Maximizing Utility of the Cloud

2009 Bio-IT World Europe

Chris Dagdigian, chris@bioteam.net

BioTeam Inc.

BIOTEAM
Enabling Science

Fair Warning

- **Giving me 60 minutes to talk is dangerous**
- **I'm somewhat infamous**
  - I speak very fast
  - Typically have an insane number of slides
- **Latest slides will be here:** http://blog.bioteam.net

BIOTEAM
Enabling Science

BioTeam Inc.

- **Independent Consulting Shop:**
  Vendor/technology agnostic

  Distributed entity - no physical office
- **Staffed by:**
  - Scientists forced to learn High Performance IT to conduct research
  - Many years of industry & academic experience
- **Our specialty:**
  Bridging the gap between Science & IT

BIOTEAM
Enabling Science

High Level Topics For Today

- **What "cloud" means to me**
  Getting our vocabulary straight

- **Current State Report**
  - Good, bad & ugly
  - Mapping informatics onto the cloud

- **An attempt at some advice**
  - Hard lessons learned
  - Some real world examples

BIOTEAM
Enabling Science

Topics - More Detail

1. Terminology
2. Blunt words:
   Cloud Computing
3. Blunt words:
   Private Clouds
4. Why I drank the Kool-Aid
5. Amazon AWS Overview
6. Cloud Sobriety
7. Cloud Security
8. State of Amazon AWS
9. AWS: Good, Bad & Ugly
10. Examples
11. Recommendations

BIOTEAM
Enabling Science

Setting The Stage

- **Burned by "OMG!! GRID Computing" Hype**
  In 2009 will try hard never to use the word "*cloud*" in any serious technical conversation. *Vocabulary matters.*

- **Understand My Bias:**
  - Speaking of "utility computing" as it resonates with *infrastructure* people
  - My building blocks are servers or groups of systems, not software stacks, developer APIs or commercial products
  - Goal: Replicate, duplicate, improve or relocate complex systems

BIOTEAM
Enabling Science

Lets Be Honest

- **Not rocket science**
- **Fast becoming accepted and mainstream**
- **Easy to understand the pros & cons**

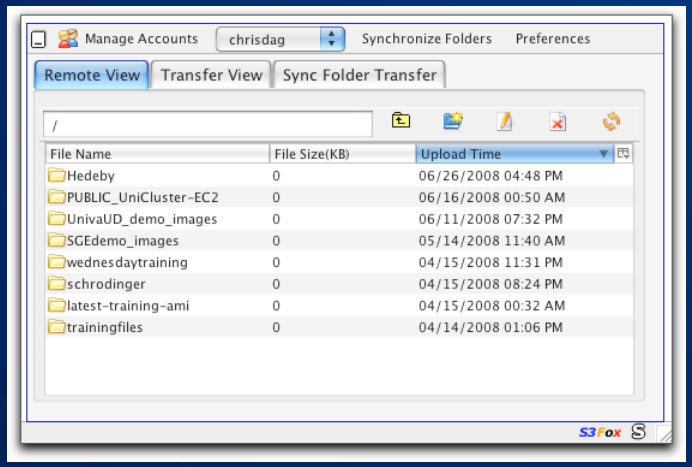

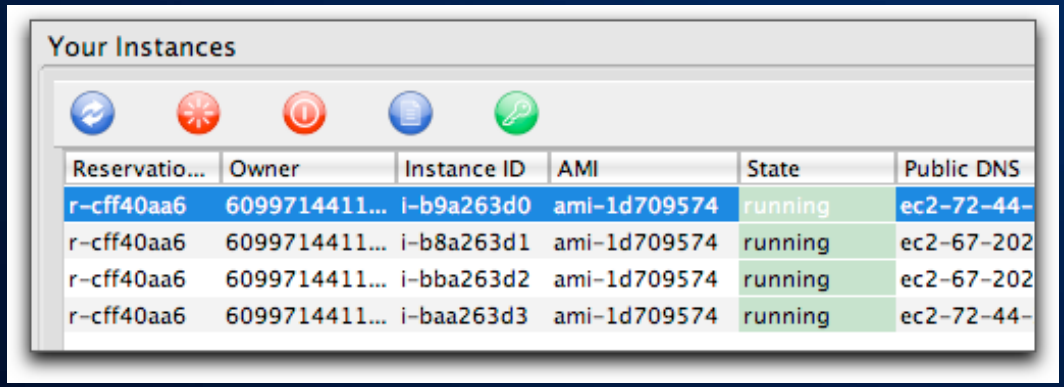

BIOTEAM
Enabling Science

While I'm Being Honest …

- Amazon Web Services *is* the cloud
  - Simple, practical, understandable and usable today by just about anyone
  - Rollout of features and capabilities continues to be impressive
- Competitors are years behind
  - … and tend to believe too much of their own marketing materials

BIOTEAM
Enabling Science

While I'm Being Honest …

- "Private Clouds" = absolute rubbish
  - … in 2009 at least
  - 98% hype & marketing, 2% usefulness
    (just like the 90's era WAN/Grid Computing days)
- There are 2 types of private clouds efforts:
  - Interesting academic papers & pilot projects
  - Heavily contrived vendor demos
  - … none of which have been extensively tested in demanding production computing environments

BIOTEAM
Enabling Science

How To Build A "Private Cloud" in 2009

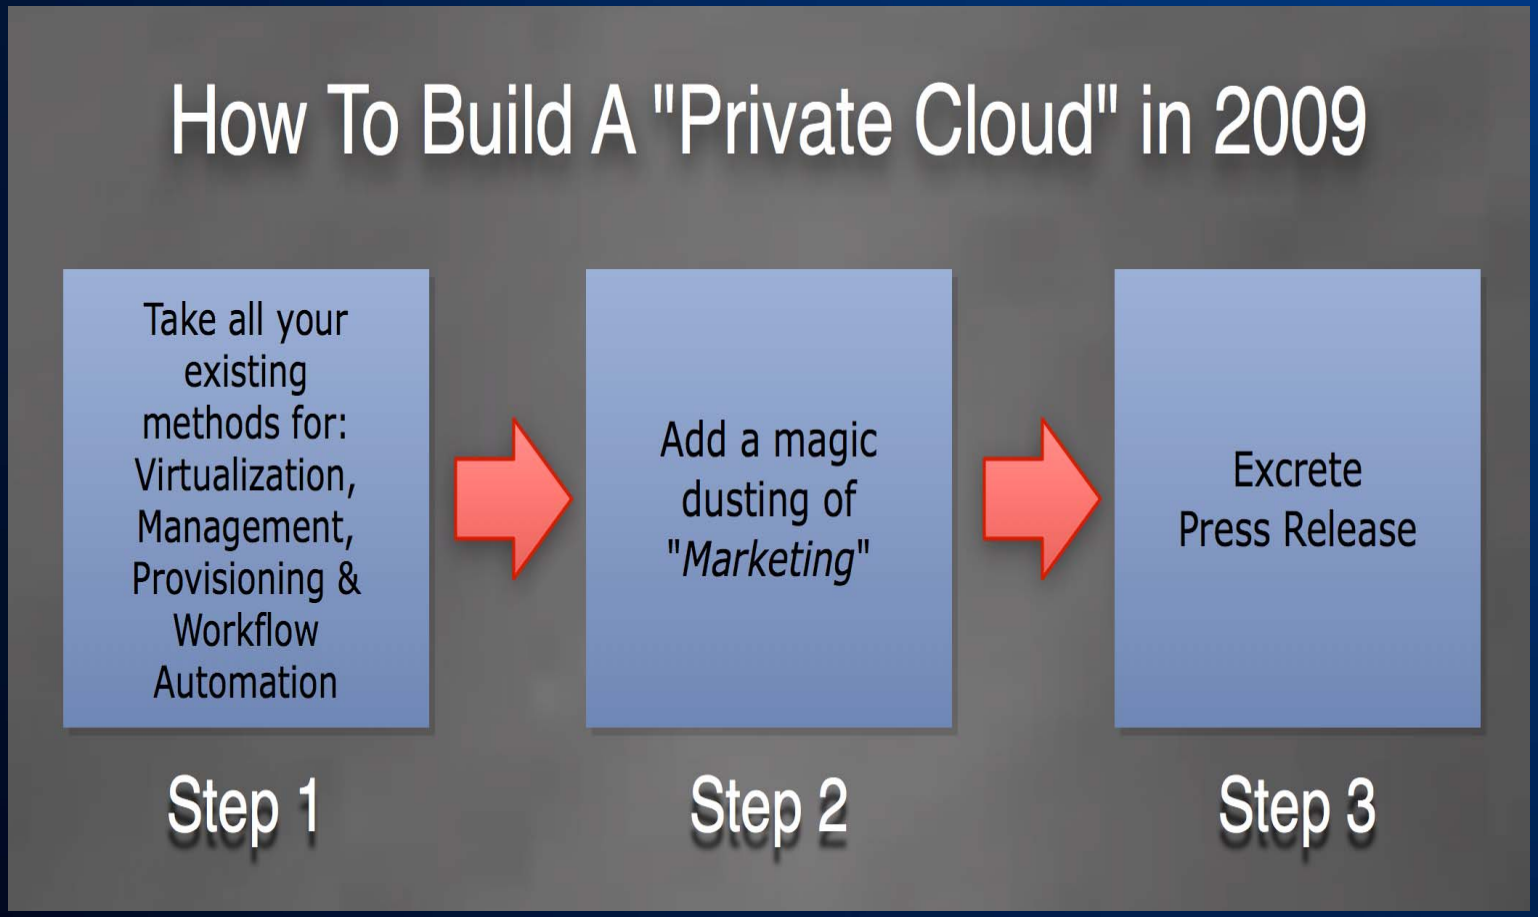

BIOTEAM
Enabling Science

Just one inconvenient truth …

- Clouds are all about motion & agility, but …
- Live migration of a running VM can usually only happen within the same subnet
  - How many of you have a flat layer 2 network spanning everything in your machine room?
    - Does that single subnet extend through all your datacenters?
- Thus We Have Our Inconvenient Fact:
  - A true “private cloud” requires extensive and possibly radical reengineering of network & hardware
  - Difficult to envision this happening in anything but a brand new environment

BIOTEAM
Enabling Science

Utility/Cloud Computing: Getting Back On Topic

Why I drank the Kool-Aid

BIOTEAM
Enabling Science

Tipping Point: Hype to Reality

- 2007: Individual staff experimentation all year
  - Including MPI applications (mpiblast)
- Q1 2008:
  - Realized that every single BioTeam consultant had *independently* used AWS to solve a customer facing problem
  - No mandate or central planning, it just happened organically

BIOTEAM
Enabling Science

BioTeam AWS Use Today

- **Running Our Business**
  Development, Prototyping & CDN
  - Effective resource for tech-centric firms
- **Grid Training Practice**
  - Self-organizing Grid Engine clusters in EC2
  - Students get root on their own cluster
- **Proof Of Concept Projects**
  - UnivaUD - UniCluster on EC2
  - Sun - SDM ‘spare pool’ servers from EC2
- **Directed Efforts on AWS**
  For ISV and Pharma clients

BIOTEAM
Enabling Science

Amazon AWS Overview

http://aws.amazon.com/products/

BIOTEAM
Enabling Science

Amazon Web Services

- A collection of agile infrastructure services available to on-demand
- New products and added features added almost monthly
- Recent enhancements:
  - Two-factor Authentication & Rotating Credentials
  - Virtual Private Cloud (“VPC”) Product
  - EC2 auto-scaling & load-balancing
  - http://aws.amazon.com/about-aws/whats-new/

BIOTEAM
Enabling Science

AWS Products/Services

- EC2 - Elastic Compute Cloud
  - *Scalable on-demand virtual servers*
- SimpleDB - Simple Database Service
  - *Simple queries on structured data*
- S3 - Simple Storage Service
  - *Bucket/object based storage*
- EBS - Elastic Block Service
  - *Persistent block storage (looks like a disk)*

BIOTEAM
Enabling Science

AWS Products/Services, cont.

- SQS - Simple Queue System
  - *Message passing service storage*
- Elastic MapReduce
  - *Hadoop on AWS*
- VPS - Virtual Private Cloud
  - *Connect your infrastructure to AWS via VPN tunnel*
  - *(more important than it sounds …)*

BIOTEAM
Enabling Science

Elastic Compute Cloud ("EC2")

- A set of APIs you can invoke to manipulate remote VM instances
- Easy to launch existing images
- Easy to build your own custom server images
- Xen instances on-demand
  - Starting at .10/hour for 32bit system
  - 64bit systems start at $.40/hour
  - Fire up as many as you need, whenever you need them
    - Many interfaces/control points
      - Mozilla plugins, CLI, Java, Perl, etc.

BIOTEAM
Enabling Science

Elastic Compute Cloud

- Why it works
  - Smart pricing
    - Server instance pricing is reasonable
    - Traffic to/from S3 storage cloud is free
    - Experimenting is dirt cheap
      - 1 week of messing around == invoice for $9 USD
      - Weeklong SGE training on big machines == $79 USD
  - Easy to use

BIOTEAM
Enabling Science

Elastic Compute Cloud

- Why it works, continued
  - Rapid rate of enhancements & new features
    - Availability zones
    - Reserved instances
    - Live credential rotation
  - Clever people can make money
    - Amazon allows reselling AMI instance images
      - I can build a specialized workflow engine and charge a small fee on top of the Amazon costs
      - All financial transactions handled by Amazon
  - Limitations are pretty obvious
    - Easy to know what workflows are or are-not EC2 friendly

BIOTEAM
Enabling Science

Amazon EC2 “Aha! Moment”

- Consider a generic 100 CPU hour research problem:
  - EC2: 10 large servers @ .40/hr for 10 hours
    - Work done in 10 HOURS at cost of $40 USD
  - EC2: 100 large servers @ .40/hr for 1 hour
    - Work done in 1 HOUR at a cost of $40 USD
- Can you do THAT in your datacenter today?

BIOTEAM
Enabling Science

Amazon S3

- Add and remove stuff into “buckets”
  - 1 byte to 5GB per object
  - Required for storage greater than 1 terabyte
- Popular with web 2.0 outfits
- Standard REST and SOAP interfaces
- BitTorrent interface as well
- Required component of EC2 usage
  - All EC2 AMI (server images) are stored in S3
- Cheap to move data in/out
- Reasonable monthly fee for persistent storage
- Free to move data within Amazon services
- Lots of interfaces

BIOTEAM
Enabling Science

Amazon S3, cont.

- Similar rapid rate of enhancements as EC2
- Hooks into Amazon CDN product ('CloudFront')
- Interesting access/download APIs
  - Including "downloader pays"
- Of significant interest to our crowd
  - Physical ingest/outgest service
  - Send your USB 2.0 or SATA device to Amazon for rapid loading of large datasets

BIOTEAM
Enabling Science

Elastic Block Store (“EBS”)

- Block storage (looks like a disk)
- 1GB to 1TB in size
- Raw block device,
  - Put your own filesystem on it
  - Do anything else that you would normally do to disk(s)
- Persistent & snapshot capable
- Mount to any EC2 instance in availability zone
- Notable enhancements:
  - Create EBS volumes from hosted AWS datasets
  - EBS snapshot share
    - Can be used to clone/create/share volume data

BIOTEAM
Enabling Science

Simple Queue Service ("SQS")

- One of the key "glue" services for workflows
  - Message passing between AMI instances
  - Cheap, flexible, reliable
  - Can add new message at any time
    - 8KB size; any format
  - Messages are locked while being processed
  - If read fails, lock is removed
    - Message free to be re-read

BIOTEAM
Enabling Science

Elastic MapReduce

- ** I have not used this service*
- Integrated Hadoop processing solution
- Has caused some controversy
- Designed to make life easier for people who do not want to custom build their own Hadoop systems within AWS

BIOTEAM
Enabling Science

Virtual Private Cloud (“VPC”)

- ** I have not used this service yet*
- Relatively new product offering
- Very interesting to me
- Solves some nasty problems with cloud-bursting and other hybrid local/cloud solutions
  - Different networks, IP address schemes and subnets can be a problem when “bridging” local and cloud systems
  - Most people doing this today implement an OpenVPN software overlay network to unify the network space
  - Amazon VPS essentially makes this a formal, supported product

BIOTEAM
Enabling Science

Cloud Sobriety

Important to think in practical terms. Utility computing has just as many negatives as positives.

BIOTEAM
Enabling Science

Cloud Sobriety

McKinsey presentation "Clearing the Air on Cloud Computing" is a must-read

- Tries to deflate the hype a bit
- James Hamilton has a nice reaction:
  - http://perspectives.mvdirona.com/

Both conclude:

- IT staff needs to understand "the cloud"
- Critical to quantify your own internal costs
- Perform your own due diligence

BIOTEAM
Enabling Science

Cloud Security

… set mindset to ‘cynical’

BIOTEAM
Enabling Science

Cloud Security Pet Peeve

- Don’t want to belittle security concerns, but …
- A whiff of hypocrisy is in the air
  - Is your staff *really* concerned or just protecting turf?
  - It is funny to see people demanding security measures that they don’t practice internally across their own infrastructure

BIOTEAM
Enabling Science

Cloud Security Pet Peeve

- My personal take:
  - Amazon, Google & Microsoft quite probably have better internal operating controls than you do
  - All of them are happy to talk as deeply as you like about all issues relating to security
  - Do your own due diligence & don't let politics or IT empire issues cloud decision making
  - *Biggest issue for me may be per-country data protection and patient privacy rules*

http://aws.amazon.com/security/

BIOTEAM
Enabling Science

State of AWS

The good, the bad, the ugly & what it means for HPC types

BIOTEAM
Enabling Science

State of Amazon AWS

New features are being rolled out fast and furious

But …

- EC2 nodes still poor on disk IO operations
- EBS service can use some enhancements
  - Many readers, one-writer on EBS volumes would be fantastic
- Poor support for latency-sensitive things and workflows that prefer tight network topologies

This matters because:

- Compute power is easy to acquire
- Life science tends to be IO bound
- Life science is currently being buried in data

BIOTEAM
Enabling Science

AWS & Internet Networking

- Can be challenging
- EC2 nodes use private IP address space
- EC2 nodes have unique public IP endpoints but do not "know" them at boot time
  - Internet data (to/from) comes via NAT
  - This breaks some software and services
  - Easy to workaround though …
    1. Query instance reservation to learn public hostname for given instance
    2. Perform DNS query on your public hostname to learn your public IP
- This is why so many cloud solutions implement their own software based VPN layers

BIOTEAM
Enabling Science

AWS & HPC Networking

- No guarantee that all your EC2 reservation instances will be allocated from the same subnet
- You really only have control over what availability zones you start your EC2 systems in
- This really freaks out OpenMPI and other HPC stacks that make implicit assumptions about subnets and the Layer 2 environment
- * Very likely to change in the future though

BIOTEAM
Enabling Science

HPC & AWS: Whole new world

- **For cluster people some radical changes**
  Years spent tuning systems for shared access
  - Utility model offers *dedicated* resources
  - EC2 not architected for *our* needs
  - Best practices & reference architectures will change
- **Current State: Transition Period**
  - Still hard to achieve seamless integration with local clusters & remote utility clouds
  - Most people are moving entire workflows into the cloud rather than linking grids
  - Some work being done on ‘transfer queues’

BIOTEAM
Enabling Science

HPC & AWS Summary

- Virtualized networking is ‘reasonable’ but there are certainly issues that need to be worked around
- Network latency can be high
- Virtualized storage I/O is far slower than anything we can do with local resources. Absolute fact.
- Still hard to share data/storage across many systems
- Inability to currently request EC2 nodes that are “close” in network topology terms is problematic (but likely to change)
- MapReduce is not a viable solution for everyone
- Amazon has a deep interest in HPC workflows, expect them to address all of our concerns

BIOTEAM
Enabling Science

Cloud Data Movement

Lessons Learned

BIOTEAM
Enabling Science

20TB Cloud Export Project

- One of my favorite ‘09 consulting projects …
  - Move 20TB scientific data out of Amazon S3 storage cloud
- What we experienced:
  - Significant human effort to swap/transport disks
  - Wrote custom DB and scripts to verify all files each time they moved
    - Avg. 22-50 MB/sec download from internet
    - Avg. 60MB/sec server to portable SATA array
    - Avg. 11MB/sec portable SATA to portable NAS array
  - At 11MB/sec, moving 20TB is a matter of *weeks*
  - *Forgot to account for MD5 checksum calculation times*
- Result:
  - Lesson Learned: data movement & handling took 5x longer than data acquisition

BIOTEAM
Enabling Science

Export via 300 mbit/sec circuit

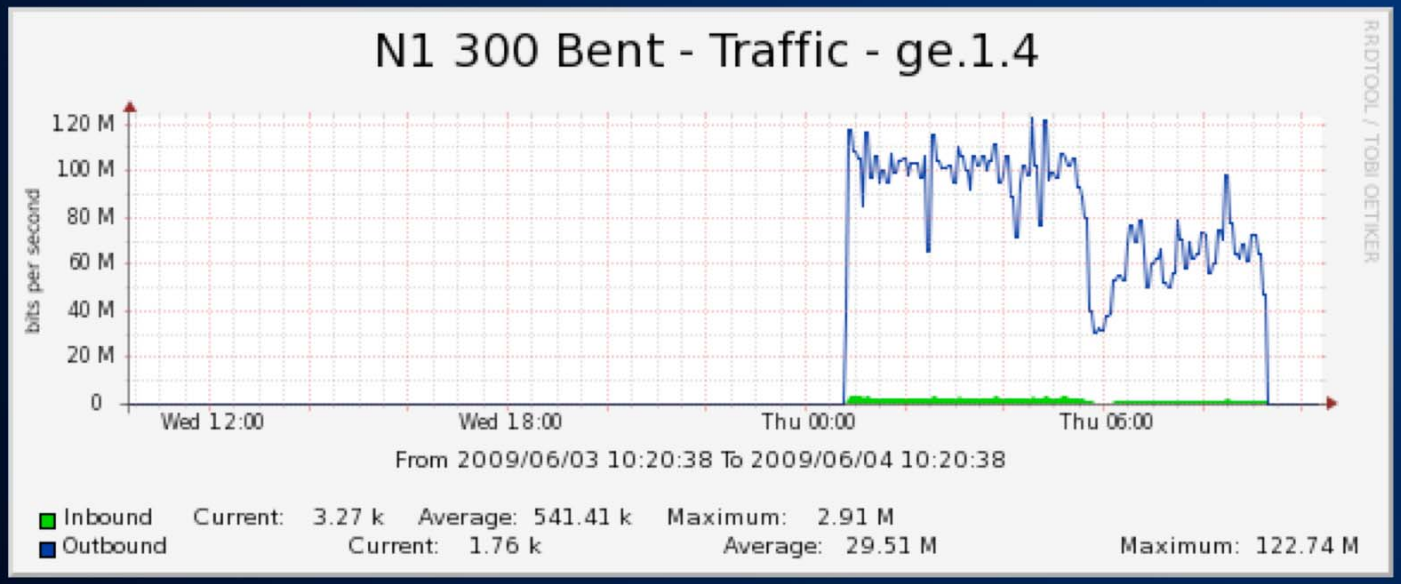

BIOTEAM
Enabling Science

Export via Gigabit internet link

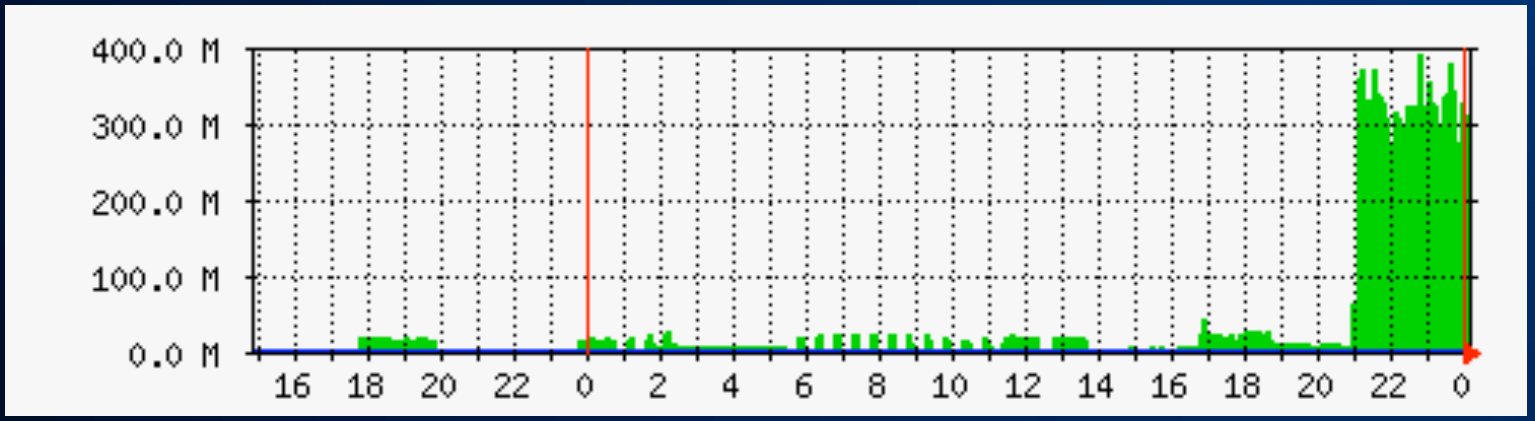

BIOTEAM
Enabling Science

Export via Gigabit internet link

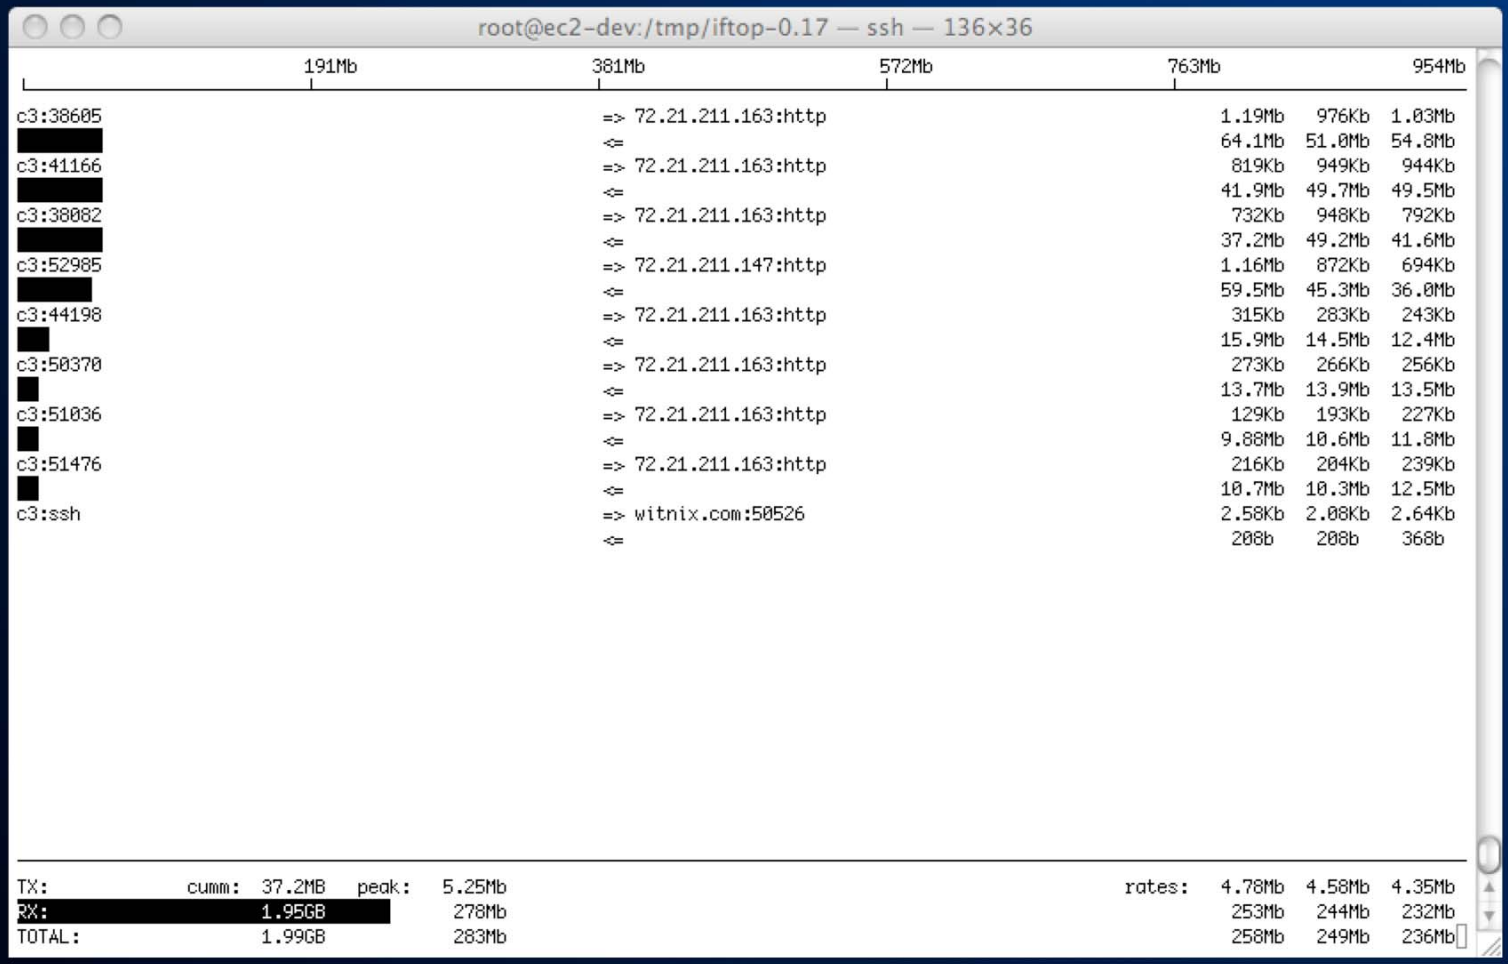

BIOTEAM
Enabling Science

S3 Bulk Download Lessons

- Your location matters, testing required
- *50 megabytes per second sustained was easily reached via both 300 mbit & 1 GbE circuits*
  - For us, we hit limitations of our download server, disks, memory and Java download code
  - Also hit limits caused by hundreds of thousands of small files to download
  - Clear that we, not Amazon were the bottleneck
- However, Guy Coats reports 10% utilization of network link when his group tested in the UK

BIOTEAM
Enabling Science

Some Real World Examples

Brief looks at some 2009 AWS projects …

BIOTEAM
Enabling Science

Rapid Prototyping & Development

- Easiest and most effective use for AWS for many of us today
- Take advantage of the absolute simplicity of rapidly deploying and destroying EC2 systems on demand
- Use this for
  - Spinning up development environments
  - Spinning up evaluation/testbeds
  - Pilot programs & training environments

BIOTEAM
Enabling Science

Prototyping & Development

- Why use AWS for this?
  - Provision new systems in minutes, not days, weeks or months
  - Spend operating funds, not capital money
  - Delegate provisioning tasks to end-users
  - BioTeam does this for training, testing & development
  - Pfizer does this and speaks publicly about it
- May be an ideal starting point for people wanting to “test the cloud”

BIOTEAM
Enabling Science

Self-organizing Compute Farms

- Build SGE/LSF clusters within the cloud for cloud-bursting, dedicated workflows or testing
- Our simple Grid Engine method
  1. Start reservation with N nodes
  2. All nodes have a firstboot script
  3. At boot, sort reservation instance names alphabetically
  4. First instance becomes SGE qmaster
  5. All other nodes know then to self-configure as execution hosts that bind to the first instance name
- Primary issue: random EC2 startup order needs to be handled

BIOTEAM
Enabling Science

Protein Engineering w/ AWS

A real pharma example

BIOTEAM
Enabling Science

Protein Engineering with AWS

- **Pfizer Biotherapeutics & Bioinnovation Center**
  - **Giles Day, Pfizer**
  - **Adam Kraut, BioTeam**
- **Problem:**
  - Antibody models can be created in a few hours on a standard workstation
  - Full-atom refinement of each model using Rosetta++ requires 1000 CPU hours
  - 2-3 months required *per-model* on existing Pfizer research cluster
  - Cluster subject to unpredictable loads

BIOTEAM
Enabling Science

Protein Engineering with AWS

- **1000 CPU Hour Antibody Refinement Problem**
  - **Using Rosetta++ (Davd Baker, UWash)**
- **Huge Opportunity for Pfizer:**
  - Deliver antibody model refinement results in *one day* rather than 2-3 months
- **Ideal AWS Candidate:**
  - CPU bound
  - Low data I/O requirements
  - Free up cluster for I/O bound workloads

BIOTEAM
Enabling Science

Protein Engineering with AWS

- Borrows heavily from RightScale & AWS published best practices
- Inbound/Outbound SQS queues
- Job specification in JSON format
- Data for each work unit in S3 buckets
- Custom EC2 AMI
- Workers pull from S3, push back when finished
- Job provenance/metadata stored in SimpleDB
- Independent work units allow dynamic allocation of Worker instances

BIOTEAM
Enabling Science

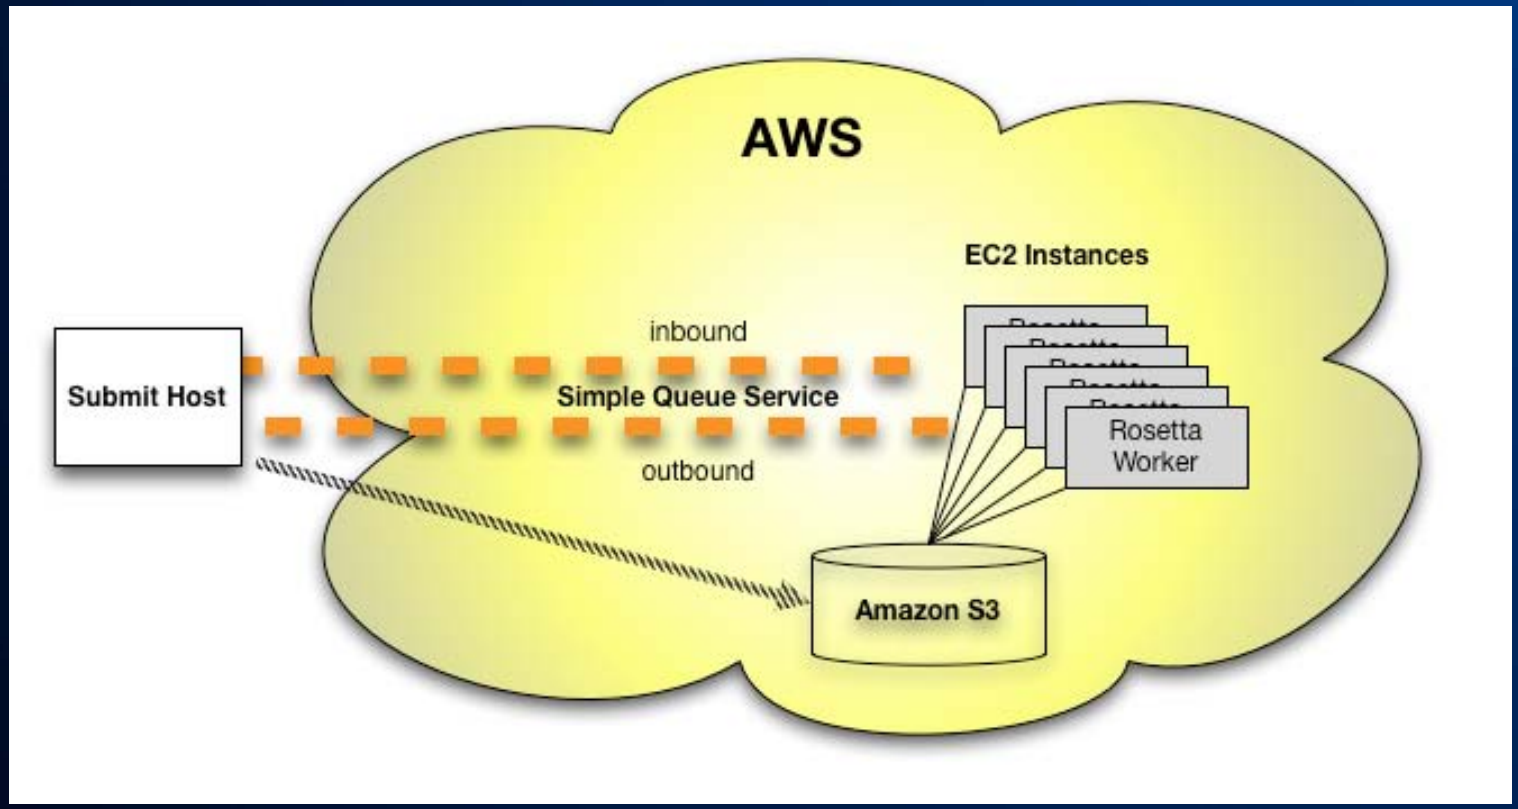

BIOTEAM
Enabling Science

Getting Hypothetical …

Potential Use-case for archival/cold storage with ability to perform re-analysis if needed

BIOTEAM
Enabling Science

Bulk Data Ingest/Export

- How do we move 1TB/day into the cloud?
  - Not very easily
  - Now that AWS Import/Export has launched we might have some options
- Our field is looking for answers
  - Need “cheap and deep” store(s)
  - Currently buried by lab instruments that produce TB/day volumes
    - Next-Gen DNA Sequencing
    - 3D Ultrasound & other imaging
    - Confocal microscopy
    - Etc.

BIOTEAM
Enabling Science

Cloud Storage

- It is quite probable that the “internet-scale” providers can provide storage far more cheaply than we can ourselves
  - *Especially if we are honest about facility, power, continuity and operational costs*
- Some people estimate cost at .80 GB/year and falling fast for Amazon and others to provide 3x geographically replicated raw storage
  - *Can you seriously match this?*
- These prices come from operating at extreme efficiency scales that we will never be able to match ourselves
- Question: how best to leverage this?

BIOTEAM
Enabling Science

When ingest problem is solved …

- I think there may be petabytes of life science data that would flock to utility storage services
  - Public and private data stores
  - Mass amount of grant funded study data
  - Archive store, HSM target and DR store
  - “Downloader Pays” model is compelling for people required to share large data sets

BIOTEAM
Enabling Science

Terabyte Wet Lab Instrument

BIOTEAM
Enabling Science

Cautionary Tale: 180TB kept on desk

The life science “data tsunami” is no joke

BIOTEAM
Enabling Science

Next-Gen & Potential AWS use

What this would mean:

- Primary analysis onsite; data moved into remote utility storage service after passing QC tests
- Data would rarely (if ever) move back
- Need to reprocess or rerun?
  - Spin up “cloud” servers and re-analyze in situ
  - Terabyte data transit not required

Summary:

- Lifesci data; 1-way transit into the cloud
- Archive store or public/private repository
- Any re-study or reanalysis primarily done in situ
- Downside: replicating pipelines & workflows remotely
- Careful attention must be paid to costs

BIOTEAM
Enabling Science

Wrapping Up

Advice for effective cloud utilization

BIOTEAM
Enabling Science

First Principal

- Economics play a critical role in cloud decisions
- You MUST have a very solid understanding of your own internal IT operating costs for CPU, network, storage & operation
- Without accurate internal cost data, cloud decisions may be made unwisely

BIOTEAM
Enabling Science

Second Principal

- Understand that this is a very hyped & trendy area
- Need to be cynical and focused on actual value
- Cloud fanatics are just as dangerous as cloud luddites
- Understand cloud strengths and weaknesses so that sensible decisions can be made about priorities and focus

BIOTEAM
Enabling Science

Third Principal

- Start small, stay targeted
- Go for the easy wins first
- But don’t fail to test out the complicated stuff
- Key areas to understand and investigate
  - AWS storage performance (S3 & EBS)
  - AWS data movement
  - AWS networking internals

BIOTEAM
Enabling Science

Fourth Principal

- Optimization matters
  - There are “good” and “bad” ways to develop & deploy on AWS
  - *Constantly re-bundling AMIs is a “bad” thing*
- Don’t reinvent the wheel if you don’t have to
  - Many interesting startup companies in this space
  - Providing dashboards, accounting, scaling, monitoring, workflow automation and administration frameworks
- Companies I watch in this space:
  - RightScale Inc.
  - Cycle Computing
  - UnivaUD

BIOTEAM
Enabling Science

End;

- Thanks!
- Any questions?
- Comments/feedback:
  - chris@bioteam.net

BIOTEAM
Enabling Science