone App Hosting
- Facebook App Hosting
- Bit-Torrent Seeding
- **Toybox**

New York Times

- Archives from 1851 - 1980
- 11 million articles, 4TB of TIFF files.
- "Months" to convert to PDF in-house
- Requested limit increase; 20 instances -> 100
- <24 hours, <\$300

Past This Point ...

- Things will now get more complex, but not much more expensive. Remember how easy the demo was!

Professor XXXL



Instance Types

Name	GB.RAM	Cores	CUs	GB.Disk	Bits	IO	Unix	Windows
m1.small	1.7	1	1	160	32	Moderate	\$0.085	\$0.12
m1.large	7.5	2	2	850	64	High	\$0.34	\$0.48
m1.xlarge	15.0	4	8	1690	64	High	\$0.68	\$0.96
m2.xlarge	17.1	2	6.5	420	64	High	\$0.50	\$0.62
m2.2xlarge	34.2	4	13	850	64	High	\$1.20	\$1.44
m2.4xlarge	68.4	8	26	1690	64	High	\$2.40	\$2.88
c1.medium	1.7	2	5	350	32	Moderate	\$0.17	\$0.29
c1.xlarge	7.0	8	20	1690	64	High	\$0.68	\$1.16

m2 = High-Memory
c1 = High-CPU

CUs = EC2 Compute Units
Prices are for "us-east" region

Instance Request Types

On-Demand

- Normal; Pay-as-you-go; zero commitment
- This is the type we will be mainly discussing

Reserved

- Pay a 1 year or 3 year lump sum, and get a box guaranteed, with a deeply discounted hourly charge.
- Lose flexibility on instance types.
- Break-even point for 1 year payment is about 11 hours every day.

Spot

- You set a maximum "bid" rate, but you pay the "spot" rate, based on current demand.
- Your instance runs while the spot is lower than your bid.
- Has been averaging about 1/3 of On-Demand prices for every instance types.

Chargeable Resources

- Instance
 - Charged per hour, rounded up;
varies with CPU/RAM/Disk
- Bandwidth
 - \$0.15/GB transfer out (\$<10TB)
 - \$0.10/GB transfer in (Free until 2010-06-30)
 - Free in&out within a Region
- CloudWatch monitoring (very optional)
 - \$0.015 per hour
- Elastic IP addresses (optional)
 - \$0 while in use
 - \$0.01 per unused address per *hour*
 - \$0.10 per 100 remaps
- Disk: EBS
 - \$0.10/GB/month
 - \$0.10 per million I/O requests
- Disk: S3
 - \$0.15/GB/month storage (\$<50TB)
 - \$0.01 per 1,000 write reqs
\$0.01 per 10,000 read reqs

S3 vs EBS

	S3: Simple Storage Service	EBS: Elastic Block Storage
Storage Cost	\$0.15 / GB / month	\$0.10 / GB / month
I/O Request Cost	\$0.10/\$0.01 per 10,000 w/r	\$0.10 per million requests
Data Durability	Guaranteed	AFR 0.1% – 0.5% for 20GB
Simultaneous Attachment	Unlimited	Single instance at a time
Storage Chunk Size	1 Byte	1 GiB
Maximum Size	5GB per object	1TB per volume, 20+ vols
Format	Amazon YDN TK	YourFS
Connection	API, BT, HTTP, unofficial s3fs	Mounted on EC2 instance
Storage Base	File	Block
File List/Search	Slow	Fast
Resize	Fast	Slow
Latency	High and Variable	Low and Consistent
User Controlled Backup	Versioning in S3	Snapshotting to S3
Atomic	Yes (but not to mirrors)	No (but very low latency)

S3 Tricks

- Temp URL for secured download
- Reverse billing for bandwidth
- Domain mapping

Initial Limits

- SMTP: 500 messages/day
- Elastic IPs: 5
- Instance count: 20
- EBS volumes: 20
- EBS snapshots: 500



Where some only see Unpleasantness...




... others see Opportunity

Windows

- Windows Server 200[38]
- m1.small is too small?
- \$+
- \$\$\$ for MS extras like SQL Server
- EBS boot recommended
- Halt, instead of Terminate; reduce boot time
- Remote Desktop, instead of SSH
- ec2-get-password, instead of keypair login

Regions and Availability Zones



Virginia, California, Ireland

[http://docs.amazonwebservices.com/
AWSEC2/latest/DeveloperGuide/index.html?
concepts-regions-availability-zones.html](http://docs.amazonwebservices.com/AWSEC2/latest/DeveloperGuide/index.html?concepts-regions-availability-zones.html)

Misc

- OpenSolaris supported!
- You can make a custom AMI for yourself, or to charge others to run (stacks, captive software)
- There's an app for that: EC2 and S3
- My Perl launch code is coming to GitHub
- EC2 value-added companies: RightScale, etc
- Boot trick: -F with #! will usually auto-run

Other types of clouds

Distributed Computing

- SETI@Home, Hadoop, BitTorrent

Software as a Service (SaaS)

- Yahoo email, Google apps.
- Application Service Providers (ASP) were the first SaaS providers.

Platform as a Service (PaaS)

- Google App Engine, Force.com, Engine Yard (Ruby on Rails)

Infrastructure as a Service (IaaS)

- Amazon EC2, Joyent, GoGrid, FlexiScale

See also

Linux Journal Issue 191, March 2010, p72-75

- Article title: *Running Ubuntu 9.10 under Amazon's Elastic Cloud*
- Given as a handout in Atlanta;
not available at <http://LinuxJournal.com/>

<http://paulstamatiou.com/how-to-getting-started-with-amazon-ec2>

- Best online starter article

<http://alestic.com/>

- Ubuntu and Debian AMIs for Amazon EC2

<http://aws.amazon.com/ec2>

- EC2 Signup

<http://docs.amazonwebservices.com/AWSEC2/latest/GettingStartedGuide>

- Amazon EC2 Getting Started Guide

<http://clouddb.info/2009/02/>

- Series of postings on different types of cloud computing
(Distributed Computing, SaaS, IaaS, PaaS, Cloud Desktop, etc)

Q&A

Changes

- Not done yet: GitHub code upload, art credits
- v 1.01 2010-03-31 Typos, Added "Changes" page, WSJ corrected to NYT, Amazon Phonecall Verification added
- v 1.00 2010-03-25 Initial version; shown to Atlanta Perlmongers; uploaded to slideshare.net/Util; slideshare failed to convert .key to .pdf.