

S3 and EC2

Brian Aker
Director of Technology

October 2007

MySQL AB

I Am A Stalker



Who I Stalk

- Web 2.0 Companies
 - Technorati, MySpace, LinkedIn, Livejournal, Facebook
- Ebay
- Yahoo
- Amazon
- Google

Once Upon A Time



iLike

“The implication is, in my view, quite clear -- the Facebook Platform is primarily for use by either big companies, or venture-backed startups with the funding and capability to handle the slightly insane scale requirements. Individual developers are going to have a very hard time taking advantage of it in useful ways.” -Marc Andreeseen

**A Purchased hard Drive costs the
the same whether full or not**

Amazon Simple Storage Service

S3

- It is HTTP (REST even...)
 - GET
 - PUT
 - DELETE
- Security (it does SSL)
- A sample URL is:
 - <http://awsamazon.com/pictures.tangent.org/myimage.jpg>
- Some HTTP meta data is stored on PUT
 - Content Length (done automatically)
 - Content Type (program occtet unless you set it)

S3 Limitations

- 100 Buckets per account
- 1 Byte to 5 Gigabytes
- Billions of Items
- Hundreds of Petabytes (Smugmugg is at 200)

**A purchased Computer Costs
the same whether on or off**

Elastic Computing Cloud

- Virtual Images on Demand
 - Images are called ‘AMI’
 - You stores these in S3
 - Yes, this is Xen
 - A few bits have been shifted in the Xen image to make an AMI, Fedora ships with tools to turn Xen images to AMI, and vice Versa

Firefox File Edit View History Bookmarks Tools Window Help
 EC2 Firefox UI

chrome://ec2ui/content/ec2ui_main_window.xul

Getting Started Latest Headlines Pedigree View For Br...

Stumble! I like it! Send to Channel: All Favorites Friends Tools

Google Search PageRank AutoLink AutoFill Send to Settings

EC2 Firefox UI

Credentials Test Account IDs About

AMIs and Instances KeyPairs Security Groups

Available AMIs

AMI ID	Manifest	State	Owner	Visibility
ami-3e836657	datawrangling-images/fc6-python-mpi-node...	available	027811143419	public
ami-befe1bd7	level22-ec2-images/ubuntu-feisty-base-2007...	available	063491364108	public
ami-60fd1809	aes-images/pyami.manifest.xml	available	084307701560	public
ami-2eba5f47	aes-images/convertvideo.manifest.xml	available	084307701560	public
ami-6ee70207	virtualmin-gpl/image.manifest.xml	available	093590521311	public
ami-4ff71226	virtualmin-gpl-rpm/image.manifest.xml	available	093590521311	public
ami-a58a6fcc	rbuilder-online/ec2-suse-0.0.1-x86_10585.im...	available	099034111737	public
ami-2a947143	rbuilder-online/sigtest2-1.0.0-x86_10734.img...	available	099034111737	public
ami-e795708e	rbuilder-online/sugarcrm-os-jp-4.5.0h-x86_1...	available	099034111737	public
ami-b1896cd8	rbuilder-online/freedomail-0.0.1-x86_10313.i...	available	099034111737	public

Launch Permissions

(Empty panel)


Your Instances

Reservation ID	Owner	Instance ID	AMI	State	Public DNS	Private DNS	Key	Groups	Reason	Idx

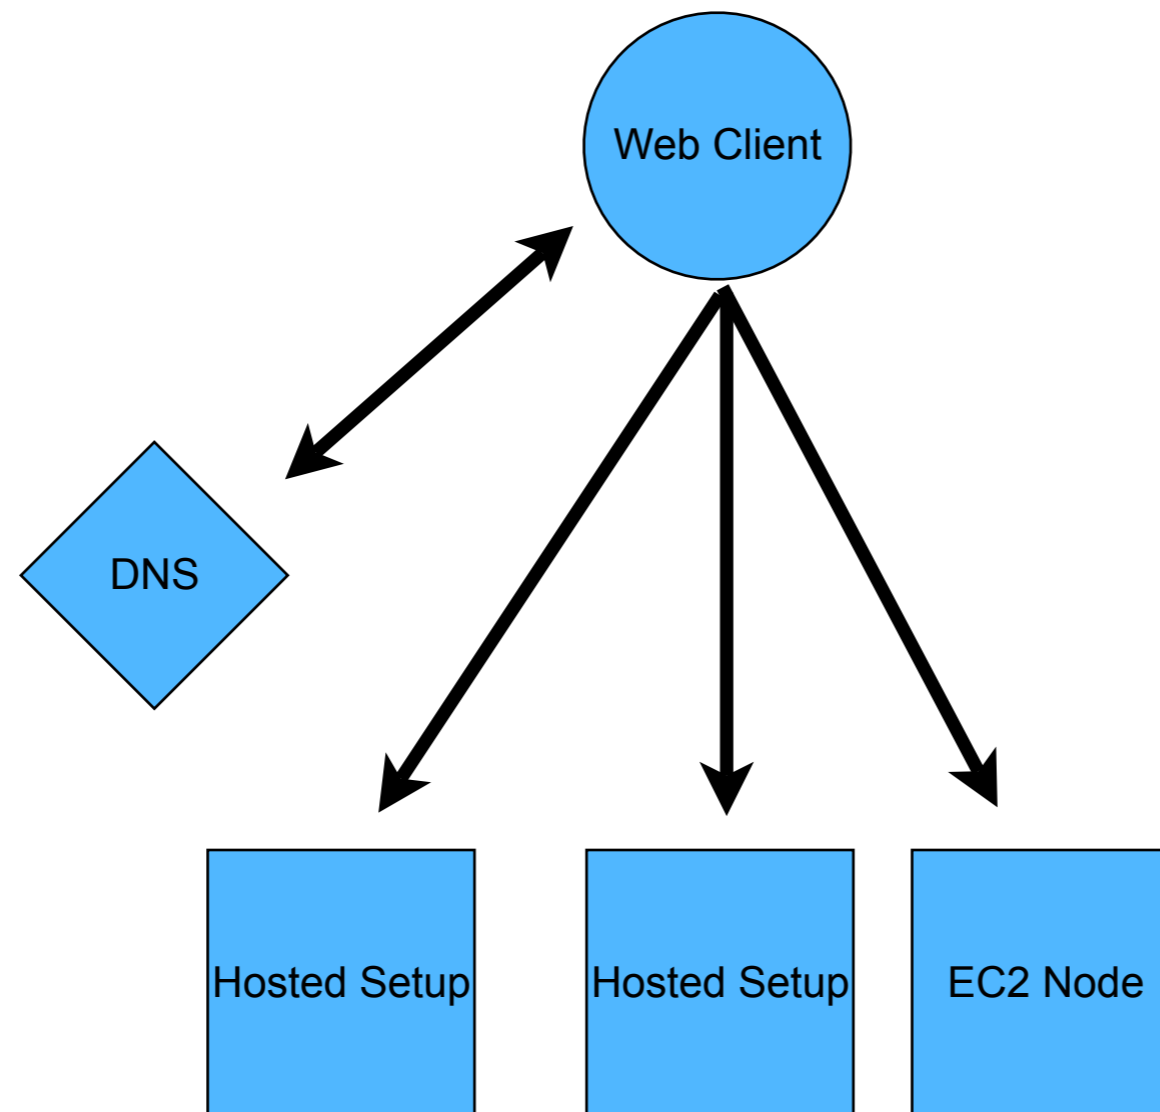
Done

Deployment Patterns


Advanced



A little more common...



Content Serving...



Issues for Databases?

- Not Durable
- Yes, MySQL has multiple storage engines and there is a S3 engine available.

S3 Engine

- Uses MySQL's Loadable Engine Interface (5.1)
- Buckets -> Tables
- Rows -> Items
- Fields -> Encoded (blob by default, XML is coming)

- More Information Can be Found
 - <http://mark.atwood.name/code/mysql-awss3/>
 - (Mark has also written some tools for working with S3 as well <http://hg.fallenpegasus.com/s3-tools/>)

Can I backup my EC2?

- Always store the AMI
- LVM!
- Zmanda
- (Want to be a consumer, try out JungleDisk)

Is there more?

- Amazon Queue Service
- Mechanical Turk



Questions?

(brian@tangent.org | <http://krow.livejournal.com>)