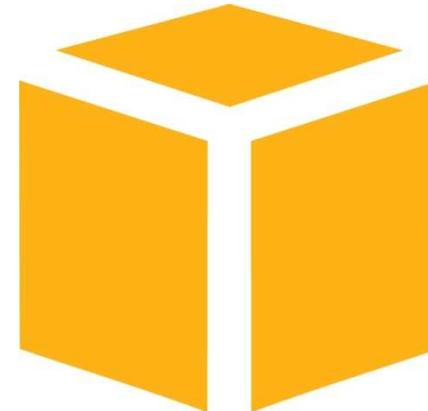


Amazon Web Services: Building Blocks for True Internet Applications

Jeff Barr

Senior Web Services Evangelist

jbarr@amazon.com



Who am I?



- ❏ Software development background
- ❏ Early RSS: Headline Viewer & Syndic8.com
- ❏ Programmable applications and web sites
- ❏ Microsoft Visual Basic and .Net Teams
- ❏ Startup / venture consultant
- ❏ Over 5 Years with Amazon:
 - ❏ Senior Developer
 - ❏ Senior Web Services Evangelist
 - ❏ World traveler – Wiki powered

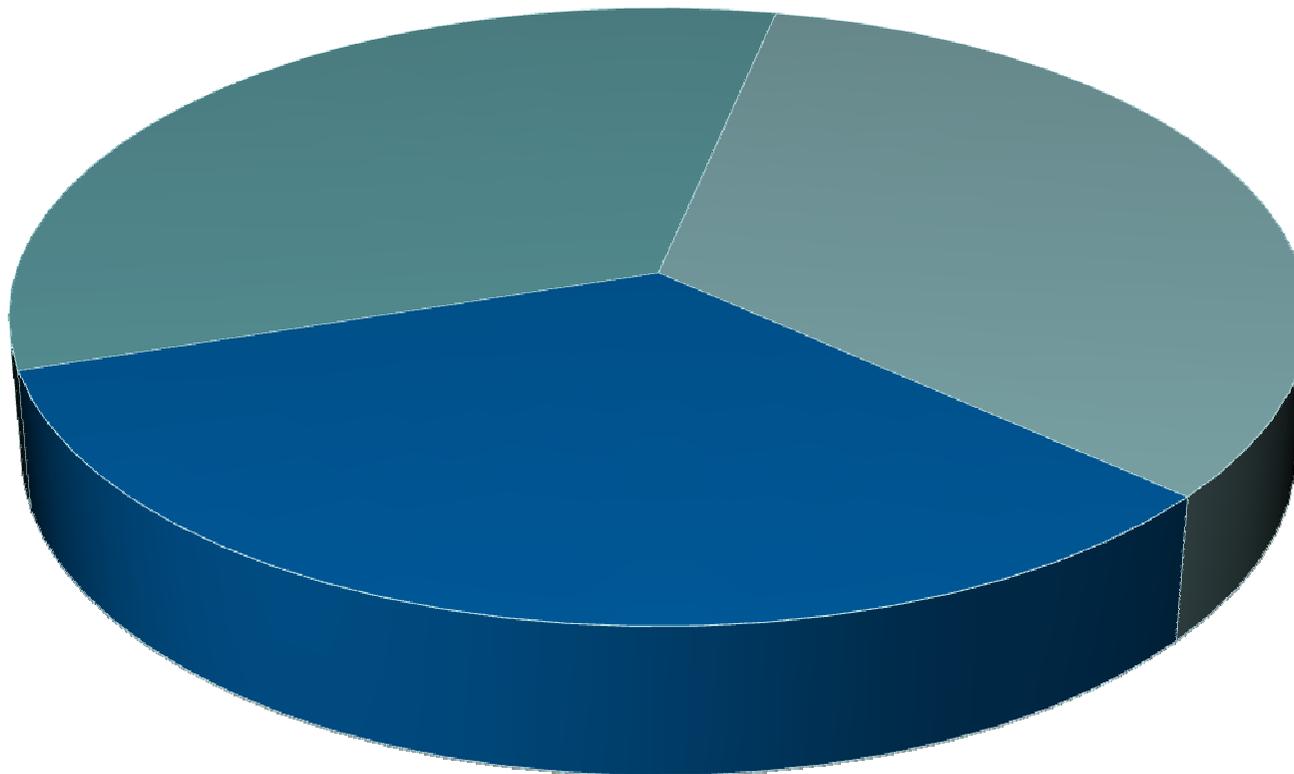


Survey Says...



- ❑ Are you an Amazon retail customer?
- ❑ Have you heard of the Amazon Web Services?
- ❑ Have you used them?
- ❑ Which ones (Associates, EC2, S3, SQS)?
- ❑ Have you seen me speak before?
- ❑ Could you give this talk?

Today's Talk...



- Technical
- Business
- Customer

Amazon Web Services



- ❏ Application building blocks
- ❏ Stable APIs
- ❏ Proven Amazon infrastructure
- ❏ Focus on innovation and creativity
- ❏ Long-term investment

Issues Facing Developers



📦 70% of Web Development Effort is:

- 📦 Data Centers
- 📦 Bandwidth / Power / Cooling
- 📦 Operations
- 📦 Staffing



📦 Scaling is Difficult and Expensive:

- 📦 Large Up-Front Investment
- 📦 Invest Ahead of Demand
- 📦 Load is Unpredictable



Dream or Nightmare?

- ❏ Rapid, unexpected customer demand/growth
- ❏ Unpredictable load
- ❏ Slashdot/Digg/TechCrunch
- ❏ “Success disaster”
- ❏ Seasonal spikes



Reality: Cloud Computing

- 📦 Scale capacity on demand
- 📦 Turn fixed costs into variable costs
- 📦 Clean APIs and conceptual models
- 📦 Always available
- 📦 Cost-effective
- 📦 Reduced time to market
- 📦 Focus on product



Utility Computing Services Menu

- Amazon Simple Queue Service
- Amazon Simple Storage Service
- Amazon Elastic Compute Cloud
- Amazon Flexible Payments Service
- Amazon SimpleDB
- Amazon DevPay



Amazon Simple Storage Service

S3

Amazon Simple Storage Service



- Object-Based Storage
- 1 B – 5 GB / object
- Fast, Reliable, Scalable
- Redundant, Dispersed
- 99.99% Availability Goal
- Private or Public
- Per-object URLs & ACLs
- US & European Locations

**\$.15 per GB
per month
storage**

**\$.01 for 1000 to
10000 requests**

**\$.10 - \$.18 per
GB data transfer**

Amazon S3 Concepts



- ❏ **Objects:**
 - ❏ Opaque data to be stored (1 byte ... 5 Gigabytes)
 - ❏ Authentication and access controls

- ❏ **Buckets:**
 - ❏ Object container – any number of objects
 - ❏ 100 buckets per account

- ❏ **Keys:**
 - ❏ Unique object identifier within bucket
 - ❏ Up to 1024 bytes long
 - ❏ Flat object storage model

- ❏ **Standards-Based Interfaces:**
 - ❏ REST and SOAP
 - ❏ URL-Addressability – every object has a URL
 - ❏ BitTorrent seed



S3 API



Service:

- ListAllMyBuckets

Buckets:

- CreateBucket
- DeleteBucket
- ListBucket
- GetBucketAccessControlPolicy
- SetBucketAccessControlPolicy
- GetBucketLoggingStatus
- SetBucketLoggingStatus

Objects:

- PutObject
- PutObjectInline
- GetObject
- GetObjectExtended
- DeleteObject
- GetObjectAccessControlPolicy
- SetObjectAccessControlPolicy



S3.RB – Establish Connection



```
require 'S3'
```

```
AWS_ACCESS_KEY = '<your key>'
```

```
AWS_SECRET_ACCESS_KEY = '<your key>'
```

```
conn = S3::AWSAuthConnection.new  
(AWS_ACCESS_KEY_ID,  
AWS_SECRET_ACCESS_KEY,  
false)
```

S3.RB – Create Bucket



```
BUCKET_NAME = 'assets.example.com'  
conn.create_bucket(BUCKET_NAME)
```

S3.RB – Upload File



```
datafile = File.open(path)
Key = path.basename;

conn.put(BUCKET_NAME, key, datafile.read,
{"Content-Type" => mime,
 "Content-Length" => File.size(path).to_s,
 "x-amz-acl" => "public-read"})

http://assets.sample.com/...
```

Common Amazon S3 Use Cases

Media Sharing

NRK, Smugmug, Blingee, Ringo



Media/Software Distribution

Microsoft, Liberated Syndication, Linden Lab



Backup (Server and PC)

Mailtrust



Online Storage

MyDataBus, SendAlong



Application Storage

37 Signals, Xerox Global Services



S3 Firefox Organizer - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

chrome://s3fox/content/s3foxWindow.xul

Bloglines |... Jeff Barr's... Twitter Pageflake... Digg / News How to ins... Home | Ub... Official Lin... amazon s... Amazon ... Kednos PL... S3 Firefo...

Manage Accounts AWS DR

C:\temp

File Name	File Size(...)	Modified Time
gestures.bmp	441	05/09/2007 11:05 PM
gestures.GIF	21	05/09/2007 11:04 PM
SQS.bmp	50	05/09/2007 09:50 PM
S3.bmp	50	05/09/2007 09:47 PM
ec2.bmp	50	05/09/2007 09:45 PM
Thumbs.db	60	05/08/2007 02:23 AM
SM-00.bmp	787	05/08/2007 00:02 AM
SM-01.bmp	787	05/08/2007 00:01 AM
Smugmug_1.bmp	3146	05/07/2007 11:47 PM
Smugmug_2.bmp	2606	05/07/2007 11:44 PM
yelx.JPG	10	05/02/2007 06:48 AM
redx.jpg	6	05/02/2007 06:32 AM
raul_y_david.jpg	155	04/23/2007 07:50 AM
2007_Q1_Evangelism_Plan-Jeff.ppt	16636	04/04/2007 07:54 PM
aws_blog.txt	1297	02/12/2007 00:10 AM
estate.png	188	02/05/2007 09:47 PM
Snowbooks	0	05/13/2007 09:56 PM
LL_Meeting	0	05/08/2007 02:23 AM
BarcelonaSmall	0	05/08/2007 02:23 AM

/chessturkj

File Name	File Size(KB)	Upload Time
bb.gif	1	08/08/2006 01:50 AM
bn.gif	2	08/08/2006 01:50 AM
bk.gif	2	08/08/2006 01:50 AM
bp.gif	1	08/08/2006 01:50 AM
bq.gif	2	08/08/2006 01:50 AM
br.gif	1	08/08/2006 01:50 AM
hnc_lockup.gif	1	08/08/2006 01:50 AM
lgoIE.gif	3	08/08/2006 01:50 AM
lwb.gif	1	08/08/2006 01:50 AM
lwk.gif	2	08/08/2006 01:50 AM
lwn.gif	2	08/08/2006 01:50 AM
lwp.gif	1	08/08/2006 01:50 AM
lwq.gif	2	08/08/2006 01:50 AM
lwr.gif	1	08/08/2006 01:50 AM
question.xml	1	08/08/2006 10:05 PM
commons-codec-1.3.jar	47	08/18/2006 01:35 AM
serializer.jar	189	08/28/2006 11:32 PM
xalan.jar	3079	08/28/2006 11:32 PM
xercesImpl.jar	1204	08/28/2006 11:33 PM

➔
➔

File Name	Type	From	To	Progress	Status

Clear

Done

PageRank Alexa

Amazon Elastic Compute Cloud

EC2



This is Not EC2...



http://en.wikipedia.org/wiki/Image:Sinclair_ZX81.jpg

Amazon Elastic Compute Cloud



- Virtual Compute Cloud
- Root-level System Access
- Elastic Capacity
- Management API
- Scale in Minutes
- Multiple Instance Sizes
- Network Security Model

**\$.10-\$.80 per
server hour**

**\$.10 - \$.18 per GB
data transfer**

EC2 Instance Types



1.7 GB RAM
1 EC2 Compute
Unit
160 GB
32-bit platform

\$0.10/Hr



7.5 GB RAM
4 EC2 Compute
Units
850 GB
64-bit platform

\$0.40/Hr



15 GB RAM
8 EC2 Compute
Units
1690 GB storage,
64-bit platform

\$0.80/Hr

Amazon EC2 Concepts

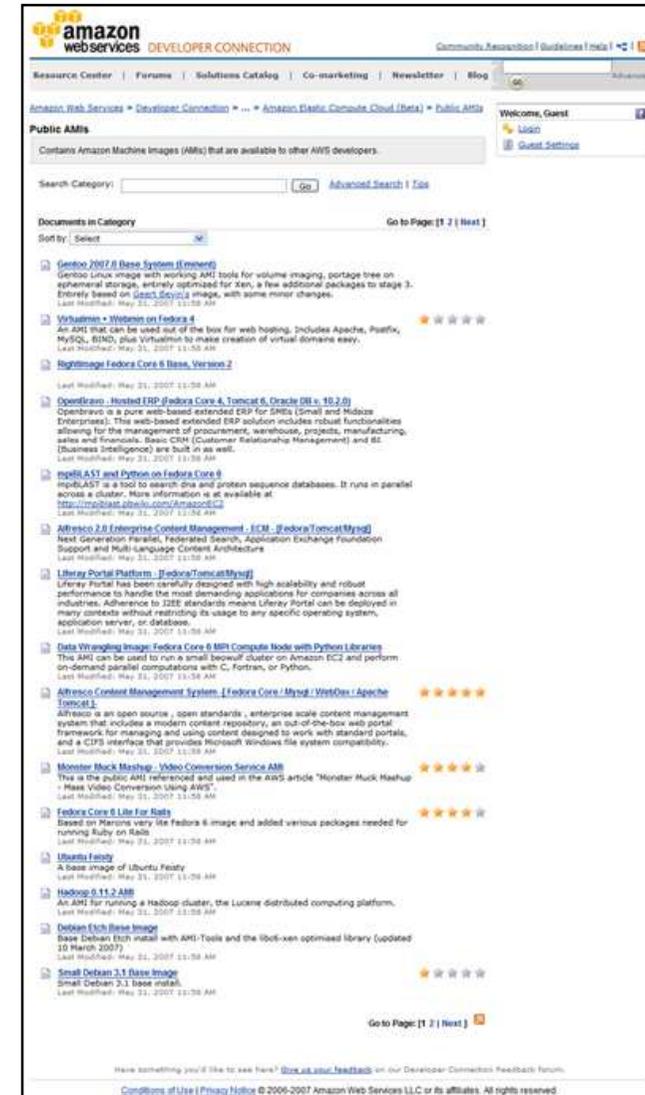


- Amazon Machine Image (AMI):
 - Bootable root disk stored in S3
 - Pre-defined or user-built
 - Catalog of user-built AMIs
 - OS: Fedora, Centos, Gentoo, Debian, Ubuntu, Windows Server
 - App Stack: LAMP, mpiBLAST, Hadoop

- Instance:
 - Running copy of an AMI
 - Launch in less than 2 minutes
 - Start/stop programmatically

- Network Security Model:
 - Explicit access control
 - Security groups

- Inter-service bandwidth is free



EC2 API



Images:

-  RegisterImage
-  DescribeImages
-  DeregisterImage

Instances:

-  RunInstances
-  DescribeInstances
-  TerminateInstances
-  GetConsoleOutput
-  RebootInstances

Keypairs:

-  CreateKeyPair
-  DescribeKeyPairs
-  DeleteKeyPair

Image Attributes:

-  ModifyImageAttribute
-  DescribeImageAttribute
-  ResetImageAttribute

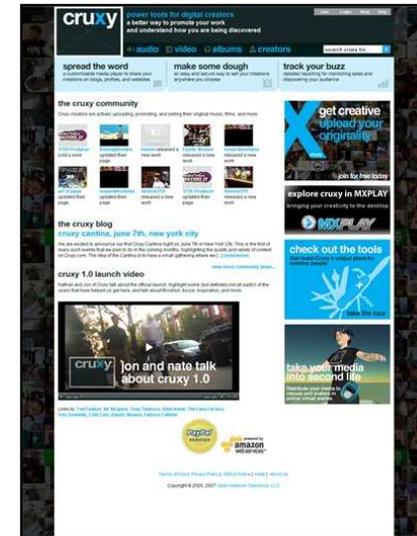
Security Groups:

-  CreateSecurityGroup
-  DescribeSecurityGroups
-  DeleteSecurityGroup
-  AuthorizeSecurityGroupIngress
-  RevokeSecurityGroupIngress

Amazon EC2 At Work



- Startups
 - Cruxy – Media transcoding
 - Podango – Podcast Management
- Science / Research:
 - Hadoop / MapReduce
 - mpiBLAST
- Load-Management and Load Balancing Tools:
 - Rightscale
 - WeoCeo
- Fortune 500 clients:
 - High-Impact, Short-Term Projects
 - Development Host



“Prorated Super Computing Fun” – The New York Times



HOME PAGE MY TIMES TODAY'S PAPER VIDEO MOST POPULAR TIMES TOPICS

The New York Times
Tuesday, February 12, 2008 **Open**

WORLD U.S. N.Y. / REGION BUSINESS TECHNOLOGY SCIENCE HEALTH SPORTS OPIN

OPEN

All the Code That's Fit to `printf()`

[Back to front page »](#)

November 1, 2007, 5:30 pm

Self-service, Prorated Super Computing Fun!

By DEREK GOTTFRID
TAGS: AWS, EC2, HADOOP, MAPREDUCE, S3

As part of [eliminating TimeSelect](#), [The New York Times](#) has decided to make all the public domain articles from 1851-1922 available free of charge. These articles are all in the form of images scanned from the original paper. In fact from 1851-1980, all 11 million articles are available as images in PDF format. To generate a PDF version of the article takes quite a bit of work — each article is actually composed of numerous smaller TIFF images that need to be scaled and glued together in a coherent fashion.



Magic EC2 / S3 Button

Previously we had generated all the PDFs dynamically. This approach had worked reasonably well, but with the strong possibility of a significant traffic increase we started to rethink things. Clearly, pre-generating all the articles and statically serving them would be a great option. Pretty quickly I thought about how we could do this (and have some fun along the way,

- 1851-1922 Articles
- TIFF -> PDF
- Input: 11 Million Articles (4TB of data)
- What did he do?
 - 100 EC2 Instances for 24 hours
 - All data on S3
 - Output: 1.5 TB of Data
 - Hadoop, iText, JetS3t

EC2 Firefox UI - Mozilla Firefox

File Edit View History del.icio.us Bookmarks Tools Help

chrome://ec2ui/content/ec2ui_main_window.xul world map

Mail Daily Weekly Blogs Forums wikipedia WRT 54 GG Router (...) Press It - Jeff Barr's ... WEBoggle Technorati + - AJAX Translator post to del.icio.us WTP Digg Generator >>

Credentials Tools Jeff Account IDs About

AMIs and Instances KeyPairs Security Groups

Available AMIs

AMI ID	Manifest	State	Owner	Visibility
ami-9c22c7f5	wowzapro132/image.manifest.xml	available	632925041470	public
ami-af32d7c6	wowzamediaserverproj1.0.3-x86_64.manifest.xml	available	632925041470	public
ami-ac32d7c5	wowzamediaserverproj1.0.3-i386.manifest.xml	available	632925041470	
ami-d02ecbb9	wowzamediaserverproj1.0.2-x86_64.manifest.xml	available	632925041470	public
ami-d12ecbb8	wowzamediaserverproj1.0.2-i386.manifest.xml	available	632925041470	public
ami-3406e35d	workspace.globus.org/osg08/osgheadnode.manifest.xml	available	971725951873	public
ami-2eee0b47	workspace.globus.org/osg07/osgheadnode.manifest.xml	available	971725951873	public
ami-3def0a54	workspace.globus.org/osg06/osgworker.manifest.xml	available	971725951873	public
ami-7eee0b17	workspace.globus.org/osg06/osgheadnode.manifest.xml	available	971725951873	public
ami-3bef0a52	workspace.globus.org/osg05/osgworker.manifest.xml	available	971725951873	public
ami-3aef0a53	workspace.globus.org/osg05/osgheadnode.manifest.xml	available	971725951873	public
ami-66ef0a0f	workspace.globus.org/osg04/osgworker.manifest.xml	available	971725951873	public
ami-aeec09c7	workspace.globus.org/osg03/osgworker.manifest.xml	available	971725951873	public

Launch Permissions

Empty panel for launch permissions.

Your Instances

Reservation ID	Owner	Instance ID	AMI	State	Public DNS	Private DNS	Key	Groups	Reason	Idx	Type	Local Launch Time
Empty table for instances.												

Done

PageRank Alexa Compete S3 Max Open Notebook

Amazon Simple Queue Service

SQS

Amazon Simple Queue Service



- Scalable Queuing
- Elastic Capacity
- Reliable, Simple, Secure

Inter-process messaging, data buffering, scalable architecture component

**\$.01 per
10000
messages**

**\$.10 - \$.18 per
GB data transfer**

SQS API



Queues:

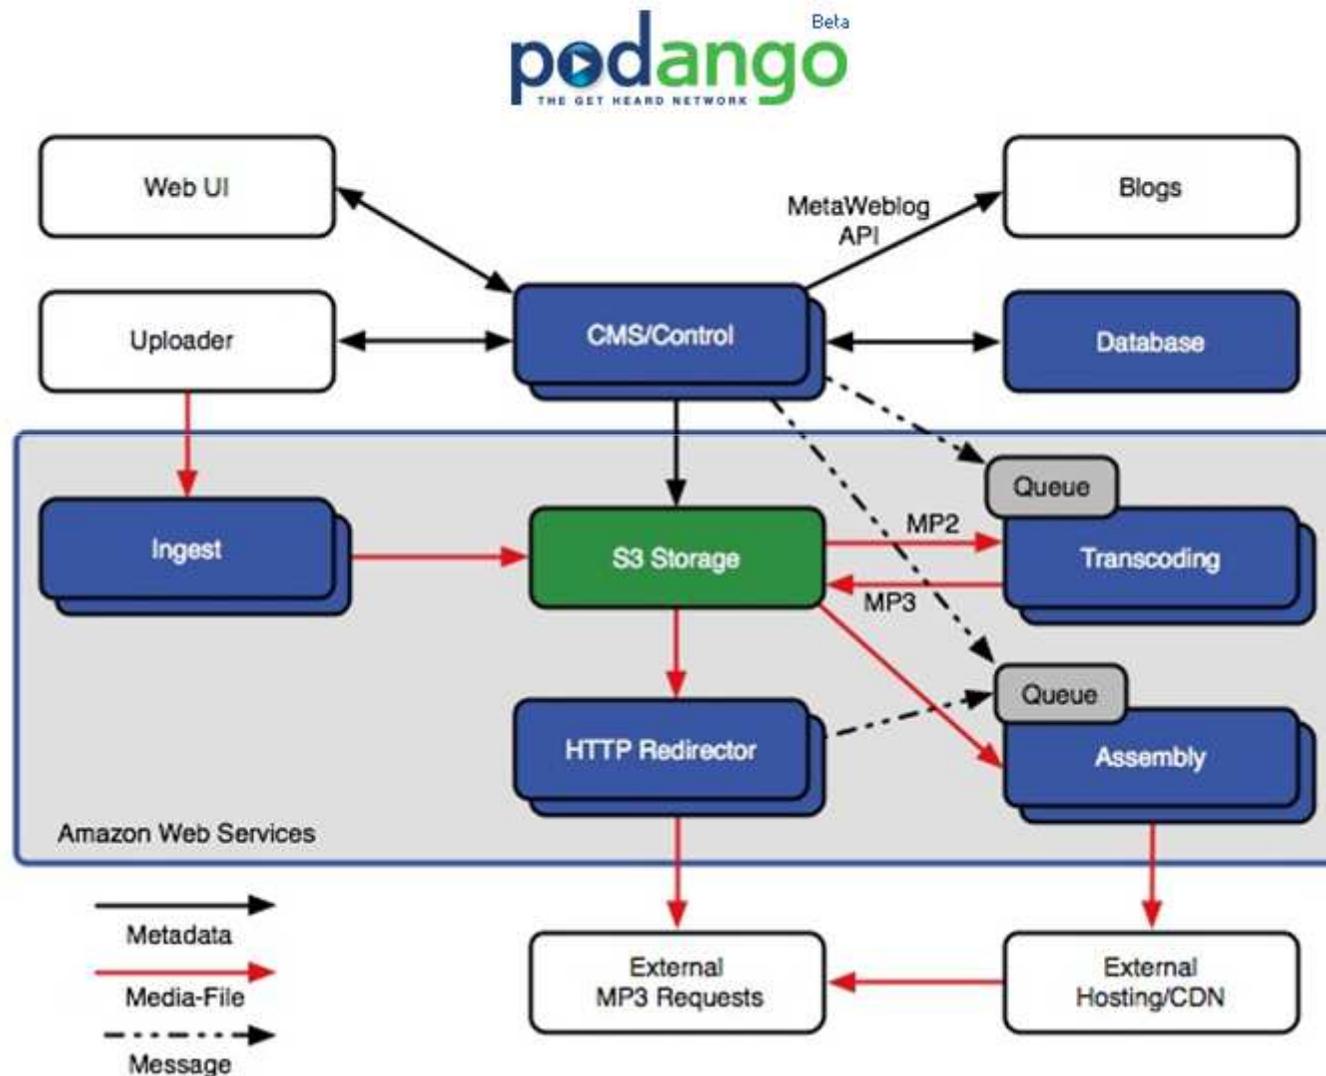
-  ListQueues
-  DeleteQueue

Messages:

-  SendMessage
-  ReceiveMessage
-  DeleteMessage



Amazon SQS At Work



Amazon SimpleDB

SDB

Amazon SimpleDB Service



- Distributed Data Store
- Structured Storage
- Fully Indexed
- Redundant
- Scalable
- Elastic Capacity
- Query Language

\$1.50 per GB
per month

\$.10 - \$.18 per GB
data transfer

\$.14 per CPU hour
(query processing)

SimpleDB Concepts



Domain:

- Collection of similar items
- Query language
- Any number of items per domain (10 GB beta limit)
- 100 domains per account

Item:

- Collection of key-value pairs (attributes)
- Multiple values per attribute
- Up to 256 attributes per item
- Up to 1024 bytes per value

Billing:

- Data storage
- CPU utilization
- Data storage

itemID	description	color	material
123	Sweater	Blue, Red	
456	Dress shirt	White, Blue	
789	shoes	Black	Leather

SimpleDB API



Domains:

-  CreateDomain
-  ListDomains
-  DeleteDomain

Items:

-  PutAttributes
-  GetAttributes
-  Query

Query Language (samples):

-  ['Title' = 'The Right Stuff']
-  ['Number of Pages' < '00310']
-  ['Rating' = '***' or 'Rating' = '*****']
-  ['Year' > '1950' and 'Year' < '1960' or 'Year' starts-with '193' or 'Year' = '2007']
-  'Keyword' = 'Frank Miller' union ['Rating' starts-with '*****']

Amazon SimpleDB Scenarios

- 📦 Data-Sharing between Social Networks
- 📦 Meta-data
 - 📦 Store files in S3 and metadata in SimpleDB
- 📦 Lookup
 - 📦 Contacts, Directories and Address maps
- 📦 Logging
 - 📦 ClickStreams, Status Data
- 📦 Catalogs
 - 📦 Product catalogs
- 📦 Other
 - 📦 Emails/IM/Feeds



Other Amazon Infrastructure Services



Flexible Payments Service:

-  Move money between any two people or systems.
-  Credit cards, bank accounts, Amazon Payments.
-  One-time, multiple, or recurring transactions.
-  Payment processing language (Gatekeeper).
-  Aggregated transactions (micropayments).

DevPay:

-  Wrap custom business models around S3 and EC2.
-  Set custom prices for each charging unit.
-  Charge for your applications.

Web Services Billing



amazon.com | Jeff's Store | Make Money | See All 34 Product Categories | Your Account | Cart | Wish List | Help

Program Overview | Marketplace | Associates | Advantage | **Web Services** | Paid Placements | On-Demand Publishing

Search: Help [input type="text"] [GO] [GO] Web Search [input type="text"] [GO]

Welcome, Jeff @ AWS. [Sign Out](#) [Your Web Services Account](#)

Account Activity
Account Number: 8892-7910-8296
[Go to Last Month's Statement](#)

Billing Statement: June 1, 2007
Billing Cycle for this Report: May 1 - May 31, 2007

Usage Charges	Rate	Usage	Totals
Alexa Web Information Service View/Edit Service	\$0.00015 per request (\$0.15 for 1,000 requests)	0 Requests	0.00
		Usage Report	0.00
Amazon Simple Storage Service View/Edit Service	\$0.20 per GB of data transferred (details)	3,415 GB	0.69
	\$0.15 per GB-Month of storage used (details)	9,222 GB-Mo	1.39
		Usage Report	2.08
Developer Connection View/Edit Service	\$0.10 per 1000 messages	0 Messages	0.00
		Usage Report	0.00
Subtotal			\$ 2.08
Taxes			\$ 0.00
Charges for June 1, 2007			\$ 2.08

(Charged to your credit card on June 1, 2007) Visit the [Amazon Web Services FAQs](#) to learn more about web services pricing models and billing.
 All web services are sold by Amazon Web Services LLC

Summary of Last Month's Statement: April 1 - April 30, 2007

Total Charges for April 1 - April 30, 2007 **\$ 1.71**
[Go to Last Month's Statement](#)

View a Previous Statement:
 Current Billing Period: [dropdown] [Go] [Printer Friendly Version](#)

[Conditions of Use](#) | [Privacy Notice](#) © 1996-2007, Amazon.com Inc. or its affiliates

Amazon Simple Storage Service			
View/Edit Service	\$0.20 per GB of data transferred (details)	3,415 GB	0.69
	\$0.15 per GB-Month of storage used (details)	9,222 GB-Mo	1.39
		Usage Report	2.08

Q & A



Thank You!



Jeff Barr

Senior Web Services Evangelist

jbarr@amazon.com

Pointers

Portal

<http://aws.amazon.com>

Blog

<http://aws.typepad.com>

EC2

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

S3

<http://aws.amazon.com/s3>

Resource Center

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

Forums

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