



Amazon Web Services

Virtual Servers & Cloud Computing

Paradigm Shifts

- The Mainframe (1950s~1980s)
 - Dumb Terminals
 - Data on the server
 - High CapEx
- Client-Server (1980s~present)
 - Powerful Clients
 - Data on the client+server
 - Low CapEx
- Cloud Computing (2000s~present)
 - Client could be anything with an internet connection
 - Data on the server
 - No CapEx, everything is rented

Head in the Cloud

- Discrete Parts
 - Database
 - Web Server
 - Reports Server
- Engineer Atomic Parts
- Availability, Scalability and Redundancy
- Detect Load > Increase Capacity
 - Reactive rather than Proactive
 - Keeps costs at the minimum
 - No wasted computation or storage

Your Options

- The Rackspace Cloud (formerly Mosso)
- Azure Services Platform
- Joyent and the other small fry
- Amazon Web Services

A many-headed beast

- Elastic Compute Cloud (EC2)
- Simple Storage Service (S3)
- Simple Queueing Service (SQS)
- Simple Notification Service (SNS)
- Elastic MapReduce
- Elastic Load Balancing (ELB)
- Simple DB
- Relational Database Service (RDS)
- CloudFront

Regions

- US - East Coast
- US - West Coast
- EU
- Singapore

EC2

- A virtual server
- Pay by HOUR of usage
- Windows or Linux or Anything else you want
- Boots up in about 5 minutes from scratch, faster from a backup
- Manage everything via the API
- *Boot up a server*

EC2: Size Matters

- Small - 32-bit, 1 Core, 1.7 GB
- Large - 64-bit, 4 Cores, 7.5 GB
- Extra Large - 64-bit, 8 Cores, 15 GB
- Double Extra Large - 64-bit, 13 Cores, 34.2 GB
- Quadruple Extra Large - 64-bit, 26 Cores, 68.4 GB
- Many more types

S3 + CloudFront

- Lots and lots of Storage
- Store files from 1 byte to 5 GB
- Unlimited Files = Unlimited Storage
- Scales Automatically
- CloudFront is an instant Content Distribution Network that sits on top of S3
- Edge Locations all over the world
- Incredibly fast delivery of content

S3: Buckets

- Files are stored in Buckets
- Buckets have users
- You can set access privileges on a per-file basis
- You can set access to public
- Many FTP clients now support S3
- Your entire application doesn't need to be on Amazon to use S3
- *Show S3 on Transmit*

SQS

- Queue Messages
- Messages can contain up to 8KB of text
- Messages get locked when being read
- Queues can be shared with anyone securely

SNS

- Centers around Topics
- Can set Access Rights for Topics
- Applications “subscribe” to topics
- Anyone can send notifications (ACL permitting)
- When a notification is received it is sent to subscribers
- Subscribers set “end-points”
- End-points can be email, a URL or SQS
- Designed to work with AWS services

Elastic MapReduce

- Super scalable data processing
- Uses Hadoop
- Ideal for indexing, log analysis, data compression, business intelligence, etc.
- Build your data processing application in Pig or Hive
- Alternatively if it's more complicated use Cascading, Java, Ruby, Perl, Python, PHP, R, or C++

Elastic MapReduce: Jobs

- Upload your data and your app to S3
- Fire an Elastic MapReduce job
- Monitor the job through the API or the console
- Output goes back to S3
- Amazon handles scaling and instances as necessary
- No extra hours of usage
- Unbelievably fast and cheap

ELB

- Availability Zones
 - Each Amazon Region has multiple Availability Zones
 - Each Zone is independent
 - BC + DR
- Load Balances requests between instances and over availability zones
- Supports Session Stickiness

Simple DB

- Very Simple Database Storage
- Think of it as a giant Excel sheet in the sky
- Super Scalable
- Uses SQL-like Syntax, doesn't support a lot of the keywords
- No relational data
- No strong table definition
- All fields are strings, can internally query CSVs
- Quick and dirty
- Gets the job done

RDS

- Managed MySQL Instances
- Can automatically replicate for failover
ACROSS availability zones
- Adding read-slaves shortly

Elastic Everything

- Elastic IPs can be provisioned directly and associated with EC2 instances
- Security is built-in with an IP firewall
- PEMs and PPKs
 - Privacy Enhanced Mail (X.509 Security Certificate) - *NIX
 - Putty Private Key - Windows

Tools

- Elastic Fox
- CloudBerry Explorer
- No good tools for Simple DB
 - None I'd recommend, anyway
- Install the APIs from Amazon
- CLI is strong for administering Linux machines
- Great for building scripts to automate
- Putty and WinSCP are your friends
- Get comfortable with the lingo and the CLI

When do I use this stuff?

- EC2
 - Web Servers
 - Application Servers
 - MS SQL Servers
- RDS
 - MySQL Servers
- Simple DB
 - Transaction Tables
 - Log Tables
- S3
 - Static Files
 - Streaming Content (+CloudFront)



Thanks!

for links and a copy of this presentation go to

<http://vizkid.com/extentia>