

Beyond the Genome: Cloud-scale computing demo

Michael Schatz, Ben Langmead, & James Taylor

Sept. 19, 2011

Beyond the Genome

Beyond the Genome Challenge

<http://schatzlab.cshl.edu/data/btg11.tgz>

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

The goal is to identify a viral sequence insertion into a human cancer exome. To keep it tractable, we will only use genes on chromosome 22, and only exons > 500bp long.

If you have questions, tweet #btg11

Submit your solution to: mschatz@cshl.edu

The subject line should be: **BTG2011 human_gene virus_name**

The body should contain all the steps you took to identify the gene and virus. If at all possible, please include the exact commands used. Winners will be selected by first correct answer (name of gene, name of virus) and for reproducibility. You must be registered and present at Beyond the Genome 2011 to win. The judges decisions are final. Rules are subject to change at anytime.



Amazon Web Services

<http://aws.amazon.com>

- All you need is a credit card, and you can immediately start using one of the largest datacenters in the world
- Elastic Compute Cloud (EC2)
 - On demand computing power
 - Support for Windows, Linux, & OpenSolaris
 - Starting at 8.5¢ / core / hour
- Simple Storage Service (S3)
 - Scalable data storage
 - 10¢ / GB upload fee, 15¢ / GB monthly fee
- Plus many others



EC2 Architecture

- Very large pool of machines
 - Effectively infinite resources
 - High-end servers with many cores and many GB RAM
- Machines run in a virtualized environment
 - Amazon can subdivide large nodes into smaller instances
 - You are 100% protected from other users on the machine
 - You get to pick the operating system, all installed software



Amazon Machine Images



- A few Amazon sponsored images
 - Suse Linux, Windows
- Many Community Images & Appliances
 - CloudBioLinux: Genomics Appliance
 - Crossbow: Hadoop, Bowtie, SOAPsnp
 - Galaxy: CloudMan
- Build you own
 - Completely customize your environment
 - You results could be totally reproducible

Amazon S3

- S3 provides persistent storage for large volumes of data
 - Very high speed connection from S3 to EC2 compute nodes
 - Public data sets include `s3://1000genomes`
- Tiered pricing by volume
 - Pricing starts at 14¢ / GB / month
 - 5.5¢ / GB / month for over 5 PB
 - Pay for transfer out of Amazon
- Import/Export service for large volumes
 - FedEx your drives to Amazon



Getting Started

<http://docs.amazonwebservices.com/AWSEC2/latest/GettingStartedGuide/>

The screenshot shows a web browser window titled "Amazon Elastic Compute Cloud" with the URL <http://docs.amazonwebservices.com/AWSEC2/latest/GettingStartedGuide/>. The page header includes the Amazon Elastic Compute Cloud logo and the text "Getting Started Guide (API Version 2010-08-31)". A navigation sidebar on the left lists several topics, with "Get Started with EC2" highlighted. The main content area features a "Documentation Feedback" link, a "Welcome" message, and the section "Get Started with EC2". Below this section is a paragraph explaining Amazon EC2 and a flowchart illustrating the steps: "Sign up for EC2" leads to "Launch instance", which then branches into "Connect to Linux/UNIX instance" and "Connect to Windows instance", both of which lead to "Terminate instance". At the bottom of the page, there is a "Get Started" button with a play icon.

Amazon Elastic Compute Cloud
Getting Started Guide (API Version 2010-08-31)

Documentation Feedback

Welcome

Get Started with EC2

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that enables you to launch and manage Linux/UNIX and Windows server instances in Amazon's data centers. You can get started with Amazon EC2 by following the tasks shown in the following diagram. You'll primarily use the AWS Management Console, a point-and-click web-based interface.

```
graph LR; A[Sign up for EC2] --> B[Launch instance]; B --> C[Connect to Linux/UNIX instance]; B --> D[Connect to Windows instance]; C --> E[Terminate instance]; D --> E;
```

This guide walks you through launching and connecting to your first Amazon EC2 instance. To start, click the following **Get Started** button.

[Get Started](#)

Signing Up

The screenshot shows the Amazon Web Services website in a browser window. The browser's address bar displays `http://aws.amazon.com/`. The page features the AWS logo and navigation links for [Sign in to the AWS Management Console](#), [Create an AWS Account](#), and [English](#). A horizontal menu includes [AWS](#), [Products](#), [Developers](#), [Community](#), [Support](#), and [Account](#).

The main content area highlights the **Introducing Amazon Simple Notification Service** with the text: "Enable applications, end-users, and devices to instantly send and receive notifications from the cloud." Below this is a "Learn More..." link and a diagram showing a server icon connected to a cloud, which then connects to icons for an email, a globe (WWW), and a mobile phone.

On the right side, there is a call to action: "Sign up for a free Amazon Web Services Account" with a prominent **Sign Up Now** button.

The **Get Started** section includes a **Business Managers** heading and the text: "Learn how Amazon Web Services enables you to reach business goals faster:". Below this is a list of links: [Solutions & Use Cases](#), [Security Center](#), [Economics Center](#), [Case Studies](#), [Service Health Dashboard](#), [Solution Providers](#), and [Videos & Webinars](#).

The **News & Events** section has three tabs: **What's New?**, **Media Coverage**, and **Upcoming Events**. Under "What's New?", there are three news items:

Date	Event/Update
Oct 14, 2010	Amazon Elastic Load Balancing Adds Support for HTTPS
Oct 12, 2010	AWS Management Console Adds Support for Amazon SNS
Oct 05, 2010	Amazon EC2 Running SUSE Linux Now Available

There are also two "Upcoming Events" listed:

Date	Event
Oct 05, 2010	Announcing Read Replicas, Lower High Memory DB Instance Prices for Amazon RDS
Sep 29, 2010	Announcing the AWS SDK for PHP
Sep 21, 2010	Oracle Certifies Enterprise Software on Amazon EC2

At the bottom of the news section, there is an [RSS](#) link and a [View all](#) link.

AWS Management Console

The screenshot displays the AWS Management Console interface for the Amazon EC2 service. The browser address bar shows the URL <https://console.aws.amazon.com/ec2/home>. The top navigation bar includes the AWS logo, navigation links (Products, Developers, Community, Support, Account), and user information (Welcome, CBCB, Settings, Sign Out).

The main content area is titled "Amazon EC2 Console Dashboard" and is divided into several sections:

- Navigation:** A sidebar on the left provides a menu for navigating through various EC2 resources, including INSTANCES (Instances, Spot Requests), IMAGES (AMIs, Bundle Tasks), ELASTIC BLOCK STORE (Volumes, Snapshots), and NETWORKING & SECURITY (Elastic IPs, Security Groups, Placement Groups, Load Balancers, Key Pairs). The region is set to "US East".
- Getting Started:** A yellow box provides instructions on how to launch an Amazon EC2 instance, with a prominent "Launch Instance" button. A note specifies that instances will launch in the US East (Virginia) region.
- My Resources:** A summary of resources in the US East region, including 0 Running Instances, 0 Elastic IPs, 0 EBS Volumes, 0 EBS Snapshots, 2 Key Pairs, 27 Security Groups, 0 Load Balancers, and 0 Placement Groups. A "Refresh" button is available.
- Service Health:** A table showing the current status of the Amazon EC2 service in the US East - N. Virginia region, which is operating normally. A link is provided to view complete service health details.
- Related Links:** A list of links for documentation, all EC2 resources, forums, feedback, and reporting an issue.

At the bottom of the page, the copyright notice reads: © 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | [Feedback](#) | [Support](#) | [Privacy Policy](#) | [Terms of Use](#) | An [amazon.com](#) company

Running your First Cloud Analysis

1. Pick your AMI
 - Machine Image: Operating System & Tools
2. Pick your instance type & quantity
 - Micro - High-Memory Quadruple Extra Large
3. Pick your credentials
 - SSH Keys
4. Configure your Firewall
 - Protect your servers
5. Launch!

I. Pick your AMIs






The screenshot shows the AWS Management Console interface with the 'Request Instances Wizard' modal open. The wizard is at the 'CHOOSE AN AMI' step. The background shows the console navigation pane on the left and the wizard's progress bar at the top. The wizard has five tabs: 'CHOOSE AN AMI', 'INSTANCE DETAILS', 'CREATE KEY PAIR', 'CONFIGURE FIREWALL', and 'REVIEW'. Below the tabs, there is a list of AMIs with their logos, names, descriptions, and 'Select' buttons.

Request Instances Wizard Cancel

CHOOSE AN AMI INSTANCE DETAILS CREATE KEY PAIR CONFIGURE FIREWALL REVIEW

Choose an Amazon Machine Image (AMI) from one of the tabbed lists below by clicking its **Select** button.

Quick Start My AMIs Community AMIs

	Basic 32-bit Amazon Linux AMI 1.0 (AMI Id: ami-3ac33653) Amazon Linux AMI Base 1.0, EBS boot, 32-bit architecture with Amazon EC2 AMI Tools.	Select
	Basic 64-bit Amazon Linux AMI 1.0 (AMI Id: ami-38c33651) Amazon Linux AMI Base 1.0, EBS boot, 64-bit architecture with Amazon EC2 AMI Tools.	Select
	SUSE Linux Enterprise Server 11 32-bit (AMI Id: ami-e0a35789) SUSE Linux Enterprise Server 11 Service Pack 1 basic install, EBS boot, 32-bit architecture with Amazon EC2 AMI Tools preinstalled; Apache 2.2, MySQL 5.0, PHP 5.3, Ruby 1.8.7, and Rails 2.3	Select
	SUSE Linux Enterprise Server 11 64-bit (AMI Id: ami-e4a3578d) SUSE Linux Enterprise Server 11 Service Pack 1 basic install, EBS boot, 64-bit architecture with Amazon EC2 AMI Tools preinstalled; Apache 2.2, MySQL 5.0, PHP 5.3, Ruby 1.8.7, and Rails 2.3	Select
	Getting Started on Microsoft Windows Server 2008 (AMI Id: ami-c5e40dac) Microsoft Windows Server 2008 R1 SP2 Datacenter edition, 32-bit architecture, Microsoft SQLServer 2008 Express, Internet Information Services 7, ASP.NET 3.5.	Select

AMI Launch Index: 0 Elastic IP: -
Root Device: /dev/sda1 Root Device Type: ebs

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | Feedback | Support | Privacy Policy | Terms of Use | An amazon.com company

CloudBioLinux

The screenshot shows the AWS Management Console interface. The main navigation bar includes the AWS logo, navigation links (Products, Developers, Community, Support, Account), and user information (Welcome, CBCB, Settings, Sign Out). The secondary navigation bar shows various AWS services, with Amazon EC2 selected. The 'Request Instances Wizard' modal is open, displaying the 'CHOOSE AN AMI' step. It features a progress bar and five steps: CHOOSE AN AMI, INSTANCE DETAILS, CREATE KEY PAIR, CONFIGURE FIREWALL, and REVIEW. Below the progress bar, there is a search bar with 'cloudbiolinux' entered and a 'Viewing: All Images' dropdown. A table lists four CloudBioLinux AMIs with their IDs, root devices, manifests, and platforms. Each row has a 'Select' button with a right-pointing arrow.

AMI ID	Root Device	Manifest	Platform	
ami-0af91263	ebs	678711657553/CloudBioLinux 32bit 20100716	Other Linux	Select
ami-4e57a227	ebs	678711657553/CloudBioLinux 64bit 20100929	Other Linux	Select
ami-6953b200	instance-store	jcvicloudbiolinux/JCVI-Cloud-BioLinux.manifest.xml	Other Linux	Select
ami-879c75ee	ebs	678711657553/CloudBioLinux 20100507	Other Linux	Select

2. Pick your Instance Type

Request Instances Wizard [Cancel]

CHOOSE AN AMI | **INSTANCE DETAILS** | CREATE KEY PAIR | CONFIGURE FIREWALL | REVIEW

Provide the details for your instance(s). You may also decide whether you want to launch your instances as "on-demand" or "spot" instances.

Number of Instances: **Availability Zone:**

Instance Type:

Note, launching a t1.micro instance will require you to specify an Elastic IP address.

Type	CPU Units	CPU Cores	Memory
Large (m1.large)	4 ECUs	2 Cores	7.5 GB
Extra Large (m1.xlarge)	8 ECUs	4 Cores	15 GB
High-Memory Extra Large (m2.xlarge)	6.5 ECUs	2 Cores	17.1 GB
High-Memory Double Extra Large (m2.2xlarge)	13 ECUs	4 Cores	34.2 GB
High-Memory Quadruple Extra Large (m2.4xlarge)	26 ECUs	8 Cores	68.4 GB
High-CPU Extra Large (c1.xlarge)	20 ECUs	8 Cores	7 GB

Launch Instances Request Spot Instance Launch On-Demand Instance

Block Devices: /dev/sda1=vol-7c8a8415:attached:2010-10-20T02:41:57.000Z:true

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | [Feedback](#) | [Support](#) | [Privacy Policy](#) | [Terms of Use](#) | An [amazon.com](#) company

3. Pick your Credentials

The screenshot shows the AWS Management Console interface. A modal dialog titled "Request Instances Wizard" is open, showing the "CREATE KEY PAIR" step. The wizard has a progress bar with five steps: CHOOSE AN AMI, INSTANCE DETAILS, CREATE KEY PAIR (current), CONFIGURE FIREWALL, and REVIEW. Below the progress bar, there is a paragraph of text explaining key pairs. The main content area has three sections: "Choose from your existing Key Pairs" (unselected), "Create a new Key Pair" (selected), and "Proceed without a Key Pair" (unselected). Under "Create a new Key Pair", there are two numbered steps: "1. Enter a name for your key pair:*" with a text input field containing "mschatz" and "(e.g., jdoekey)", and "2. Click to create your key pair:*" with a yellow button labeled "Create & Download your Key Pair". Below the button is a tip: "Save this file in a place you will remember. You can use this key pair to launch other instances in the future or visit the Key Pairs page to create or manage existing ones." At the bottom of the dialog are "Back" and "Continue" buttons. The background shows the console navigation pane on the left and the top navigation bar with the user name "Welcome, CBCB".

Request Instances Wizard

CHOOSE AN AMI | INSTANCE DETAILS | **CREATE KEY PAIR** | CONFIGURE FIREWALL | REVIEW

Public/private key pairs allow you to securely connect to your instance after it launches. To create a key pair, enter a name and click **Create & Download your Key Pair**. You will then be prompted to save the private key to your computer. Note, you only need to generate a key pair once - not each time you want to deploy an Amazon EC2 instance.

Choose from your existing Key Pairs

Create a new Key Pair

1. Enter a name for your key pair:* (e.g., jdoekey)

2. Click to create your key pair:* [Create & Download your Key Pair](#)

Save this file in a place you will remember. You can use this key pair to launch other instances in the future or visit the Key Pairs page to create or manage existing ones.

Proceed without a Key Pair

< Back | Continue >

Root Device: /dev/sda1 | Root Device Type: ebs

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | Feedback | Support | Privacy Policy | Terms of Use | An amazon.com company

4. Configure your Firewall

The screenshot shows the AWS Management Console interface with the 'Request Instances Wizard' dialog box open. The wizard is currently on the 'CONFIGURE FIREWALL' step. The dialog box contains the following information:

Request Instances Wizard [Cancel]

CHOOSE AN AMI | INSTANCE DETAILS | CREATE KEY PAIR | **CONFIGURE FIREWALL** | REVIEW

Security groups determine whether a network port is open or blocked on your instances. You may use an existing security group, or we can help you create a new security group to allow access to your instances using the suggested ports below. Add additional ports now or update your security group anytime using the Security Groups page. All changes take effect immediately.

Choose one or more of your existing Security Groups

Create a new Security Group

1. Name your Security Group:

2. Describe your Security Group:

3. Define allowed Connections

Application	Transport	Port	Source Network (IPv4 CIDR)	Actions
HTTP	TCP	80	All Internet	<input type="button" value="Remove"/>
SSH	TCP	22	All Internet	<input type="button" value="Remove"/>
Select...	-	-	All Internet Change	<input type="button" value="Add Rule"/>

< Back

Block Devices: /dev/sda1=vol-7c8a8415:attached:2010-10-20T02:41:57.000Z:true

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | Feedback | Support | Privacy Policy | Terms of Use | An amazon.com company

5. Launch!

The screenshot shows the AWS Management Console interface with the 'Request Instances Wizard' modal open. The wizard is in the 'REVIEW' step, indicated by a progress bar at the top. The configuration details are as follows:

- AMI:** Amazon Linux AMI ID ami-3ac33653 (i386)
 - Name:** Basic 32-bit Amazon Linux AMI 1.0
 - Description:** Amazon Linux AMI Base 1.0, EBS boot, 32-bit architecture with Amazon EC2 AMI Tools.
 - [Edit AMI](#)
- Number of Instances:** 1
 - Availability Zone:** No Preference
 - Instance Type:** Micro (t1.micro)
 - Instance Class:** On Demand
 - [Edit Instance Details](#)
- Monitoring:** Disabled
 - Kernel ID:** Use Default
 - RAM Disk ID:** Use Default
 - User Data:**
 - [Edit Advanced Details](#)
- Key Pair Name:** mschatz
 - [Edit Key Pair](#)
- Security Group(s):** basic
 - [Edit Firewall](#)

At the bottom of the wizard, there is a '< Back' button and a 'Launch' button with a right-pointing arrow.

Below the wizard, the 'Root Device' is set to '/dev/sda1' and the 'Root Device Type' is 'ebs'.

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | [Feedback](#) | [Support](#) | [Privacy Policy](#) | [Terms of Use](#) | An [amazon.com](#) company

Monitoring your Server

The screenshot shows the AWS Management Console interface for the Amazon EC2 service. The main content area displays a table of instances. The table has the following data:

Name	Instance	AMI ID	Root Device	Type	Status	Security Groups	Key Pair Name
empty	i-4aab6027	ami-3ac33653	ebs	t1.micro	running	crossbow46474683	gsg-keypair

Below the table, it indicates "0 EC2 Instances selected" and provides the instruction "Select an instance above".

The footer of the console contains the following text: © 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | [Feedback](#) | [Support](#) | [Privacy Policy](#) | [Terms of Use](#) | An [amazon.com](#) company

Connecting (I)

The screenshot displays the AWS Management Console interface. At the top, the browser address bar shows the URL `https://console.aws.amazon.com/ec2/home#c=EC2&s=Instances`. The navigation bar includes the AWS logo and links for Products, Developers, Community, Support, and Account. A welcome message for 'CBCB' is visible along with Settings and Sign Out options.

The main content area is titled 'My Instances' and shows a table of EC2 instances. One instance is selected, and a context menu is open over the 'Connect' option. The instance details are as follows:

Name	Instance	AMI ID	Root Device	Type	Status	Security Groups	Key Pair Name
empty	i-4aab6027	ami-3ac33653	ebs	t1.micro	running	crossbow46474683	gsg-keypair

The context menu includes the following options:

- Instance Management
 - Connect
 - Get System Log
 - Create Image (EBS AMI)
 - Add/Edit Tags
 - Launch More Like This
 - Disassociate IP Address
- Instance Lifecycle
 - Terminate
 - Reboot
 - Stop
 - Start
- CloudWatch Monitoring
 - Enable CloudWatch
 - Disable CloudWatch

The instance details pane shows the following information:

- AMI ID:** ami-3ac33653
- Security Groups:** crossbow464746833350-X-master
- Status:** running
- VPC ID:** -
- Virtualization:** para
- Reservation:** r-81ff3ceb
- Platform:** -
- Kernel ID:** aki-407d9529
- AMI Launch Index:** 0
- Root Device:** /dev/sda1
- Region:** us-east-1c
- Type:** t1.micro
- Owner:** 464746833350
- Subnet ID:** -
- Placement Group:** -
- RAM Disk ID:** -
- Key Pair Name:** gsg-keypair
- Monitoring:** disabled
- Elastic IP:** -
- Root Device Type:** ebs

At the bottom of the console, the footer contains copyright information: © 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. It also includes links for Feedback, Support, Privacy Policy, and Terms of Use, and identifies the company as An amazon.com company.

Connecting (2)

Connect Help - Secure Shell (SSH) Cancel X

Instance: i-4aab6027

To access your instance using any SSH client

1. Open the SSH client of your choice (e.g., PuTTY, terminal).
2. Locate your private key file, gsg-keypair.pem
3. Use chmod to make sure your key file isn't publicly viewable, ssh won't work otherwise:
`chmod 400 gsg-keypair.pem`
4. Connect to your instance using instance's public DNS [ec2-184-72-85-153.compute-1.amazonaws.com].

Example

Enter the following command line:

```
ssh -i gsg-keypair.pem root@ec2-184-72-85-153.com
```

Close

Optional. RAM disk associated with the image. For more information, refer to the Amazon Elastic Compute Cloud Developer Guide.

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | Feedback | Support | Privacy Policy | Terms of Use | An amazon.com. company

Calling SNPs in the Cloud 😊

```
chmod 400 mschatz.pem
```

```
scp -r -i mschatz.pem data.tgz ubuntu@ec2-174-129-123-73.compute-1.amazonaws.com:
```

```
ssh -i mschatz.pem ubuntu@ec2-174-129-123-73.compute-1.amazonaws.com
```

```
<remote>
```

```
ls
```

```
tar xzvf data.tgz
```

```
bowtie -S data/genomes/e_coli data/reads/e_coli_10000snp.fq ec_snp.sam
```

```
samtools view -bS -o ec_snp.bam ec_snp.sam
```

```
samtools sort ec_snp.bam ec_snp.sorted
```

```
samtools pileup -cv -f data/genomes/NC_008253.fna ec_snp.sorted.bam > snps
```

```
samtools index ec_snp.sorted.bam
```

```
samtools tview ec_snp.sorted.bam data/genomes/NC_008253.fna
```

```
exit
```

```
<local>
```

```
scp -i mschatz.pem ubuntu@ec2-174-129-123-73.compute-1.amazonaws.com:snps .
```

1000Genomes in the Cloud

```
s3cmd --configure
```

```
# cp data/.s3cfg .
```

```
s3cmd ls s3://1000genomes
```

```
s3cmd ls s3://1000genomes/Pilots_Bam/NA20828/
```

```
s3cmd get s3://1000genomes/Pilots_Bam/NA20828/*chr22* .
```

```
samtools view NA20828.SLX.maq.SRP000033.2009_09.chr22_1_49691432.bam
```

Terminating

The screenshot shows the AWS Management Console interface. At the top, the browser address bar displays the URL `https://console.aws.amazon.com/ec2/home#c=EC2&s=Instances`. The navigation bar includes the AWS logo and various service links like Amazon S3, Amazon EC2, Amazon VPC, etc. The main content area is titled "My Instances" and shows a table with one instance selected. A context menu is open over the instance, with the "Terminate" option highlighted. Below the table, the details for the selected instance "i-4aab6027" are displayed in a tabbed view under the "Description" tab.

Name	Instance	AMI ID	Root Device	Type	Status	Security Groups	Key Pair Name
<input checked="" type="checkbox"/>	empty	i-4aab6027	ami-2aa22652	t1.micro	running	crossbow46474683	gsg-keypair

1 EC2 Instance selected

EC2 Instance: i-4aab6027

Description | Monitoring | Tags

AMI ID: ami-2aa22652

Security Groups: crossbow464746833350-X-master

Status: running

VPC ID: -

Virtualization: para

Reservation: r-81ff3ceb

Platform: -

Kernel ID: aki-407d9529

AMI Launch Index: 0

Root Device: /dev/sda1

Region: us-east-1c

Type: t1.micro

Owner: 464746833350

Subnet ID: -

Placement Group: -

RAM Disk ID: -

Key Pair Name: gsg-keypair

Monitoring: disabled

Elastic IP: -

Root Device Type: ebs

© 2008 - 2009, Amazon Web Services LLC or its affiliates. All right reserved. | [Feedback](#) | [Support](#) | [Privacy Policy](#) | [Terms of Use](#) | An [amazon.com](#) company

Total cost: 8.5¢

Reflections

- Launching and managing virtual clusters with the AWS Console is quick and easy
 - Entirely scriptable using ec2 tools
 - iPhone App also available
- Things get really interesting on 168 cores
 - 1 week CPU = 1 hour wall

Just 3 commands to bring up a 168 core (21 node) cluster & crunch terabytes:

```
$HADOOP/src/contrib/ec2/bin/hadoop-ec2 launch-cluster HADOOP 21
```

```
$HADOOP/src/contrib/ec2/bin/hadoop-ec2 <hadoop cmd> HADOOP
```

```
$HADOOP/src/contrib/ec2/bin/hadoop-ec2 terminate-cluster HADOOP
```

Thank You!

<http://schatzlab.cshl.edu>
[@mike_schatz](#) / [#btg](#)