

Cloud Computing and Amazon Web Services

CJUG March, 2009
Tom Malaher

Agenda

- What is Cloud Computing?
- Amazon Web Services (AWS)
- Other Offerings
- Composing AWS Services
- Use Cases
- Ecosystem
- Reality Check
- Pros&Cons
- Cost
- AWS Demo
- Resources



What is Cloud Computing?

Many Definitions: Each Vendor defines it to match their service offering

From the Infrastructure Executive Council

- **Elastic**: The ability of the service delivery infrastructure to expand and contract automatically based on the capacity needs.
- **Multi-tenant**: The services are architected such that several customers share the underlying infrastructure resources without compromising the privacy and security of customer's data.
- **On-demand**: All Cloud services are available over the internet and can be consumed as and when there is a need for resources.
- **Usage-based metering**: Billing is either based on the usage of computing resources or on a subscription based model (without long-term contracts).
- **Self-service**: All services are simple to use and can be directly consumed by the user from a user interface or an API.

Variations on a Theme

Infrastructure as a Service (IaaS)
Storage, Computing

Platform as a Service (PaaS)
Microsoft Azure
Google App Engine

Software as a Service (SaaS)
SalesForce.com
Google Docs
Microsoft Office Live (note: still needs local install!)

Grid Computing



Amazon Web Services (AWS)

- Commerce-Related
 - Associates Web Service
 - Flexible Payment Service (FPS)
 - Fulfillment Web Service (FWS)
- Infrastructure-Related
 - Elastic Compute Cloud (EC2)
 - Runs Amazon Machine Images (AMIs) (Linux, Windows)
 - Elastic IP, Elastic Block Store (EBS)
 - SimpleDB (SDB)
 - Simple Storage Service (S3)
 - Simple Queue Service (SQS)
 - CloudFront
 - DevPay
- Miscellaneous
 - Alexa Web Search (Family of services)
 - Mechanical Turk
 - Public Data Sets

Other Offerings

Sun

MySQL, OpenSolaris AMIs on EC2

Google App Engine

Host specially written Python Apps in Google's Cloud

IBM

Computing on Demand (CoD) Centers

Amazon Machine Images

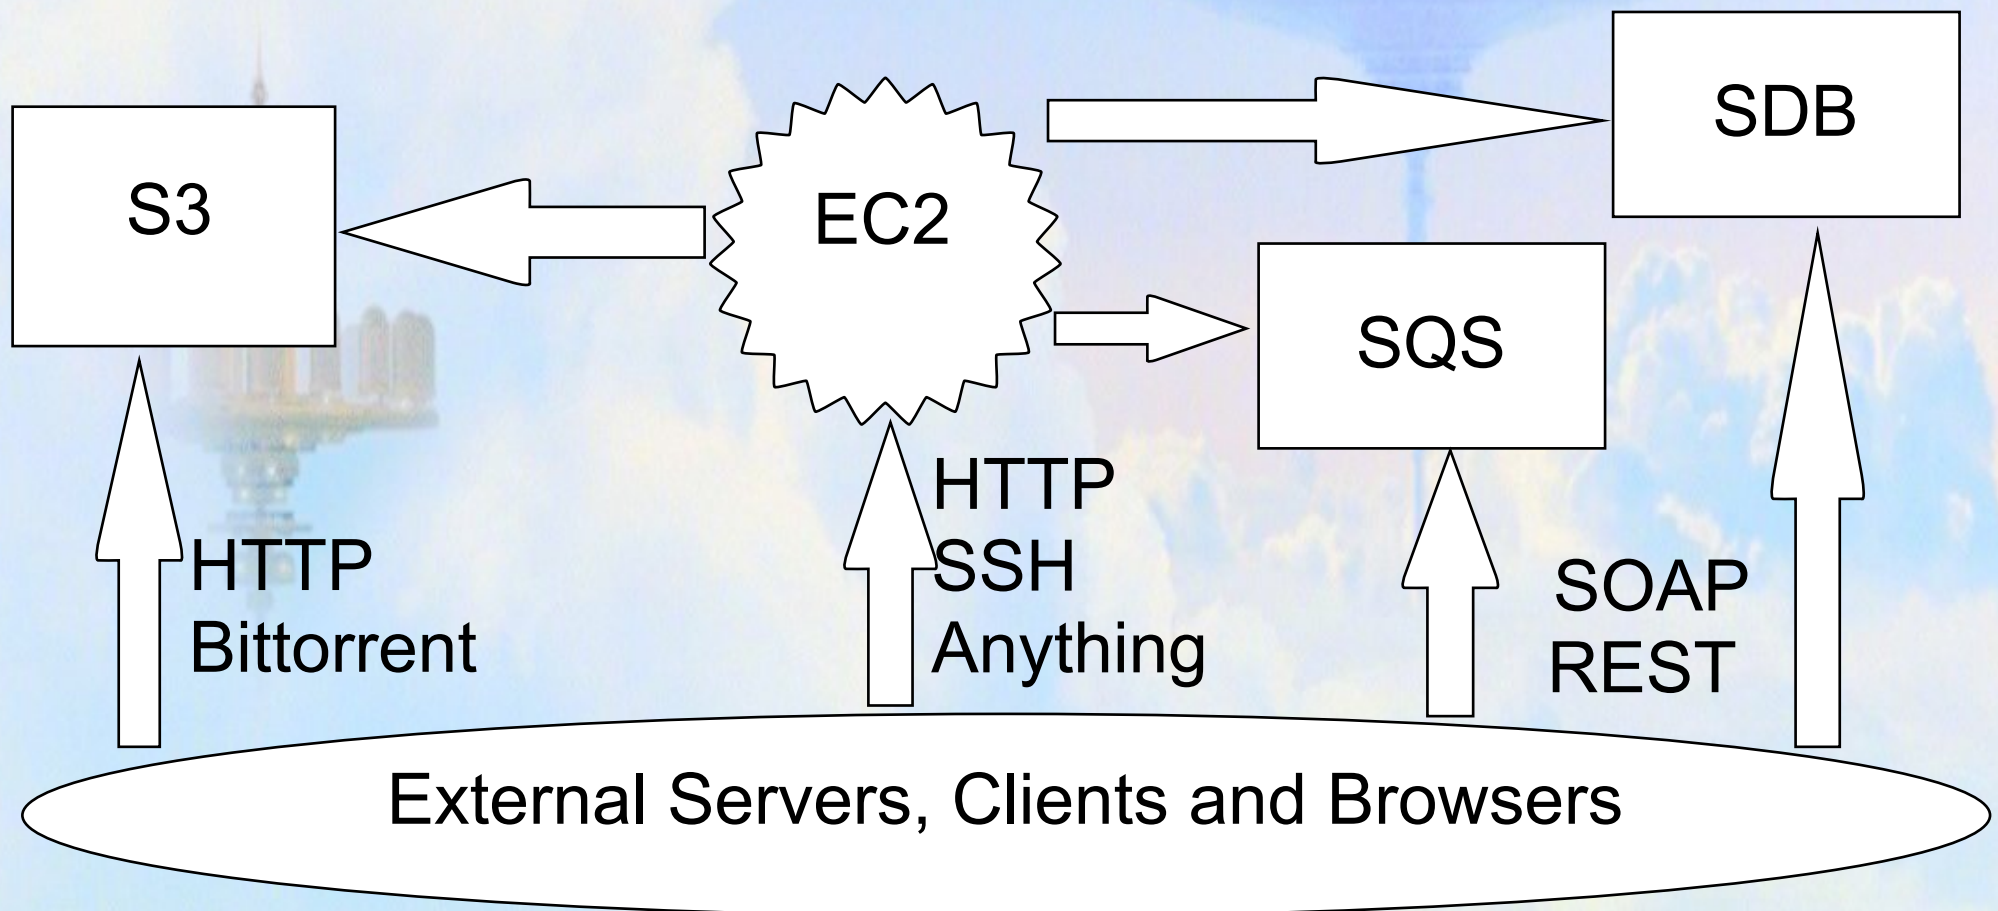
Microsoft Azure

Host specially written .NET Apps in Microsoft's Cloud

.NET, Live and SQL Services available

Composing AWS Services

Communication between AWS Services "within the cloud" is cheaper (no bandwidth charges) than crossing the cloud boundary (some services still charge per transaction).



Composing AWS Services (cont.)

Roles of each service

- EC2: Web Server, App Server, Relational DB Server
 - Can be clustered
 - Expand/contract on demand
- SDB: Non-Relational DB
 - Quick lookup
 - No Schema, Strings only!
 - No Transactions
 - Limited size: point to S3 for data
- SQS: Asynchronous de-coupled communicating processes
 - Visibility Timeout instead of transactions/acknowledgement
 - Limited message size: point to SDB or S3 for data
- S3: Large data sets(Media, Block Data)

Use Cases

eCommerce

- EC2:
 - Web (cluster)
 - App (cluster)
 - DB if needed (data warehouse, reporting, order history)
- SDB: Carts, Wish Lists, Product Catalog
- SQS: Completed Orders sent to Fulfillment
- S3: Media, Storage of:
 - Product Images/Videos
 - PDF Specifications
 - Software bundles for electronic delivery
- CloudFront
 - Delivery of above via HTTP or BitTorrent

Use Cases (cont.)

Bulk Data Processing/Map-Reduce (e.g. Image Conversion, Hadoop, Genome Analysis)

- EC2: Controller, Workers
- SQS: Controller messages, Worker responses
- SDB: Work unit metadata
- S3: Worker input data, Worker results, Intermediate Reduce Results, Final Results

Software as a Service:

- EC2: Web Application, Load Balancing
- S3: Application Data
- SDB: Metadata (User Database, ACLs, Indices)

It's an Ecosystem

Other vendors (e.g. Sun, IBM) are offering AMIs and special licensing deals for use on EC2.

There are software vendors who are writing free or commercial applications entirely on AWS.

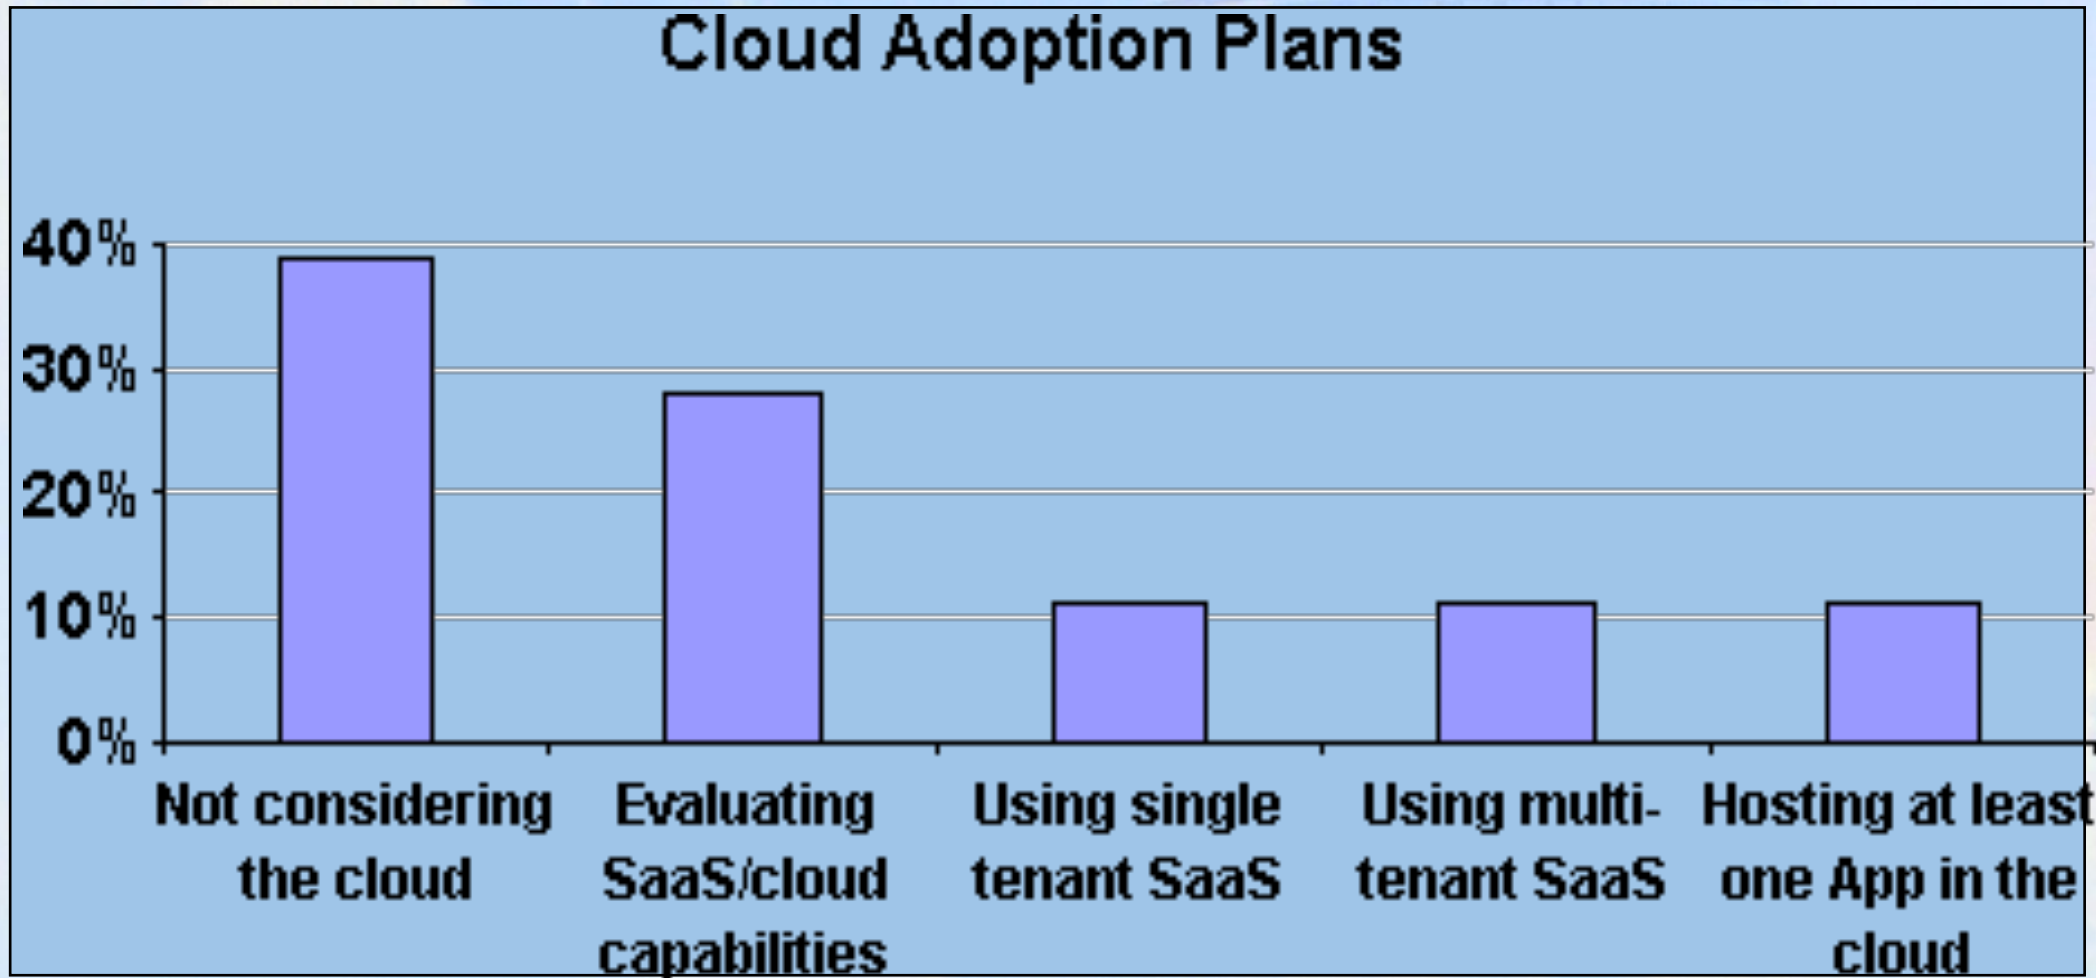
Amazon hosts a Solutions Catalog

- Backup
- Asset Management
- EC2 Management/Monitoring
- CRM
- etc.

Reality Check

Is anybody doing it ?

Source: Infrastructure Executive Council, Fall 2008



Pros

- Expand and Contract as needed (TSN March Madness, PumpkinMasters, Facebook App)
- Pay for only what you use (CPU Time, Bandwidth, Transactions)
- Leave the infrastructure management to someone else
- Amazon offers an SLA with Service Credits (for EC2, so far)

Cons

Applications not suitable for the Cloud (yet)

- Very tight SLA
- Real time application response times
- Applications requiring high bandwidth (to the end user)
- Sensitivity of the application to latency
- Security Concerns
- High Performance (CPU, RAM) Requirements

Cost!

Cost

Difficult to compare, due to threshold-based vs. pay-as-you go pricing, but:

1and1.com: Virtual Server: \$59/month
1GB RAM, 50 GB HD, 2000 GB Transfer

AWS EC2 Image: \$72/month

1.7GB RAM, 160 GB HD

If you actually transferred 2000GB in one month: \$340/month

Persistent 50 GB Storage (EBS): \$67.50/month

Worst case Total: \$479.50

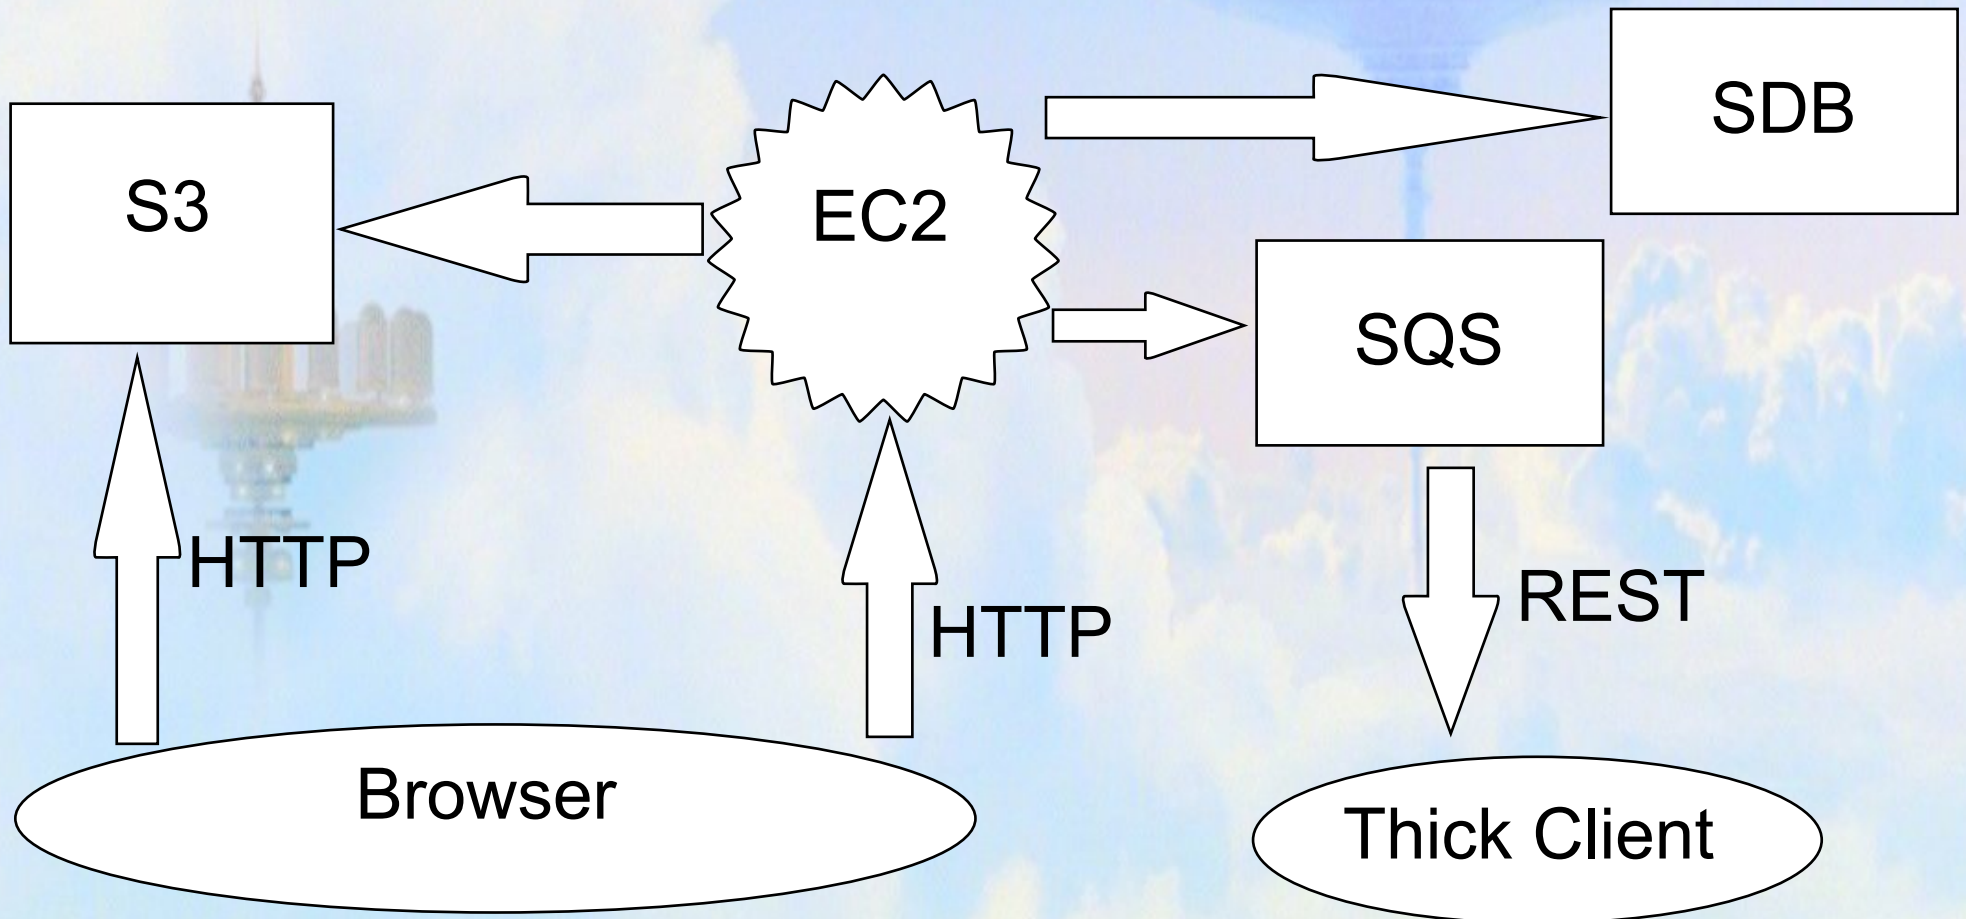
<http://calculator.s3.amazonaws.com/calc5.html>

But if you only used it for a day? \$16

AWS Demo

Task Management System:

Manager sends tasks to workers who are notified of new tasks via a desktop 'alerter' application and a web-based task list.



Resources

Sun

<http://www.sun.com/solutions/cloudcomputing/>

http://blogs.sun.com/WebScale/entry/mysql_in_the_cloud

IBM

<http://www.ibm.com/cloud/>

HP

<http://www.hpl.hp.com/research/cloud.html>

Dell

<http://www.dell.com/cloudcomputing>

Everybody

<http://gigaom.com/2008/05/28/hp-weds-cloud-and-high-performance-computing/>

Yahoo

http://research.yahoo.com/Cloud_Computing

Resources (cont.)

Google

<http://docs.google.com>

<http://code.google.com/appengine/>

Microsoft

<http://www.officelive.com>

<http://www.microsoft.com/azure/>

Amazon

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

Solutions Catalog:

<http://developer.amazonwebservices.com/connect/kbcategory.jspa?categoryID=60>

New York Times Image conversion projects

<http://open.blogs.nytimes.com/tag/aws/>