

# Scientific Computing with Amazon Web Services

Jamie Kinney

Director of Scientific Computing Amazon Web Services

[jkinney@amazon.com](mailto:jkinney@amazon.com)

[@jamiakinney](#)



# AWS Scientific Computing Team Focus

- Global “Big Science” Projects
- Enabling the “long tail of science”
- Collaborative research
- Education for our customers, Amazonians and Policy Makers
- Initial focus on:
  - Life Sciences
  - Earth Sciences
  - Astronomy/Astrophysics
  - High Energy Physics
  - Materials Science

# Why are we focusing on the Scientific Community?

- In order to meaningfully change our world for the better by accelerating the pace of scientific discovery
- Scientific computing is a profitable business for AWS
- To develop new capabilities which will benefit all AWS customers
  - Streaming data processing & analytics
  - Exabyte scale data management solutions
  - Collaborative research tools and techniques
  - New AWS regions
  - Significant advances in low-power compute, storage and data centers
  - Identify efficiencies which will lower our costs and pricing for customers
  - Push our existing services to support exabyte/exaflop scale workloads

2004

amazon.com

\$7B retail business  
~10,000 employees  
A whole lot of servers

2013




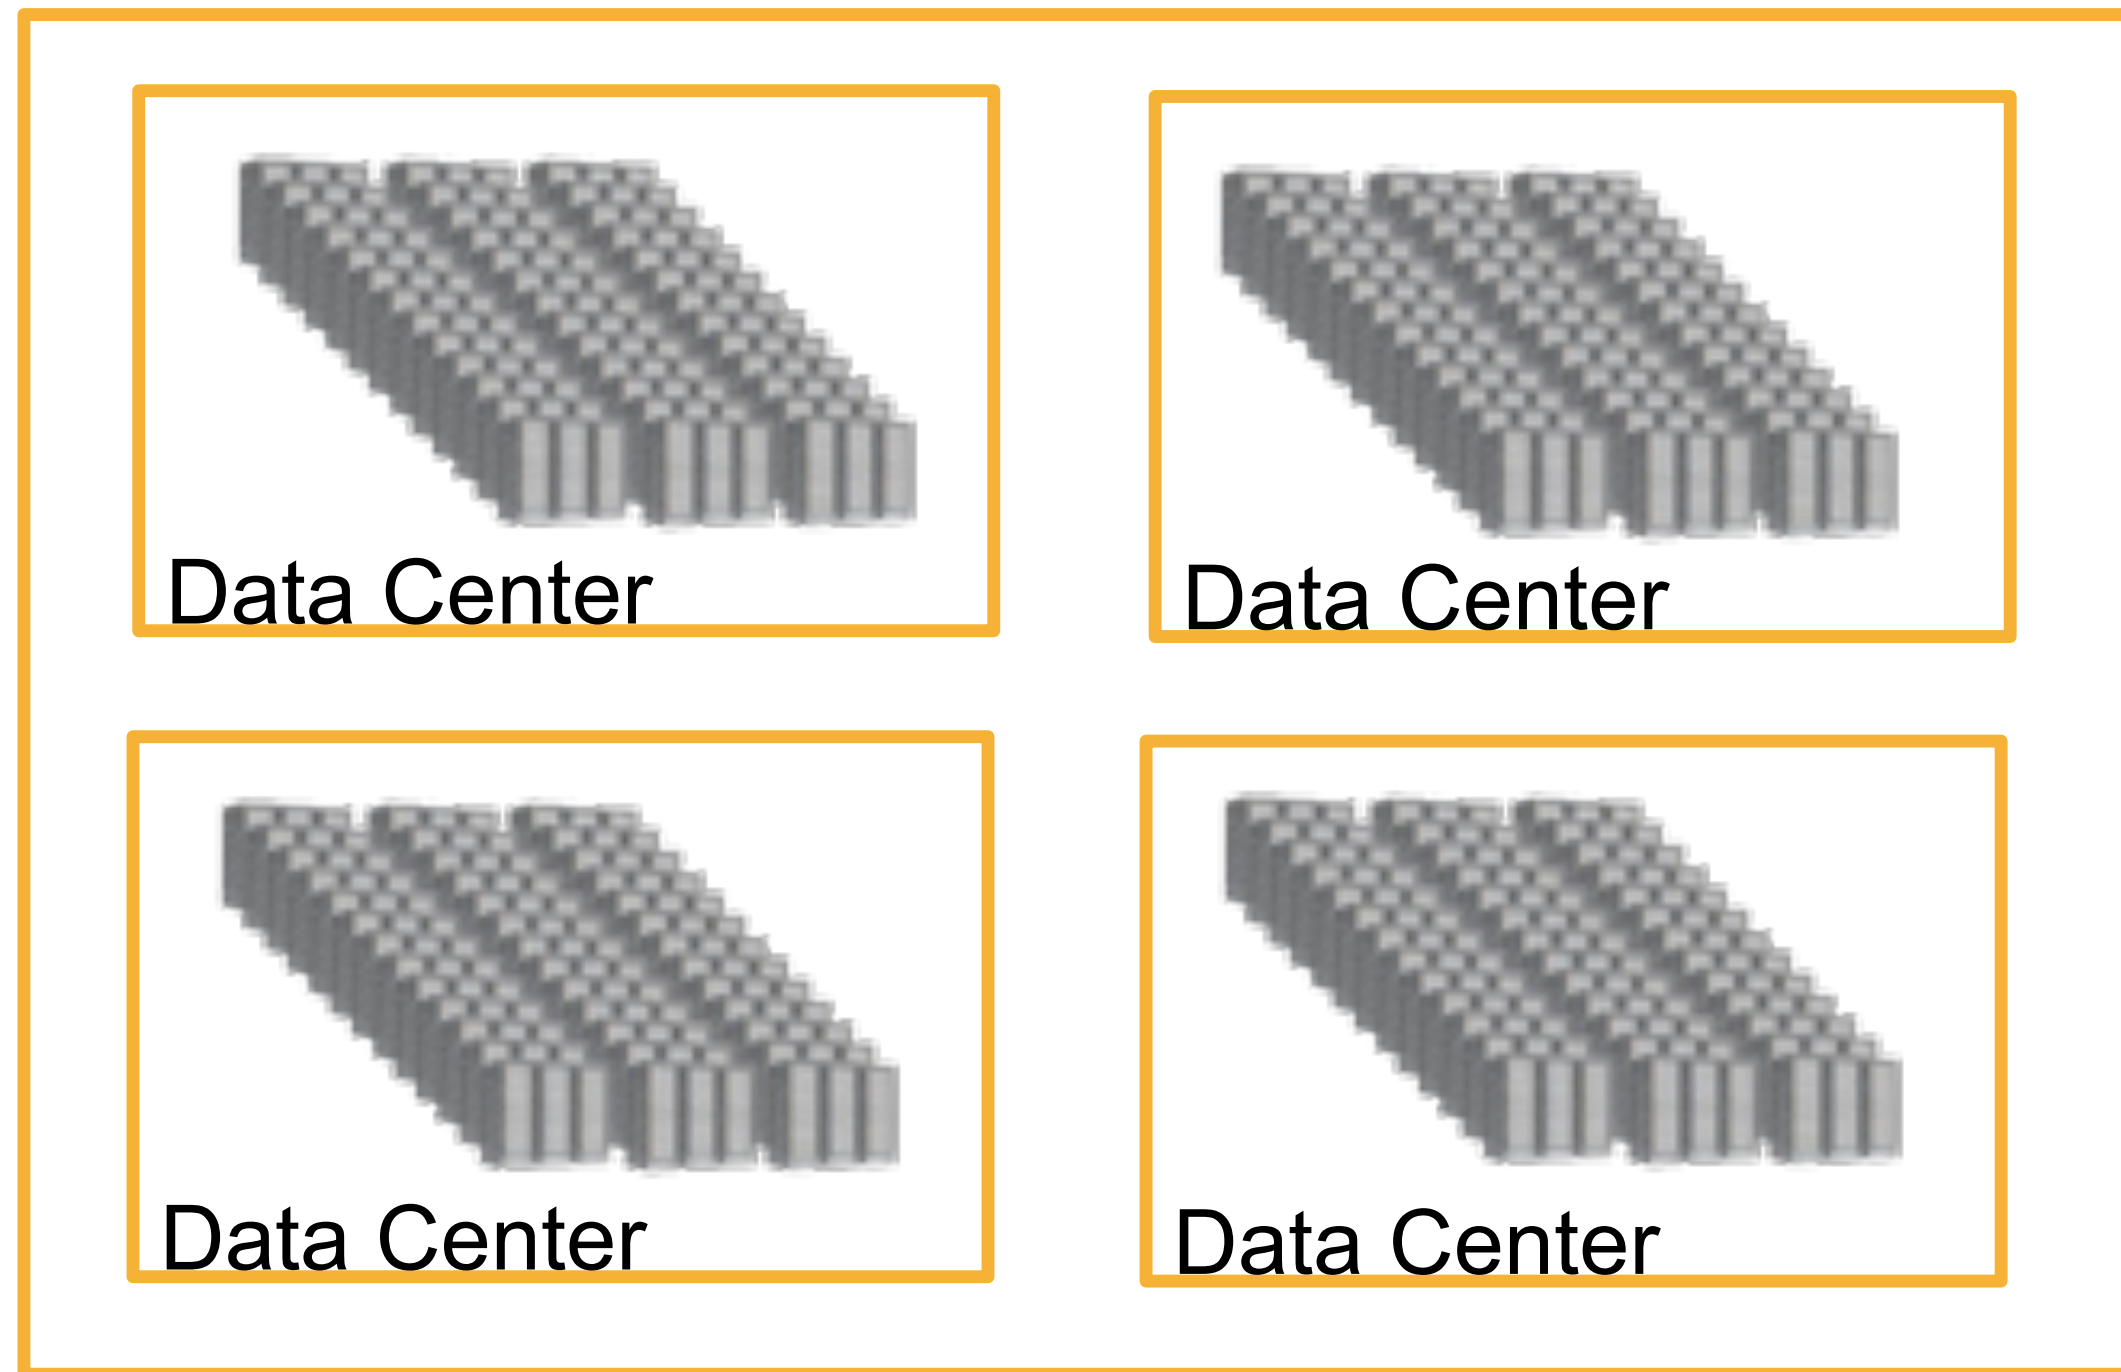
Every day, AWS adds enough  
server capacity to power this  
\$7B enterprise

Said another way...

AWS is deploying the equivalent of a top-20 supercomputer every few days


One server running for 10000 hours  
has the same cost as 10000 servers  
running for one hour.

# A Snapshot of AWS Global Capacity



- 53 AWS Edge Locations
- 11 Regions
- 28 Availability Zones
- 2 or more AZs per Region
- 1-6 Data Centers per AZ
- 50,000-80,000+ servers per DC
- Up to 102 Tbps provisioned to each DC

# EC2 Instance Type History: Increasing Customer Choice



**new**  
**existing**

**m1.small**

2006

**m1.xlarge**  
**m1.large**  
m1.small

2007

**c1.medium**  
**c1.xlarge**  
m1.xlarge  
m1.large  
m1.small

2008

**m2.2xlarge**  
**m2.4xlarge**  
c1.medium  
c1.xlarge  
m1.xlarge  
m1.large  
m1.small

2009

**cc1.4xlarge**  
**cg1.4xlarge**  
**t1.micro**  
**m2.xlarge**  
m2.2xlarge  
m2.4xlarge  
c1.medium  
c1.xlarge  
m1.xlarge  
m1.large  
m1.small

2010

**cc2.8xlarge**  
cc1.4xlarge  
cg1.4xlarge  
t1.micro  
m2.xlarge  
m2.2xlarge  
m2.4xlarge  
c1.medium  
c1.xlarge  
m1.xlarge  
m1.large  
m1.small

2011

**cr1.8xlarge**  
**hs1.8xlarge**  
**m3.xlarge**  
**m3.2xlarge**  
**hi1.4xlarge**  
**m1.medium**  
cc2.8xlarge  
cg1.4xlarge  
t1.micro  
m2.xlarge  
m2.2xlarge  
m2.4xlarge  
c1.medium  
c1.xlarge  
m1.xlarge  
m1.large  
m1.small

2012-2013


hs1.8xlarge  
m3.xlarge  
m3.2xlarge  
hi1.4xlarge  
m1.medium  
cc2.8xlarge  
cg1.4xlarge  
t1.micro  
m2.xlarge  
m2.2xlarge  
m2.4xlarge  
c1.medium  
c1.xlarge  
m1.xlarge  
m1.large  
m1.small

2014

**g2.2xlarge**  
**hs1.xlarge**  
**hs1.2xlarge**  
**hs1.4xlarge**  
c3.large  
c3.xlarge  
c3.2xlarge  
c3.4xlarge  
c3.8xlarge  
c4.large  
c4.xlarge  
c4.2xlarge  
c4.4xlarge  
c4.8xlarge  
m3.medium  
m3.large  
i2.large  
i2.xlarge  
i2.4xlarge  
i2.8xlarge  
r3.large  
r3.xlarge  
r3.2xlarge  
r3.4xlarge  
r3.8xlarge



# C4 Instance Family



Turbo to 3.5 Ghz on all cores

2.9GHz Haswell E5-2666 v3

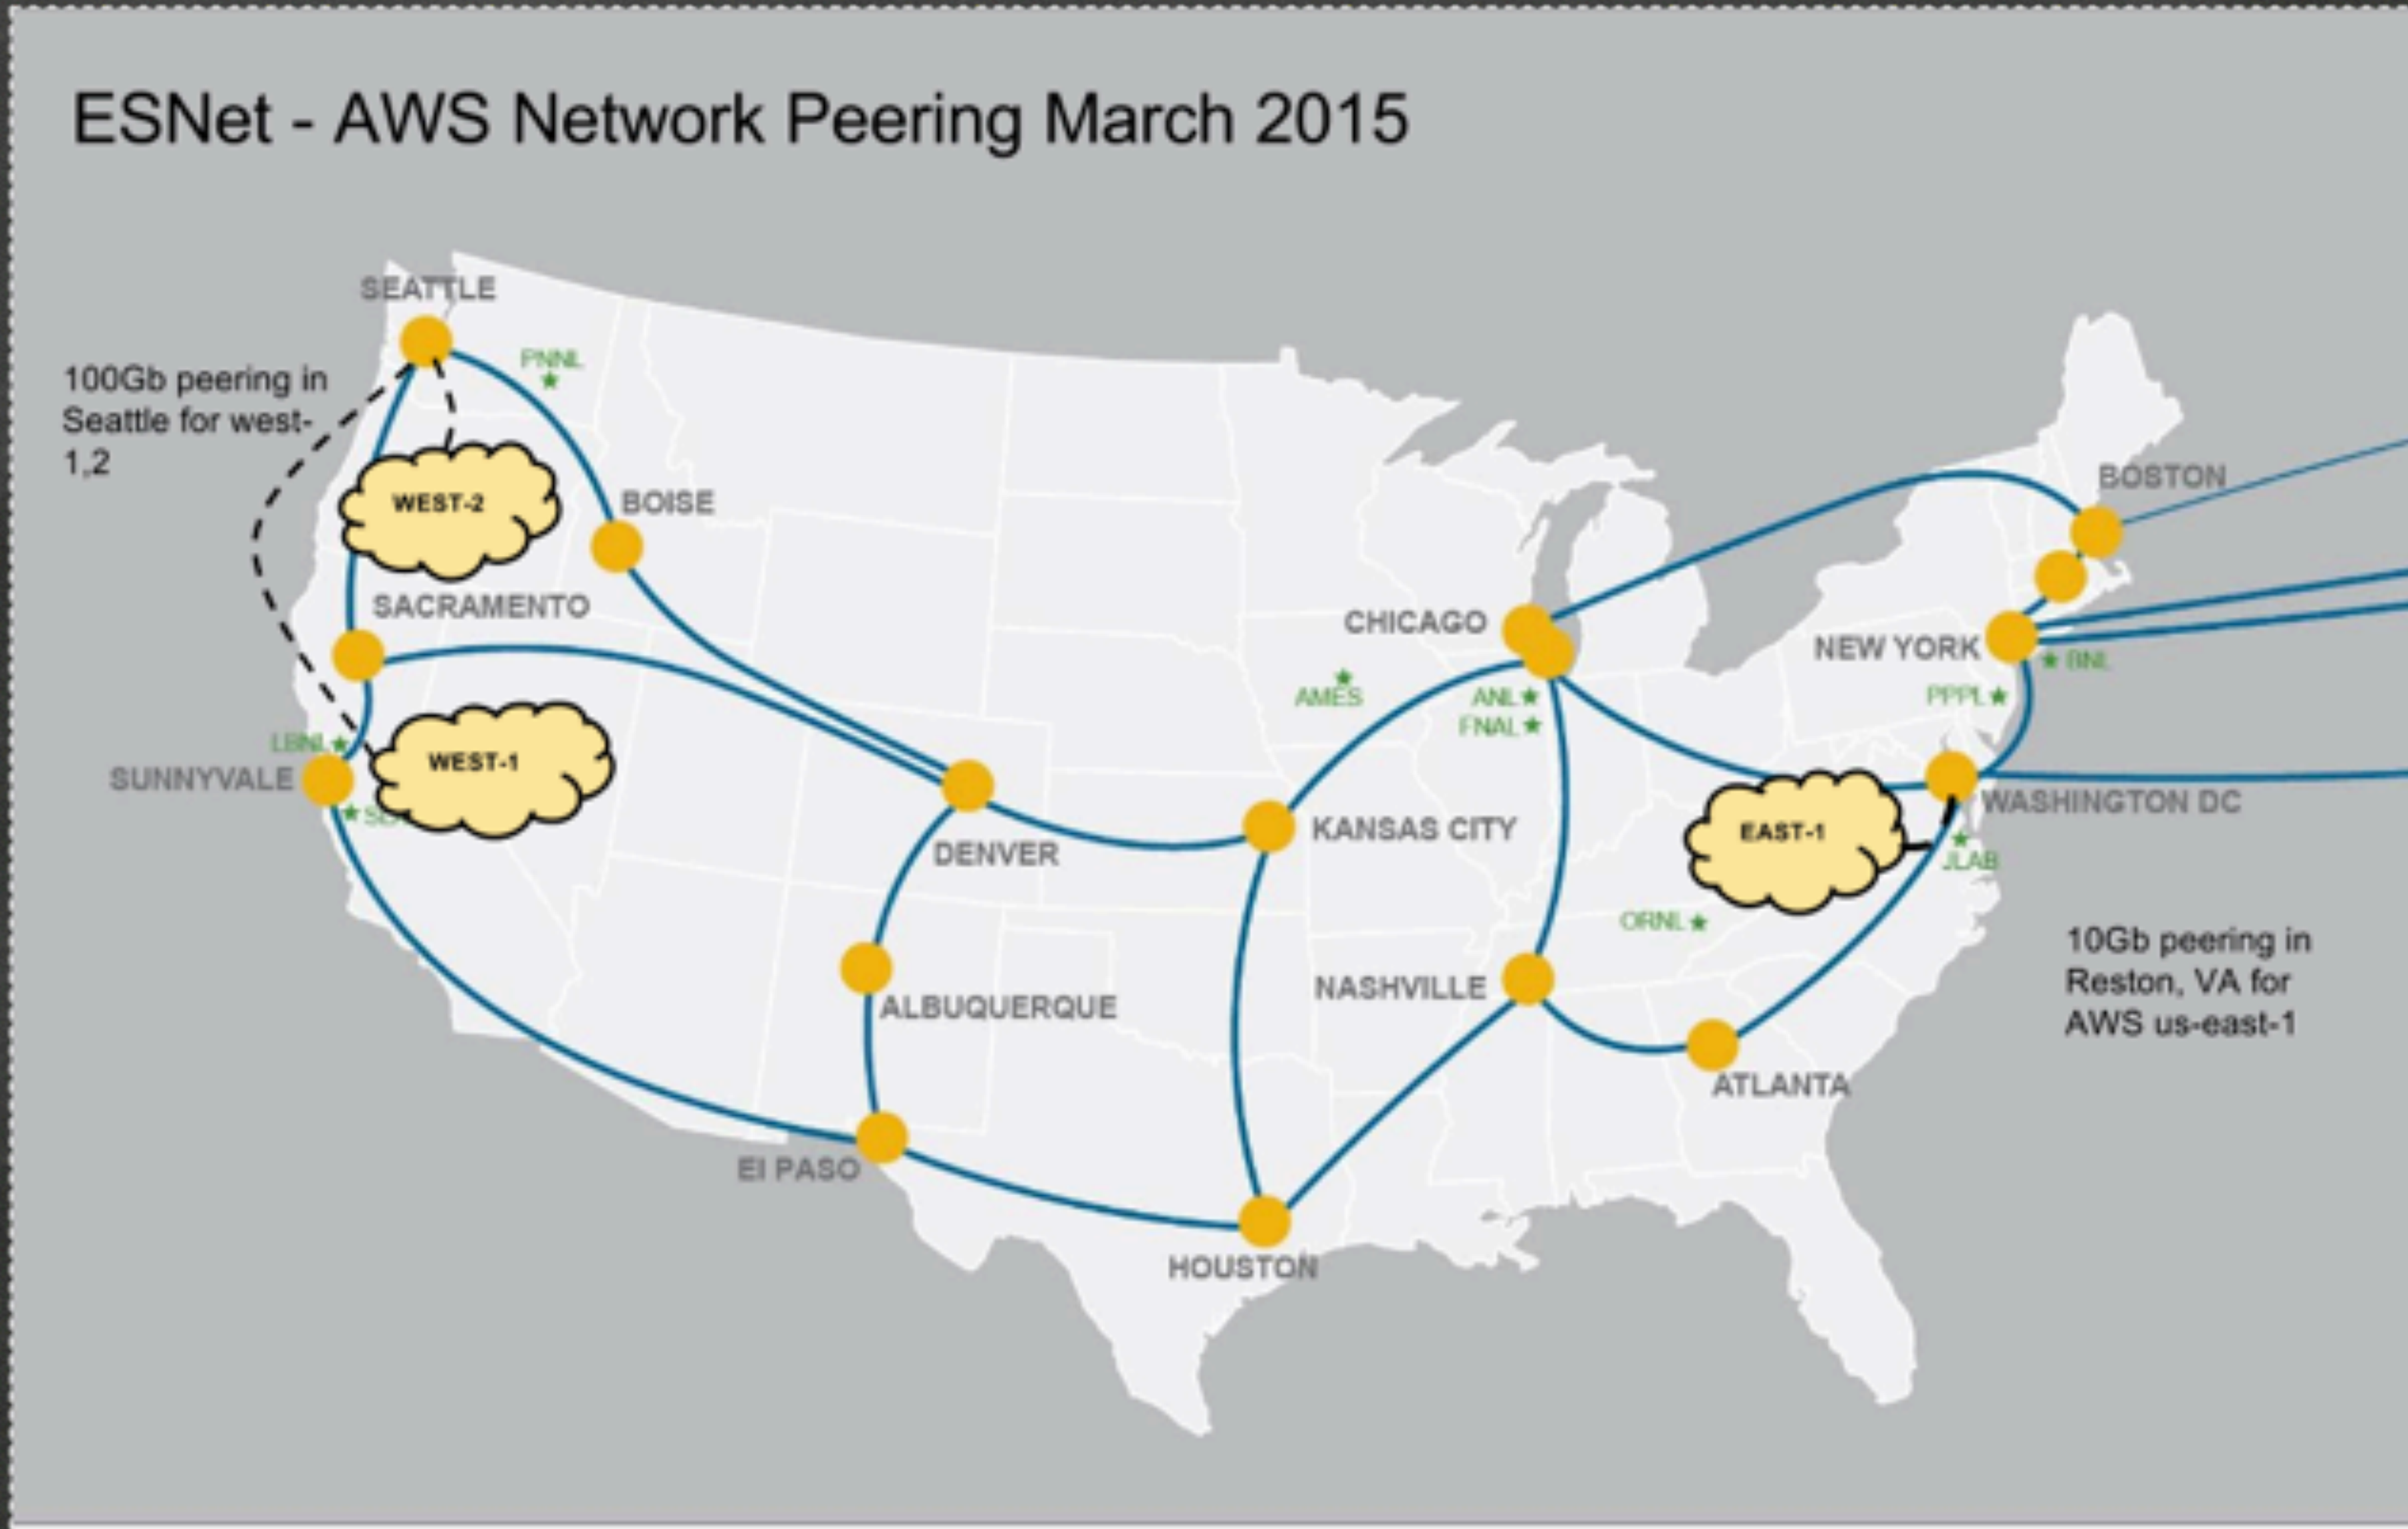
up to 60 gb RAM

Supports AVX2

<https://aws.amazon.com/blogs/aws/now-available-new-c4-instances/>

Instance Name	vCPU Count	RAM	Network Performance	Dedicated EBS Throughput	Linux On-Demand Price
c4.large	2	3.75 GiB	Moderate	500 Mbps	\$0.116/hour
c4.xlarge	4	7.5 GiB	Moderate	750 Mbps	\$0.232/hour
c4.2xlarge	8	15 GiB	High	1,000 Mbps	\$0.464/hour
c4.4xlarge	16	30 GiB	High	2,000 Mbps	\$0.928/hour
c4.8xlarge	36	60 GiB	10 Gbps	4,000 Mbps	\$1.856/hour

# AWS-ESNET Peering March 2015




# AWS Egress Waiver for Researchers


- 2013: Initial trial in Australia for users connecting via AARNET and AAPT
- 2014: Extended the waiver to include ESNET and Internet2
- 2015: Extending support to other major NRENs:
- Terms:
  - Waiving egress fees up to 15% of AWS bill, customers responsible for anything above this amount
  - Majority of traffic must transit via NREN with no transit costs
  - 15% waiver applies to aggregate usage for consolidated billing
  - Does not apply to workloads for which the egress is the service we provide...e.g. live video streaming, MOOCs, websites, etc...

# Scientific Computing Use Cases

- Science-as-a-Service
- Large-scale HTC (100,000+ core clusters)
- Large-scale MapReduce (Hadoop/Spark/Shark) using EC2 or EMR
- Small to medium-scale clusters (hundreds of nodes) for traditional MPI workloads
- Many small MPI clusters working in parallel to explore parameter space
- Small to medium scale GPGPU workloads
- Dev/test of MPI workloads prior to submitting to supercomputing centers
- Ephemeral clusters, custom tailored to the task at hand, created for various stages of a pipeline
- Collaborative research environments
- On-demand academic training/lab environments





On-demand access to effectively limitless resources



Specialized supercomputing resources




NASA Researcher

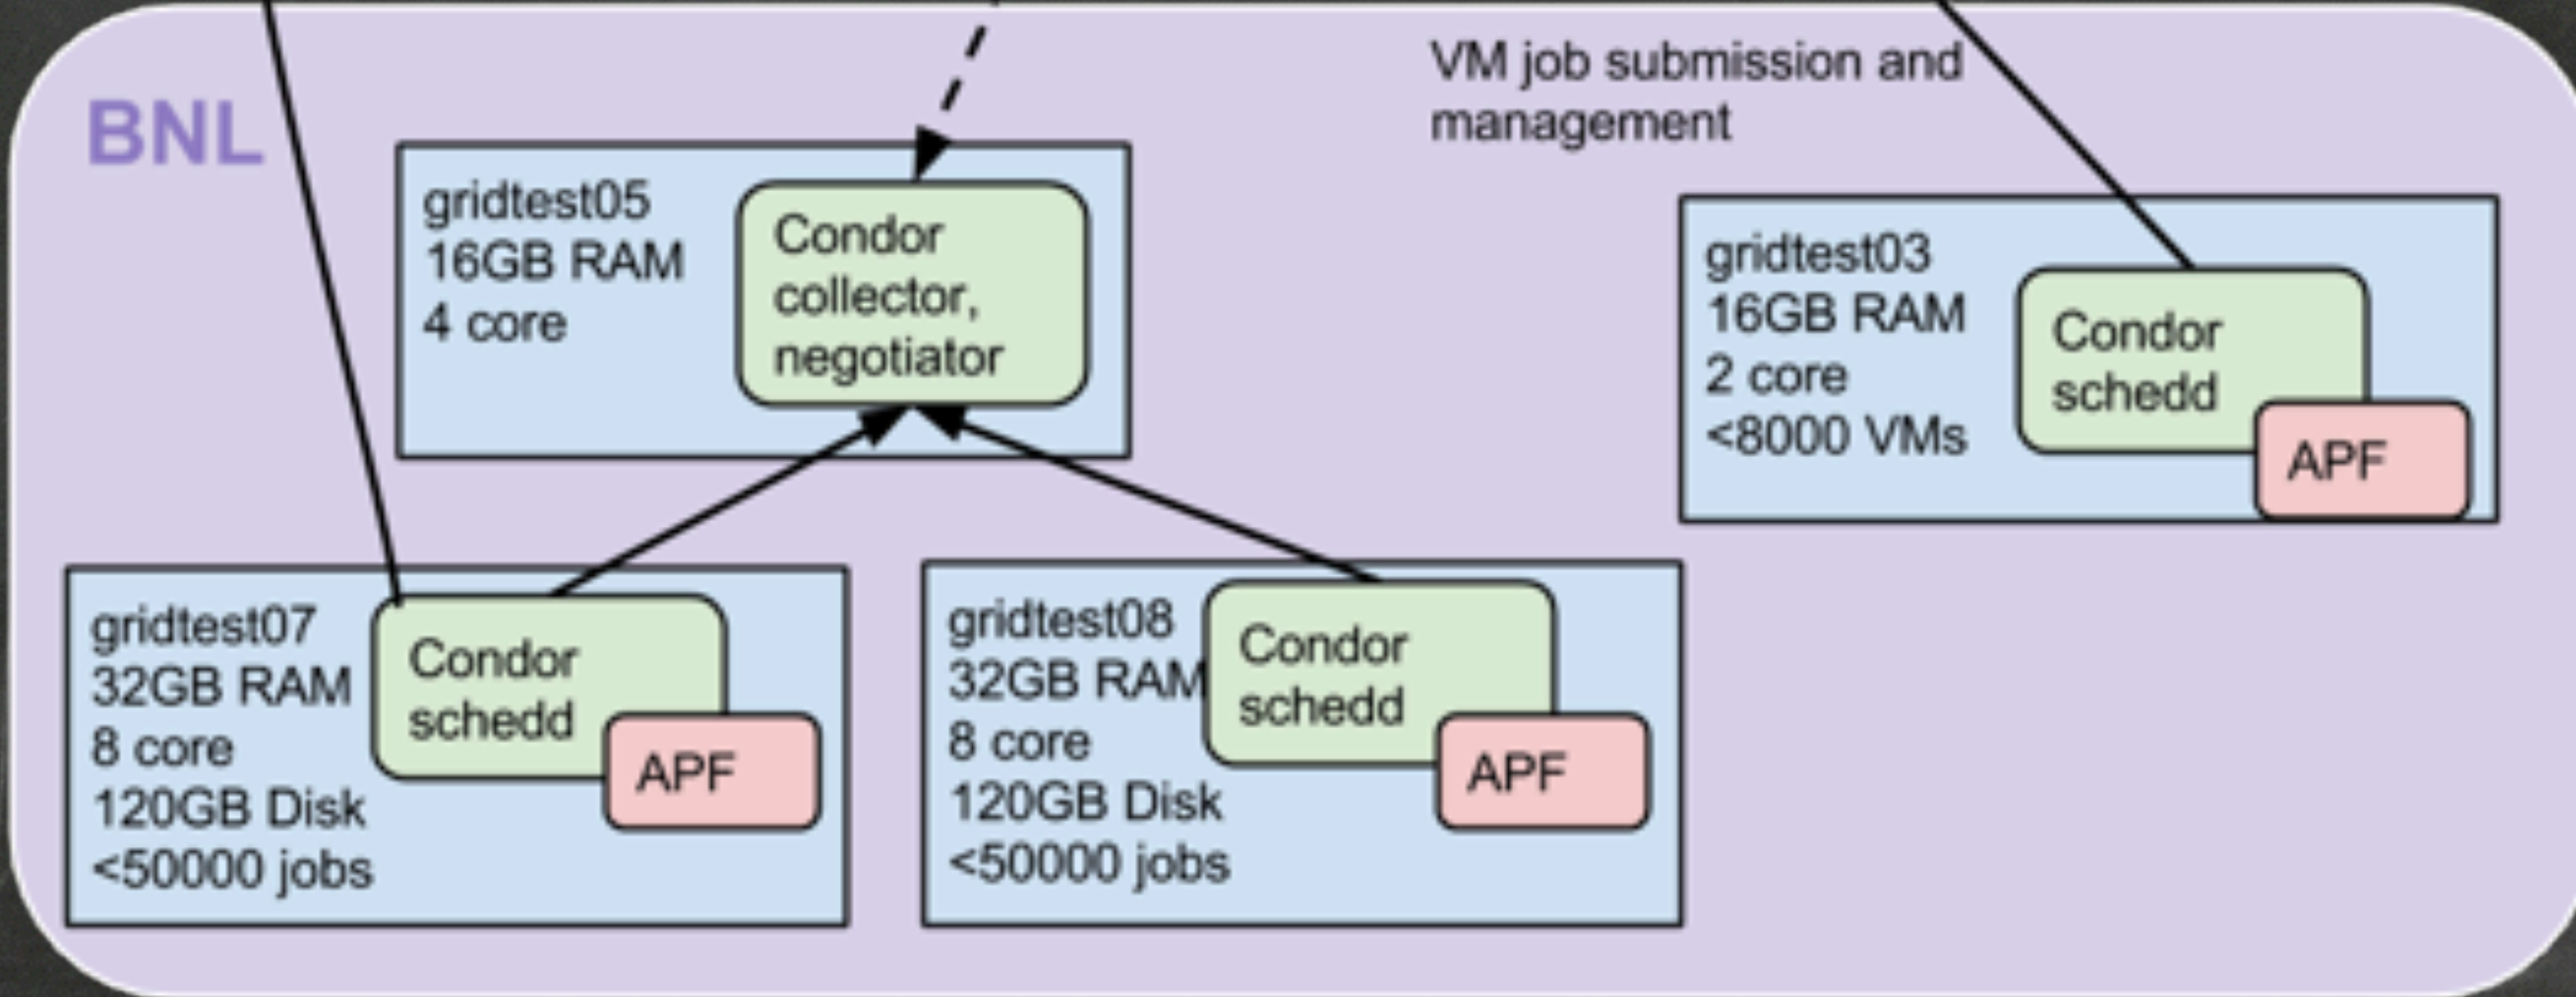



Small-scale shared compute

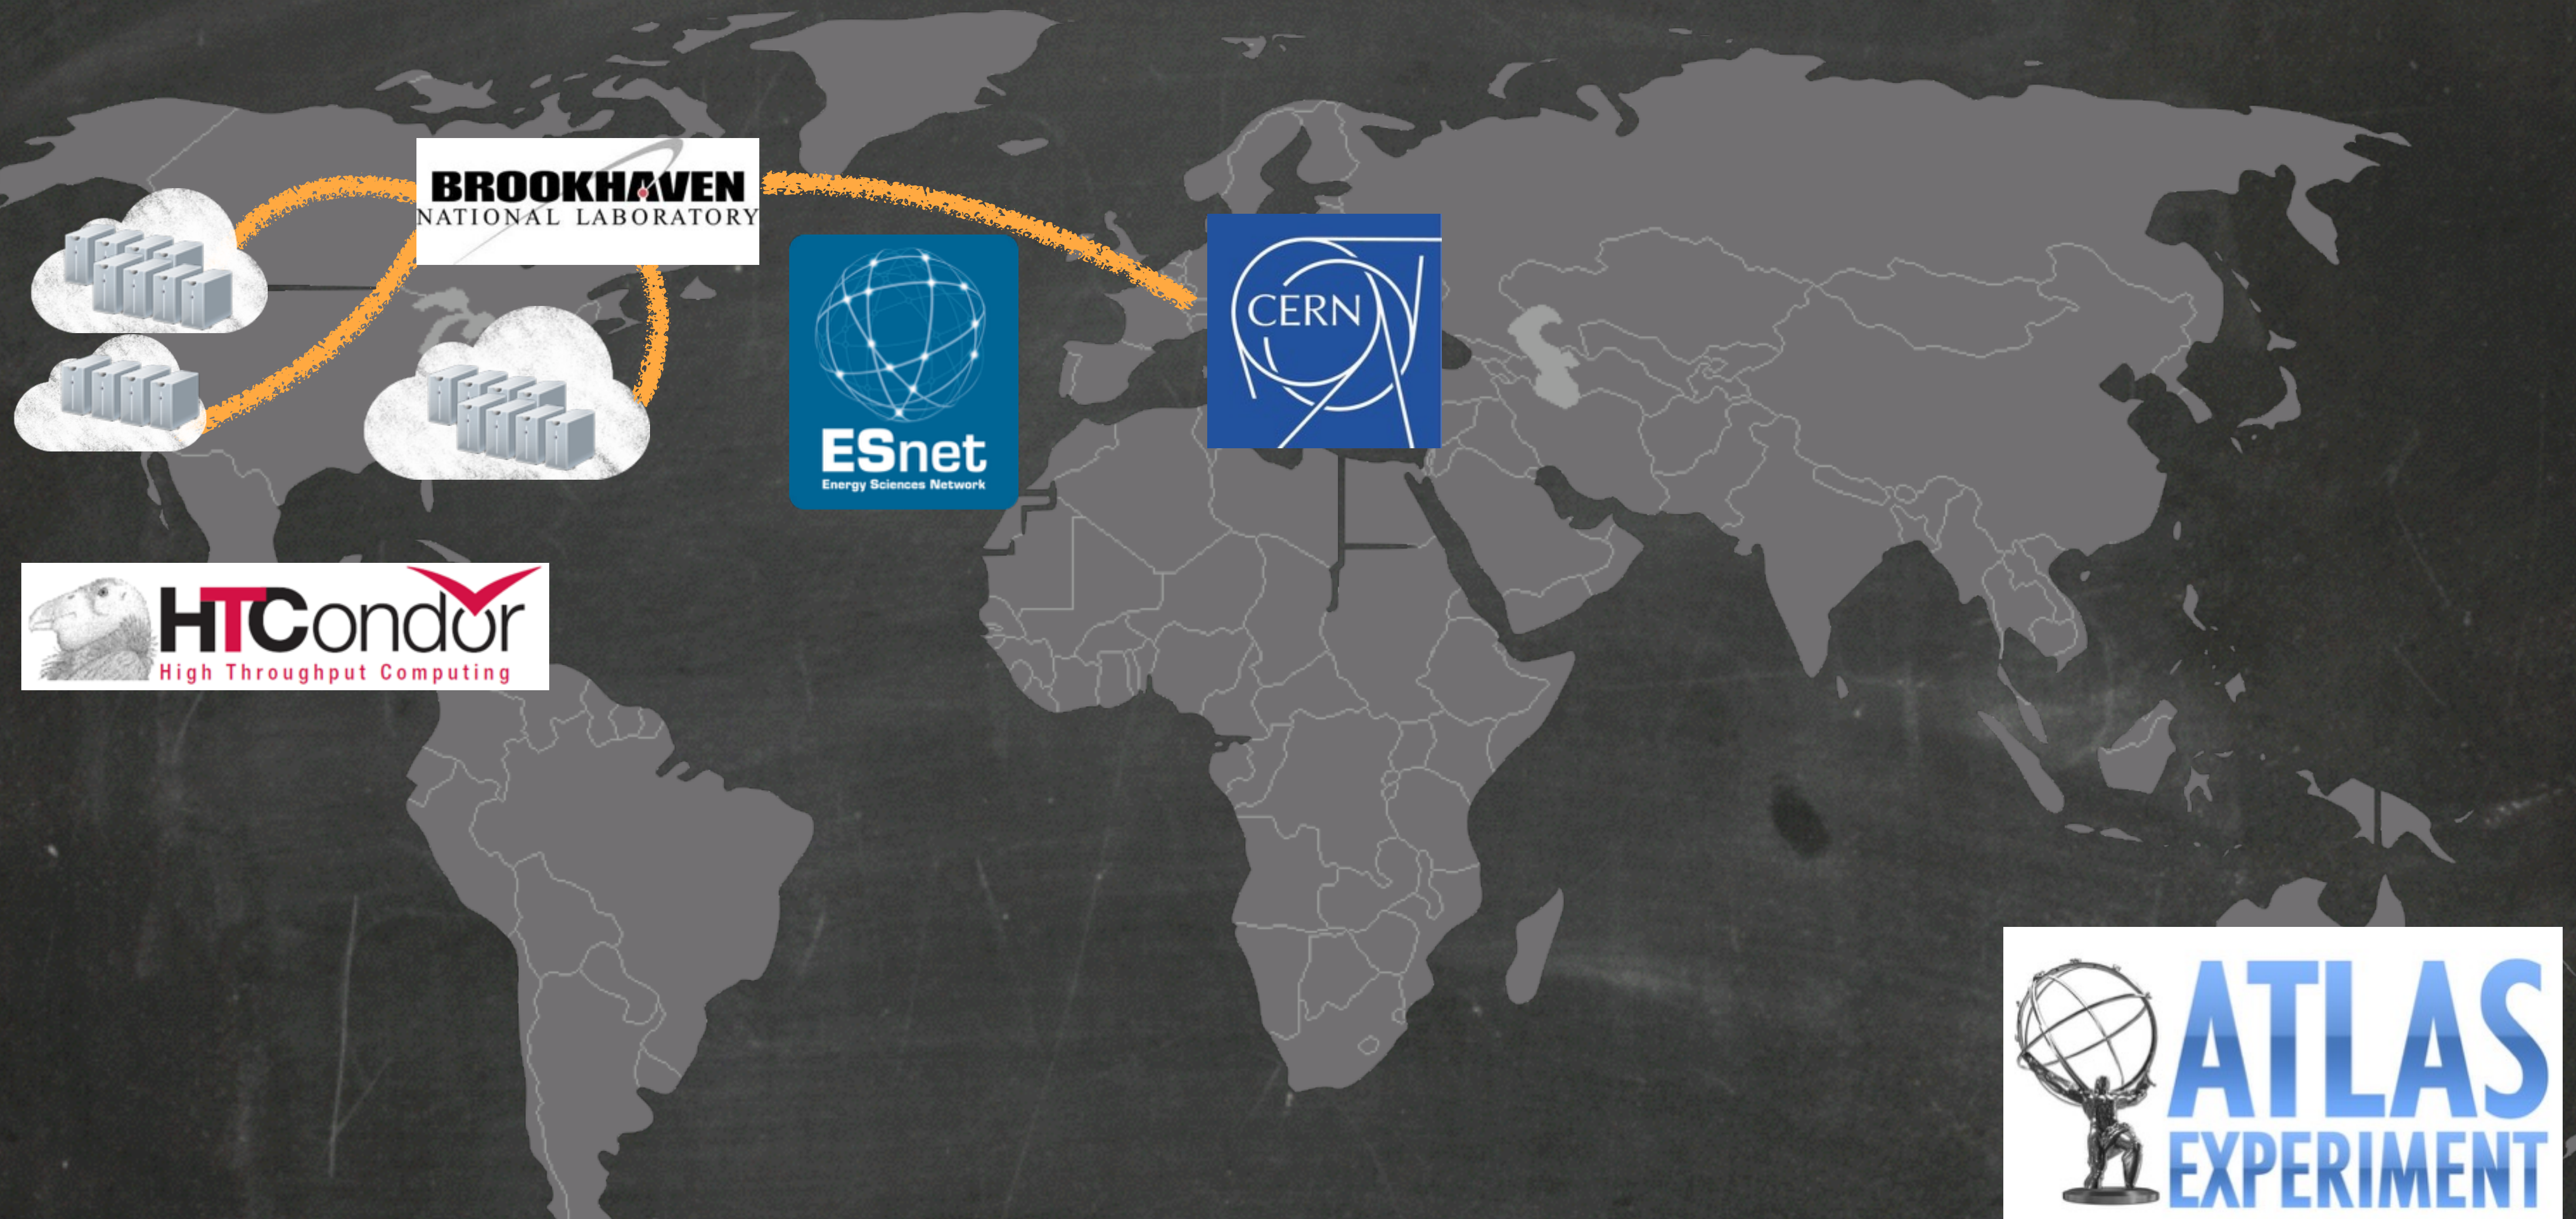


CMS Experiment at the LHC, CERN  
Data recorded: 2012-May-13 20:08:14.621490 GMT  
Run/Event: 194108 / 564224000






# Globally Distributed Compute for LHC US Atlas





# AWS Public Data Sets



[aws.amazon.com/datasets](https://aws.amazon.com/datasets)

# AWS Research Grants

**amazon web services** Sign Up My Account / Console English

AWS Products & Solutions AWS Product Information Developers Support

## AWS in Education

With AWS in Education, educators, academic researchers, and students can apply to obtain free usage credits to tap into the on-demand infrastructure of the Amazon Web Services cloud to teach advanced courses, tackle research endeavors, and explore new projects – tasks that previously would have required expensive up-front and ongoing investments in infrastructure.

Learn more about how Education programs use AWS grants:

- Berkeley RAD lab**  
Education Grant Case Study – Berkeley RAD Lab
- University of San Francisco**  
UNIVERSITY OF SAN FRANCISCO  
MS in Analytics Graduate Program
- University of California at Berkeley**  
Web 2.0 Application Development
- University of Texas at Austin**  
3 Day Startup Competition
- University of Oxford**  
UNIVERSITY OF OXFORD  
The Malaria Atlas Project

## Grants Application Process

Educators **Researchers** Students Education IT Machine Learning

### Researchers

AWS in Education will review and support selected research projects with grants that offer free access to most AWS infrastructure services. Often, large research projects require extensive compute power and storage infrastructure to complete. Now, researchers around the world have access to the global computing infrastructure and storage capacity of the AWS cloud. Instead of purchasing a large amount of hardware, researchers can get started by simply opening an AWS account. And, with services like Amazon Elastic

[AWS.amazon.com/grants](https://aws.amazon.com/grants)

Thank  
You

