Accelerating Time to Science: Transforming Research in the Cloud

Jamie Kinney - @jamiekinney
Director of Scientific Computing, a.k.a. “SciCo” – Amazon Web Services

amazon web services

Why Do Researchers Love AWS?

Time to Science
Access research infrastructure in minutes

Low Cost
Pay-as-you-go pricing

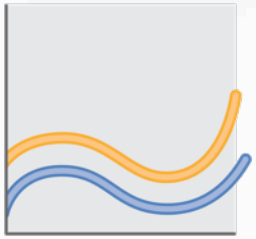

Elastic
Easily add or remove capacity

Globally Accessible
Easily Collaborate with researchers around the world

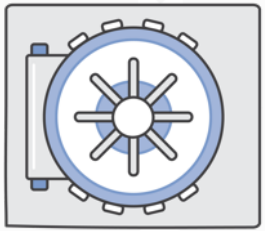

Secure
A collection of tools to protect data and privacy

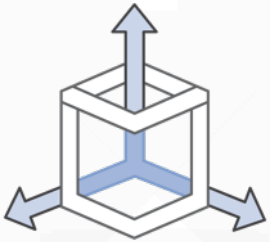

Scalable
Access to effectively limitless capacity

amazon web services

Why does Amazon care about Scientific Computing?

- In order to fundamentally accelerate the pace of scientific discovery
- It is a great application of AWS with a broad customer base
- The scientific community helps us innovate on behalf of all customers
  - Streaming data processing & analytics
  - Exabyte scale data management solutions and exaflop scale compute
  - 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 therefore reduce pricing for all AWS customers

How is AWS Used for Scientific Computing?

- High Throughput Computing (HTC) for Data-Intensive Analytics
- High Performance Computing (HPC) for Engineering and Simulation
- Collaborative Research Environments
- Hybrid Supercomputing Centers
- Science-as-a-Service
- Citizen Science

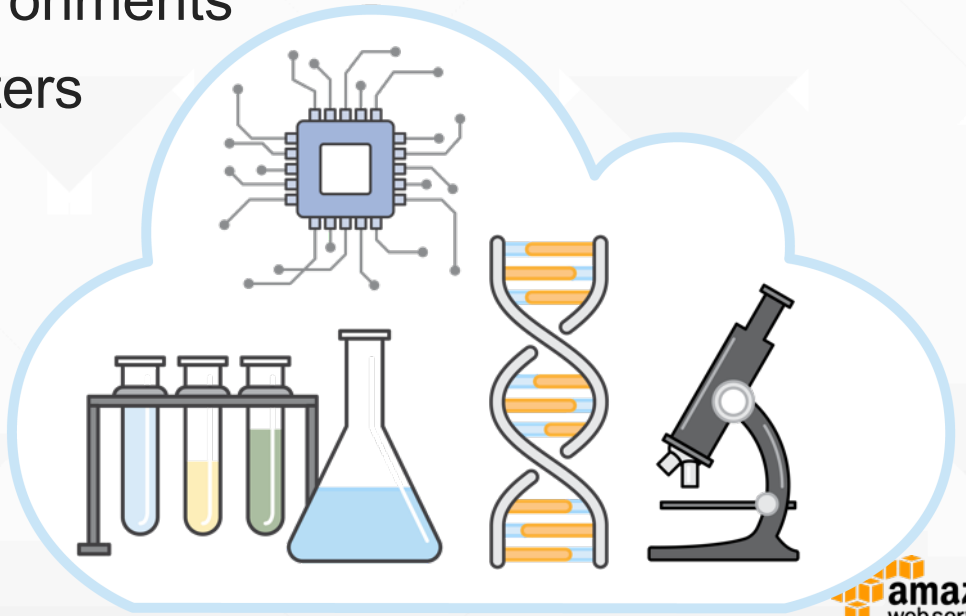

amazon web services

Research Grants

AWS provides free usage credits to help researchers:

- Teach advanced courses
- Explore new projects
- Create resources for the scientific community

aws.amazon.com/grants

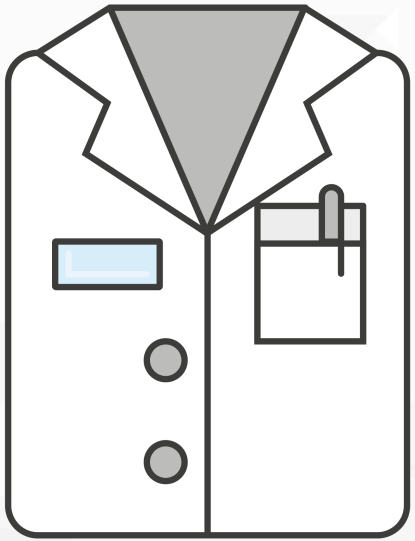

amazon web services

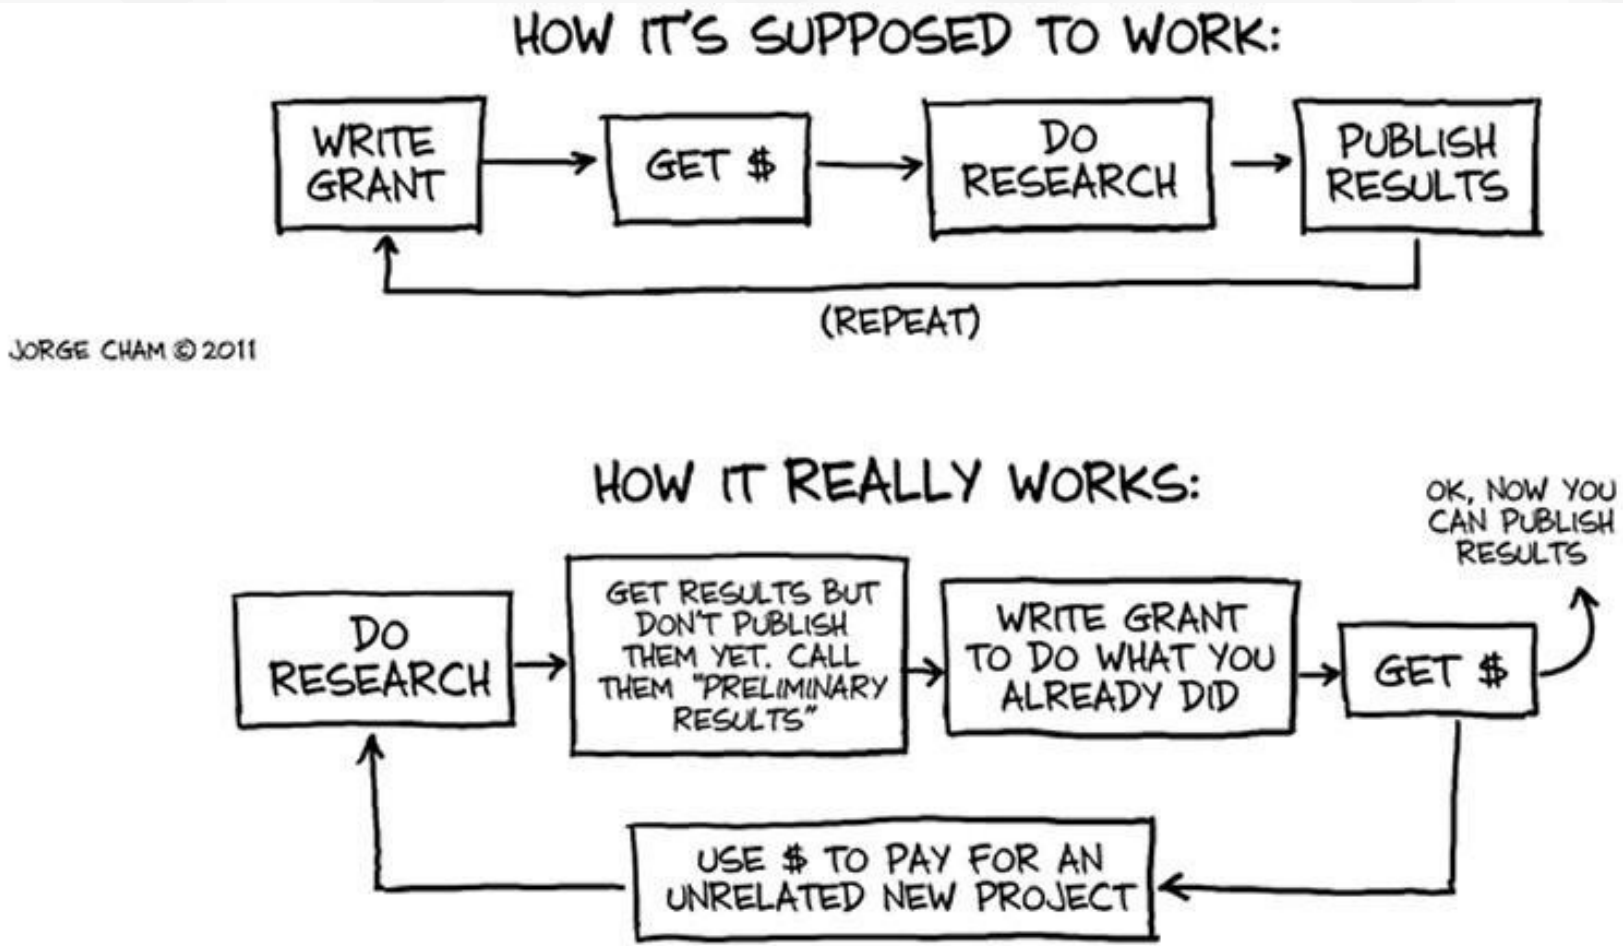

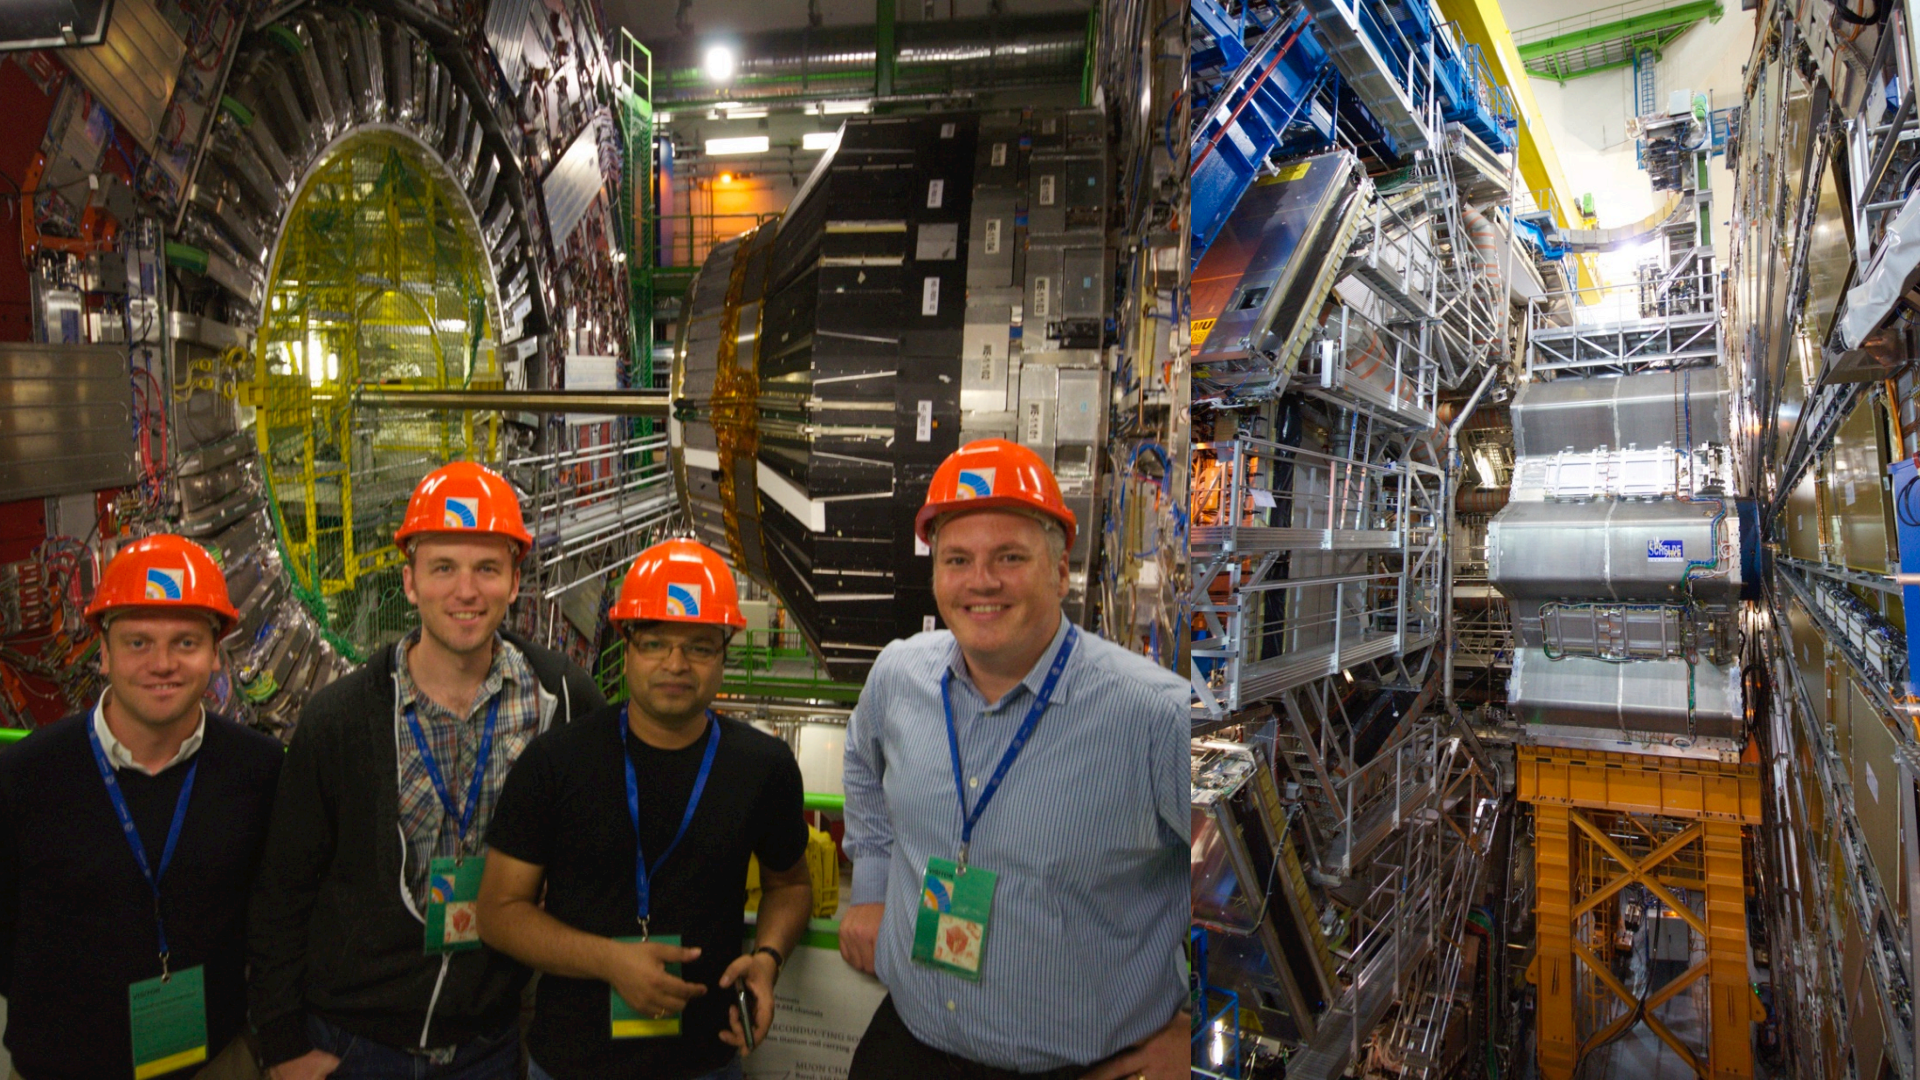

Amazon Public Data Sets

Public Data Sets

AWS hosts “gold standard” reference data at our expense in order to catalyze rapid innovation and increased AWS adoption

A few examples:

1,000 Genomes ~250 TB

Common Crawl

OpenStreetMap

Actively Developing…

Cancer Genomics Data Sets ~2-6 PB

SKA Precursor Data 1PB+

amazon web services

Nepal Earthquake

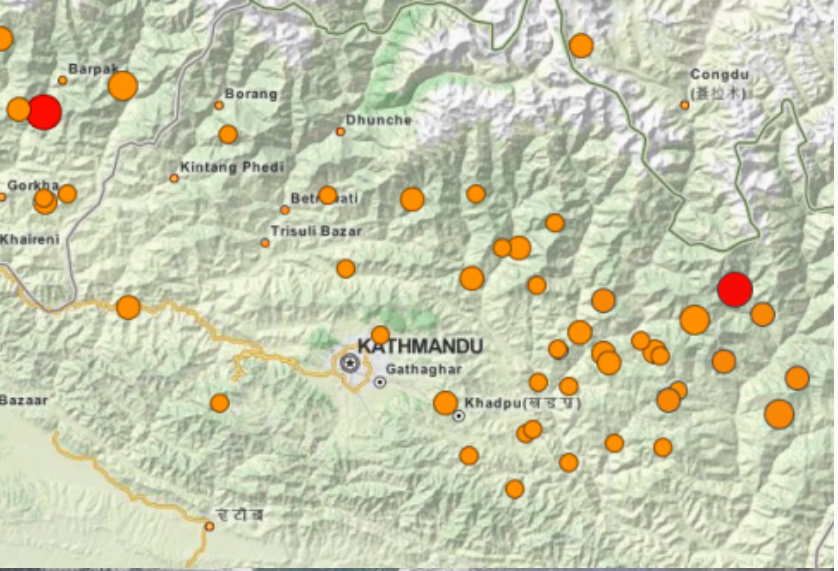

Individuals around the world are analyzing before/after imagery of Kathmandu in order to more-effectively direct emergency response and recovery efforts

HOT Humanitarian Open Street Map Team

Peering with all global research networks

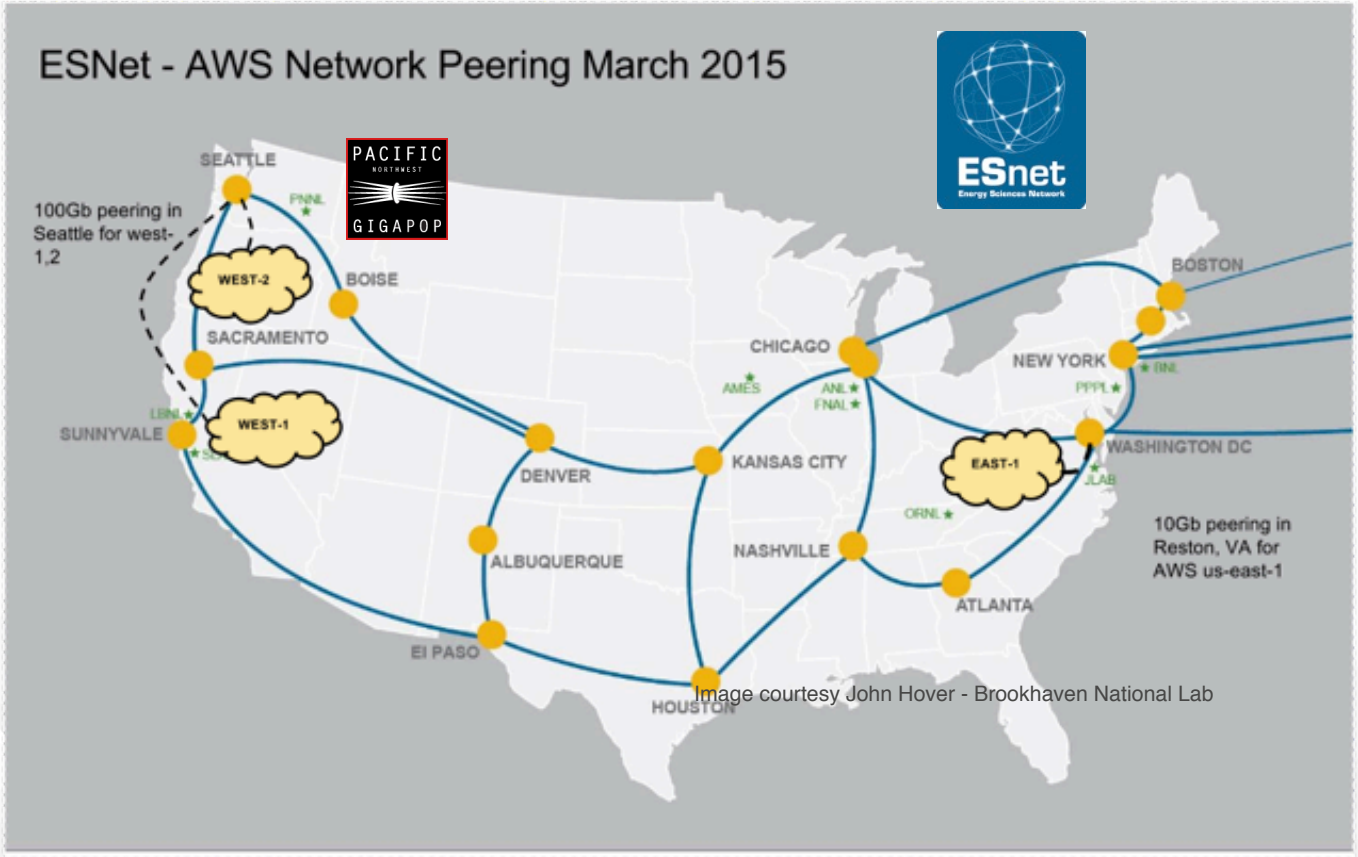

Image courtesy John Hover - Brookhaven National Lab

INTERNET 2

aarnet
Australia's Academic and Research Network

SINET4

GÉANT

amazon web services

AWS Egress Waiver for Research & Education

Timeline:

- 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:

- AWS waives egress fees up to 15% of total AWS bill, customers are responsible for anything above this amount
- Majority of traffic must transit via NREN with no transit costs
- 15% waiver applies to aggregate usage when consolidated billing is used
- Does not apply to workloads for which egress is the service we are providing (e.g. live video streaming, MOOCs, Web Hosting, etc…)
- Available regardless of AWS procurement method (i.e. direct purchase or Internet2 Net+)

Contact us if you would like to sign up!

amazon web services

Breaking news! Restricted-access genomics on AWS

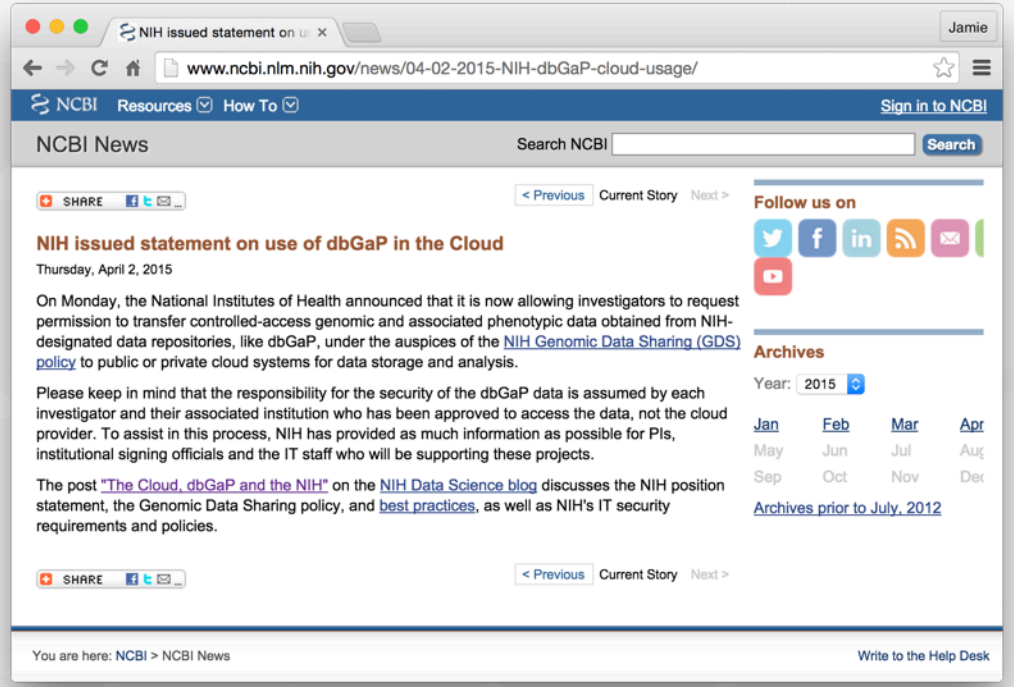

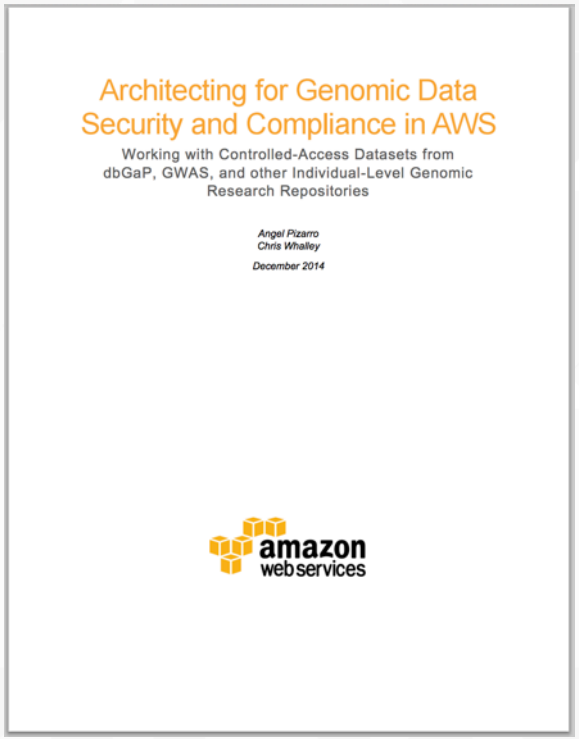

National Center for Biotechnology Information
NCBI

aws.amazon.com/genomics

amazon web services

Data-Intensive Computing

The Square Kilometer Array will link 250,000 radio telescopes together, creating the world's most sensitive telescope. The SKA will generate *zettabytes* of raw data, publishing exabytes annually *over 30-40 years.*

Researchers are using AWS to develop and test:

- Data processing pipelines
- Image visualization tools
- Exabyte-scale research data management
- Collaborative research environments

www.skatelescope.org/ska-aws-astrocompute-call-for-proposals/

amazon web services

Astrocompute in the Cloud Program

- AWS is adding 1PB of SKA pre-cursor data to the Amazon Public Data Sets program
- We are also providing $500K in AWS Research Grants for the SKA to direct towards projects focused on:
  - High-throughput data analysis
  - Image analysis algorithms
  - Data mining discoveries (i.e. ML, CV and data compression)
  - Exascale data management techniques
  - Collaborative research enablement

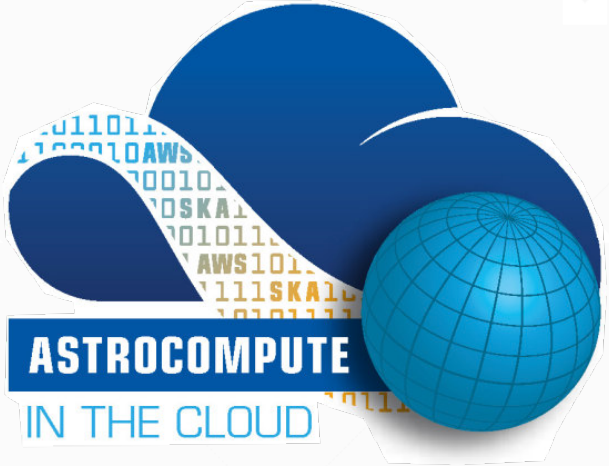

https://www.skatelescope.org/ska-aws-astrocompute-call-for-proposals/

Schrodinger & Cycle Computing: Computational Chemistry for Better Solar Power

Simulation by Mark Thompson of the University of Southern California to see which of 205,000 organic compounds could be used for photovoltaic cells for solar panel material.

Estimated computation time 264 years completed in 18 hours.

SCHRÖDINGER.

CYCLECOMPUTING

Loosely Coupled

- 156,314 core cluster, 8 regions
- 1.21 petaFLOPS (Rpeak)
- $33,000 or 16¢ per molecule

amazon web services

Some Core AWS Concepts

amazon web services

Region

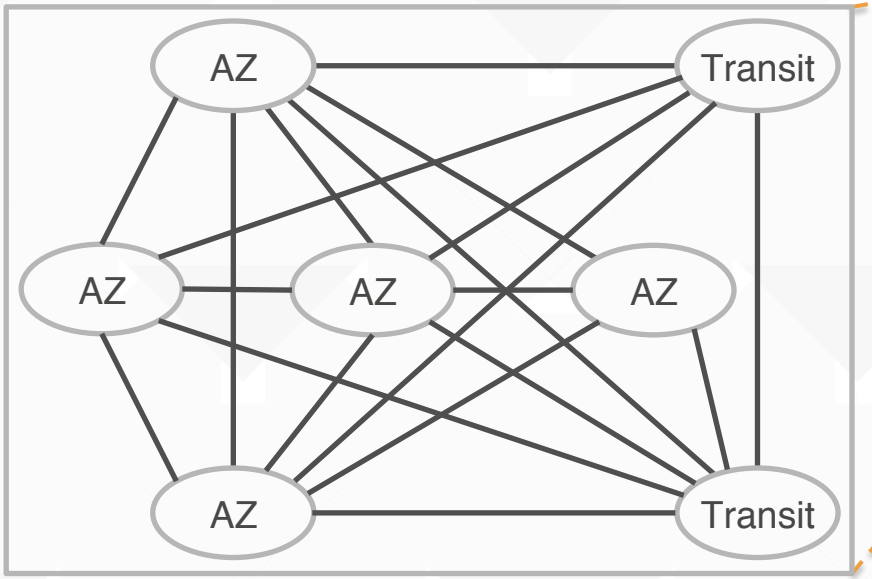

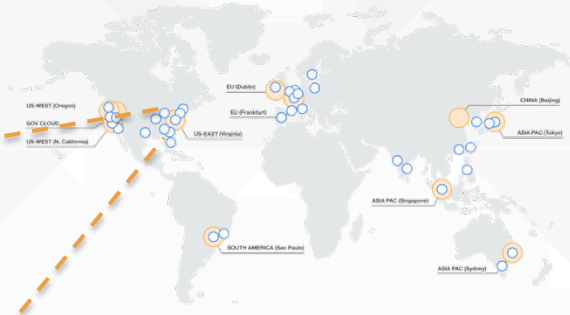

- Geographic area where AWS services are available
- Customers choose region(s) for their AWS resources
- Eleven regions worldwide

Availability Zone (AZ)

- Each region has multiple, isolated locations known as Availability Zones
- Low-latency links between AZs in a region <2ms, usually <1ms
- When launching an EC2 instance, a customer chooses an AZ
- Private AWS fiber links interconnect all major regions

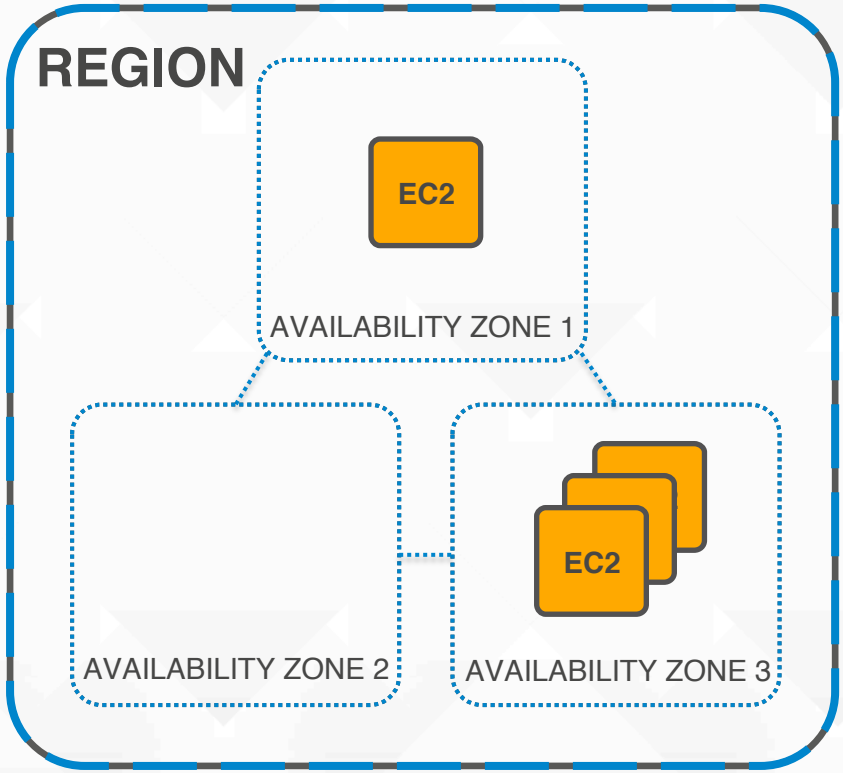

Example AWS Availability Zone

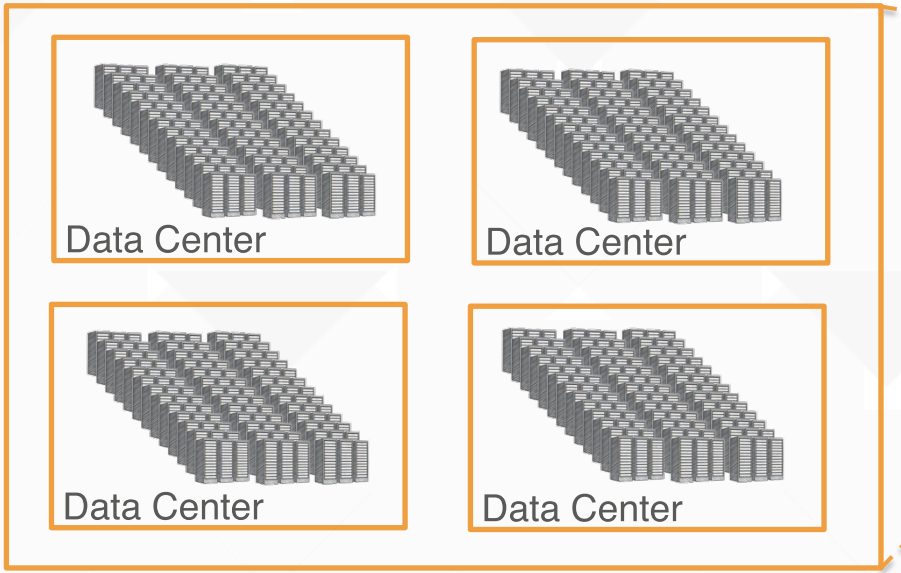

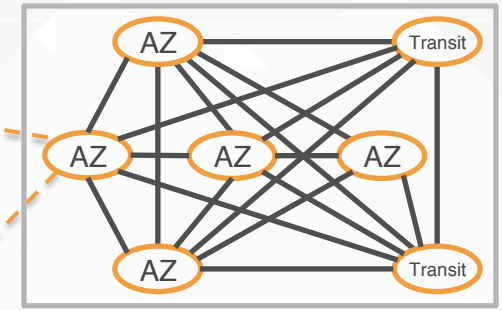

- 1 of 28 AZs world-wide
- All regions have 2 or more AZs
- Each AZ is 1 or more DC
  - No data center is in two AZs
  - Some AZs have as many as 6 DCs
- DCs in AZ less than ¼ ms apart

amazon web services

Example AWS Data Center

- Single DC typically over 50,000 servers & often over 80,000
  - Larger DCs undesirable (blast radius)
- Up to 102Tbps provisioned to a single DC
- AWS custom network equipment:
  - Multi-ODM sourced
  - Amazon custom network protocol stack

Virtual Private Cloud (VPC)

- Logically isolated section of the AWS cloud, virtual network defined by the customer
- When launching instances and other resources, customers place them in a VPC
- All new customers have a default VPC

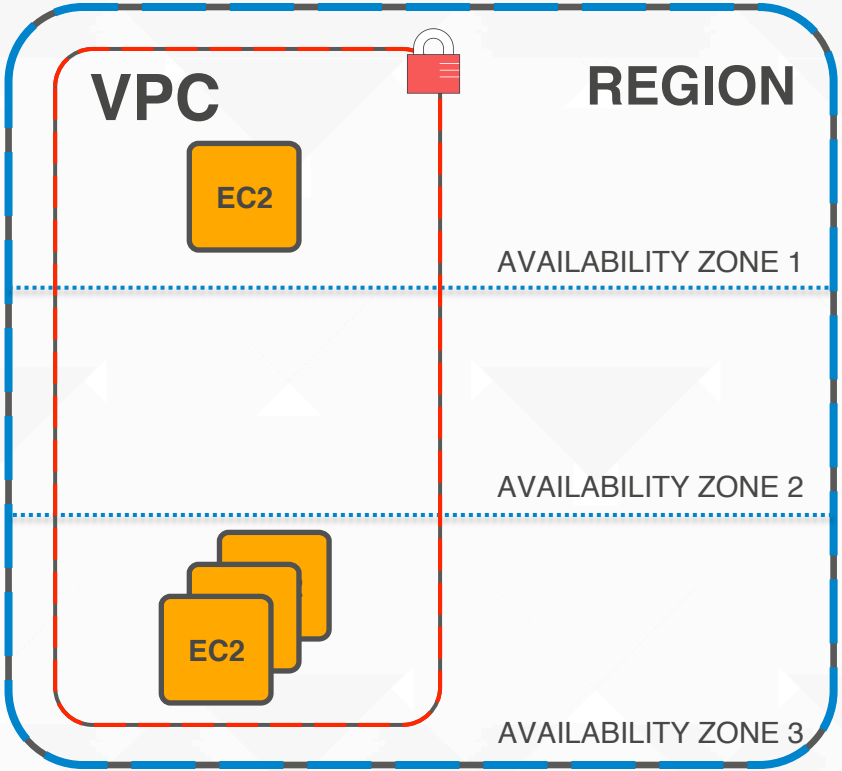

Spot Fleet

amazon web services

What is Spot?

- Name your own price for EC2 Compute
  - A market where price of compute changes based upon Supply and Demand
  - When Bid Price exceeds Spot Market Price, instance is launched
  - Instance is terminated (with 2 minute warning) if market price exceeds bid price
- Where does capacity come from?
  - Unused EC2 Instances

Product : Linux/UNIX (Amazon VPC) Instance type: m1.small Date range : 1 month Availability zone: All zones

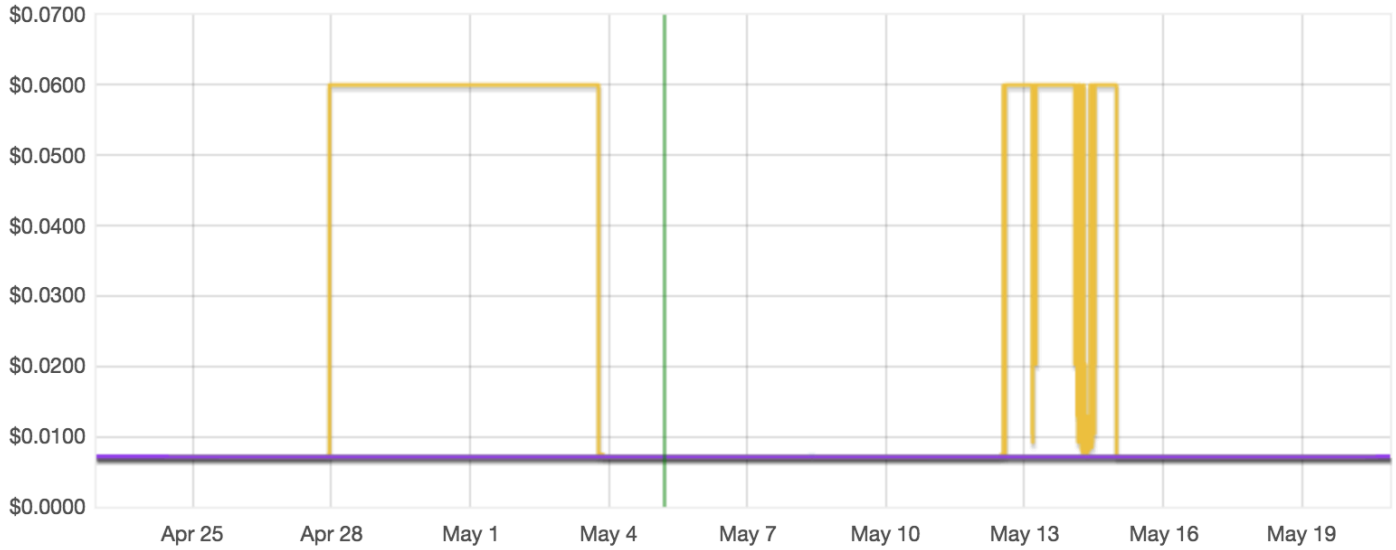

m1.small
On Demand
Price:
$0.044/hr

Availability zone	Price
us-east-1a	$0.0071
us-east-1b	$0.0071
us-east-1c	$0.0071
us-east-1d	$0.0071
us-east-1e	$0.0071
Date	May 4, 2015 at 11:58:59 PM UTC-5

amazon web services

Product : Linux/UNIX (Amazon VPC) Instance type: c3.8xlarge Date range : 1 month Availability zone: All zones

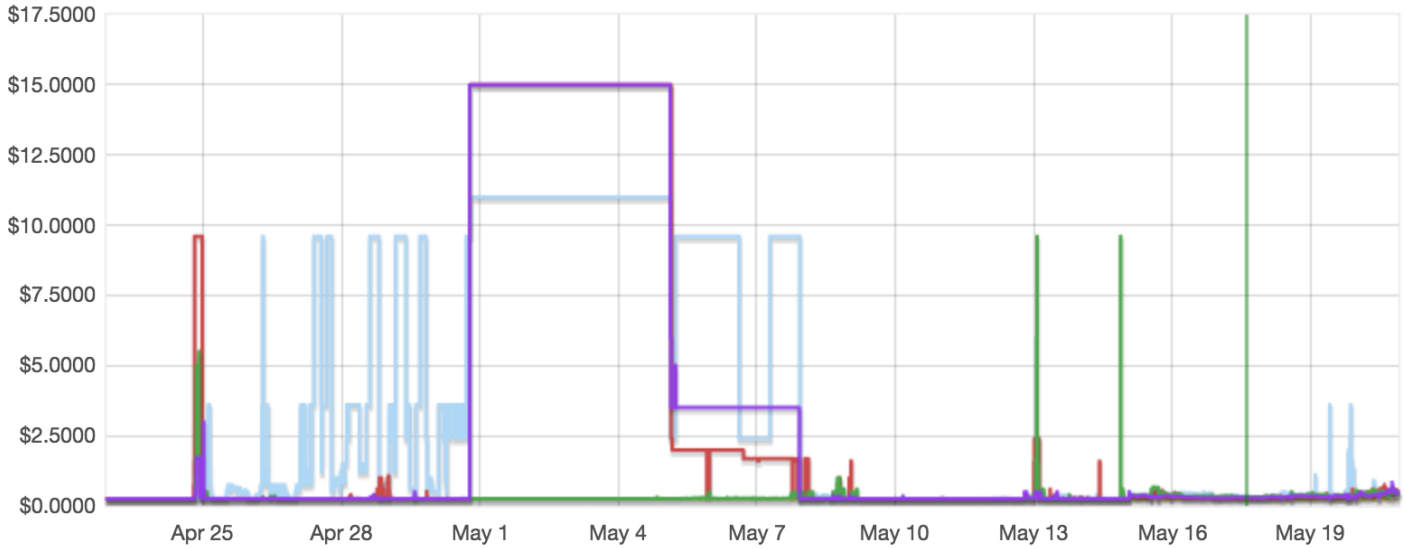

c3.8xlarge On Demand Price: $1.68/hr

Availability zone	Price
us-east-1a	$0.2560
us-east-1b	$0.3546
us-east-1c	$0.2814
us-east-1d	$0.3114
us-east-1e	$0.2903
Date	May 17, 2015 at 9:51:51 AM UTC-5

amazon web services

Product : Linux/UNIX (Amazon VPC) Instance type: cc2.8xlarge Date range : 1 month Availability zone: All zones

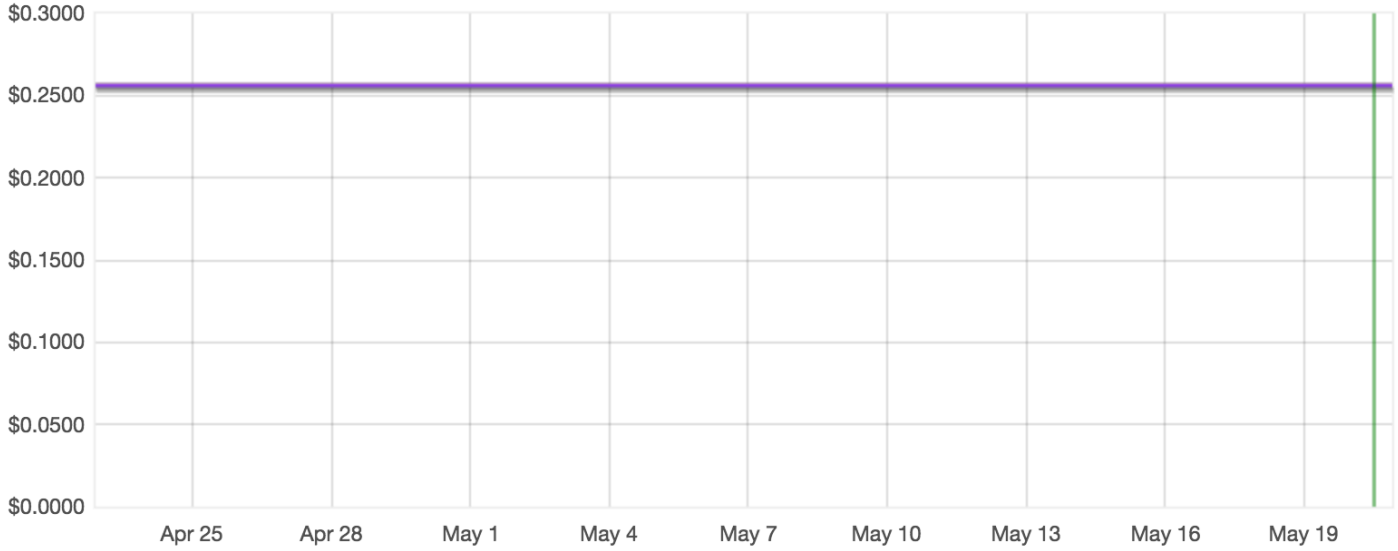

cc2.8xlarge On Demand Price: $2.00/hr

Availability zone	Price
us-east-1a	
us-east-1b	
us-east-1c	$0.2561
us-east-1d	$0.2561
us-east-1e	$0.2561
Date	May 20, 2015 at 7:23:55 AM UTC-5

amazon web services

1. zsh

```
Last login: Tue May 19 00:35:50 on ttys001
b8f6b1153503% aws ec2 describe-spot-price-history --instance-types m1.xlarge --product-description "Linux/UNIX" --start-time 2015-05-01T00:00:0 -
-end-time 2015-05-01T04:00:00 --max-items 20
SPOTPRICEHISTORY        us-east-1a      m1.xlarge       Linux/UNIX      0.480000        2015-05-01T03:23:51.000Z
SPOTPRICEHISTORY        us-east-1b      m1.xlarge       Linux/UNIX      0.032600        2015-05-01T03:23:51.000Z
SPOTPRICEHISTORY        us-east-1d      m1.xlarge       Linux/UNIX      0.032100        2015-05-01T03:23:51.000Z
SPOTPRICEHISTORY        us-east-1a      m1.xlarge       Linux/UNIX      1.500000        2015-05-01T03:17:05.000Z
SPOTPRICEHISTORY        us-east-1b      m1.xlarge       Linux/UNIX      0.175000        2015-05-01T03:17:05.000Z
SPOTPRICEHISTORY        us-east-1d      m1.xlarge       Linux/UNIX      0.040000        2015-05-01T03:17:05.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.680000        2015-05-01T03:17:05.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.033800        2015-05-01T03:10:28.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.033600        2015-05-01T03:03:41.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.033700        2015-05-01T02:57:10.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.033400        2015-05-01T02:50:46.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.033700        2015-05-01T02:44:31.000Z
SPOTPRICEHISTORY        us-east-1b      m1.xlarge       Linux/UNIX      0.032600        2015-05-01T02:44:31.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.039300        2015-05-01T02:37:55.000Z
SPOTPRICEHISTORY        us-east-1b      m1.xlarge       Linux/UNIX      0.032500        2015-05-01T02:37:55.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.350000        2015-05-01T02:31:21.000Z
SPOTPRICEHISTORY        us-east-1a      m1.xlarge       Linux/UNIX      0.480000        2015-05-01T01:31:26.000Z
SPOTPRICEHISTORY        us-east-1b      m1.xlarge       Linux/UNIX      0.032600        2015-05-01T01:31:26.000Z
SPOTPRICEHISTORY        us-east-1d      m1.xlarge       Linux/UNIX      0.032100        2015-05-01T01:31:26.000Z
SPOTPRICEHISTORY        us-east-1c      m1.xlarge       Linux/UNIX      0.480000        2015-05-01T01:31:26.000Z
NEXTTOKEN       None___20
b8f6b1153503%
```

Spot allows customers to run workloads that they would likely not run anywhere else..

But today, to be successful in Spot requires a little bit of additional effort

The Spot Experience today

- Build stateless, distributed, scalable applications
- Choose which instance types fit your workload the best
- Ingest price feed data for AZs and regions
- Make run time decisions on which Spot pools to launch in based on price and volatility
- Manage interruptions
- Monitor and manage market prices across Azs and instance types
- Manage the capacity footprint in the fleet
- And all of this while you don't know where the capacity is
- Serve your customers

UNDIFFERENTIATED HEAVY LIFTING

Making Spot Fleet Requests

- Simply specify:
  - **Target Capacity** – The number of EC2 instances that you want in your fleet.
  - **Maximum Bid Price** – The maximum bid price that you are willing to pay.
  - **Launch Specifications** – # of and types of instances, AMI id, VPC, subnets or AZs, etc.
  - **IAM Fleet Role** – The name of an IAM role. It must allow EC2 to terminate instances on your behalf.

Spot Fleet

- Will attempt to reach the desired target capacity given the choices that were given
- Manage the capacity even as Spot prices change
- Launch using launch specifications provided

Using Spot Fleet

- Create EC2 Spot Fleet IAM Role
- Requesting a fleet:
  - aws ec2 request-spot-fleet --spot-fleet-request-config file://mySmallFleet.json
- Describe fleet:
  - aws ec2 describe-spot-fleet-requests
  - aws ec2 describe-spot-fleet-requests --spot-fleet-request-ids <sfr-……..>
- Describe instances within the fleet
  - aws ec2 describe-spot-fleet-instances --spot-fleet-request-id <sfr-……>
- Cancel Spot Fleet (with termination):
  - aws ec2 cancel-spot-fleet-requests --spot-fleet-request-ids <sfr-……..> -terminate-instances

http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/spot-fleet.html

mySpotFleet.json

```
{
    "SpotPrice": "0.50",
    "TargetCapacity": 20,
    "IamFleetRole": "arn:aws:iam::123456789012:role/my-spot-fleet-role",
    "LaunchSpecifications": [
        {
            "ImageId": "ami-1a2b3c4d",
            "InstanceType": "cc2.8xlarge",
            "SubnetId": "subnet-a61dafcf"
        },
        {
            "ImageId": "ami-1a2b3c4d",
            "InstanceType": "r3.8xlarge",
            "SubnetId": "subnet-a61dafcf"
        }
    ]
}
```

Elastic File System

amazon web services

The AWS storage portfolio

Amazon S3

- Object storage: data presented as buckets of objects
- Data access via APIs over the Internet

Amazon Elastic Block Store

- Block storage (analogous to SAN): data presented as disk volumes
- Lowest-latency access from single Amazon EC2 instances

Amazon Glacier

- Archival storage: data presented as vaults/archives of objects
- Lowest-cost storage, infrequent access via APIs over the Internet

Amazon EFS

- File storage (analogous to NAS): data presented as a file system
- Shared low-latency access from multiple EC2 instances

Amazon Elastic File System

- Fully managed file system for EC2 instances
- Provides standard file system semantics
- Works with standard operating system APIs
- Sharable across thousands of instances
- Elastically grows to petabyte scale
- Delivers performance for a wide variety of workloads
- Highly available and durable
- NFS v4–based

EFS is designed for a broad range of use cases, such as…

- Content repositories
- Development environments
- Home directories
- Big data

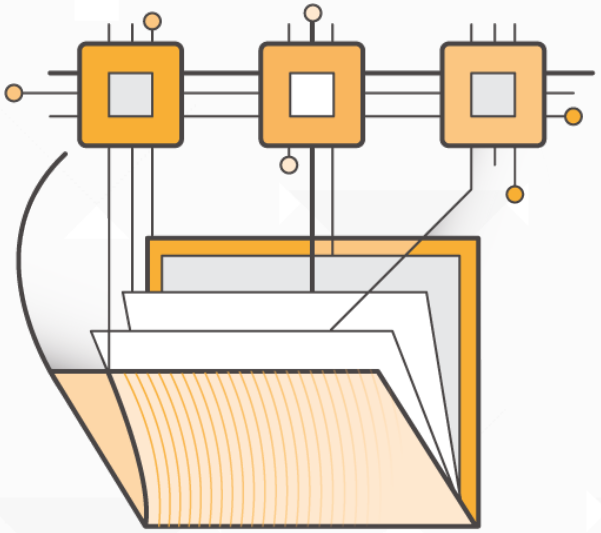

Amazon Elastic Container Service

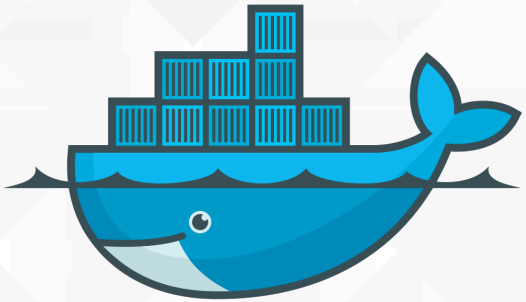

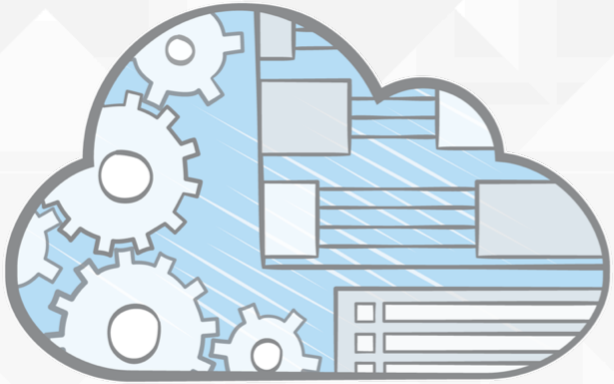

amazon web services

Key Components

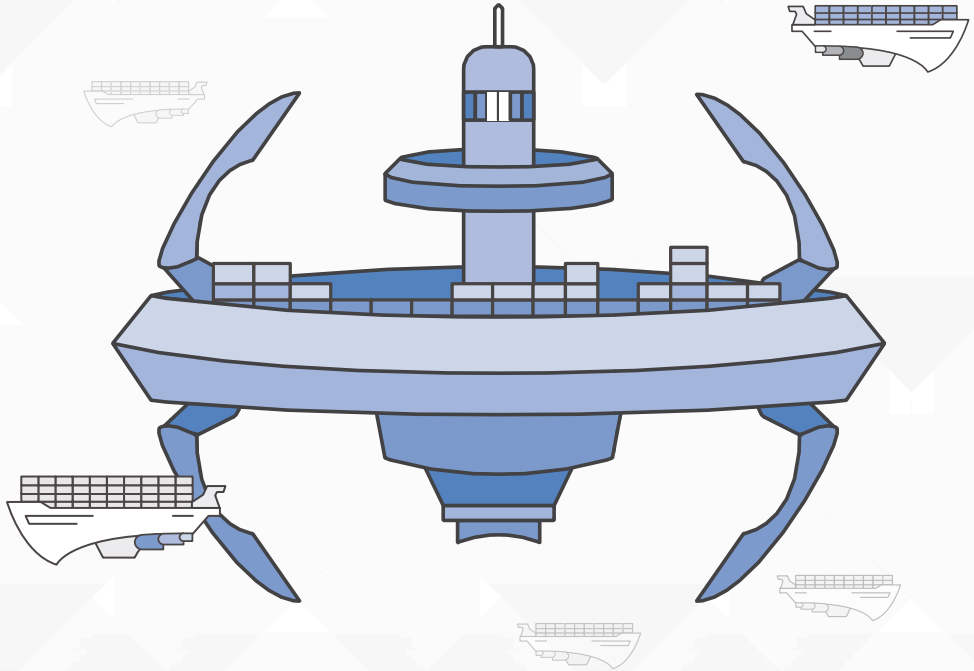

Docker Daemon

Task Definitions

Containers

Clusters

Container Instances

Typical User Workflow

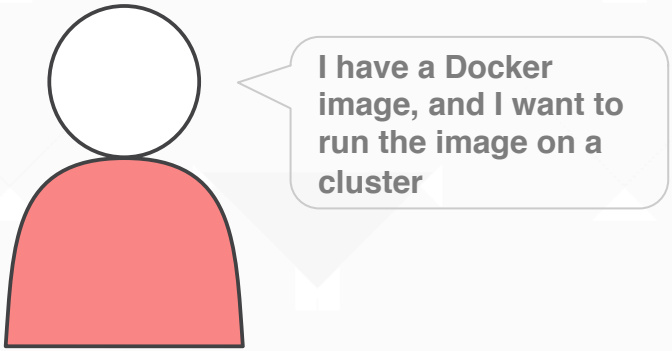

Typical User Workflow

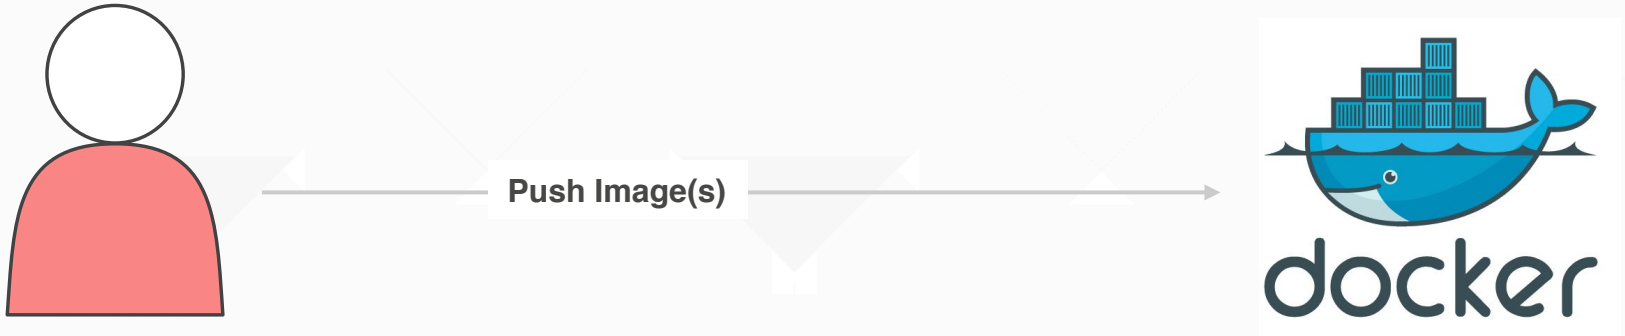

Typical User Workflow

Typical User Workflow

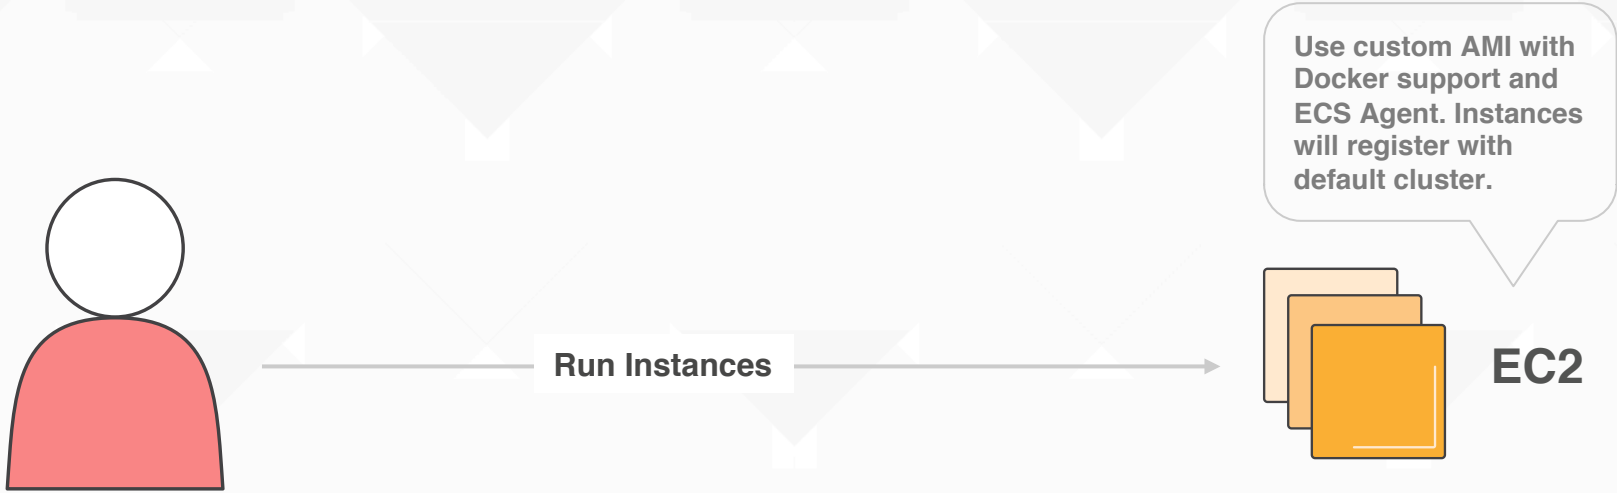

Typical User Workflow

Typical User Workflow

Typical User Workflow

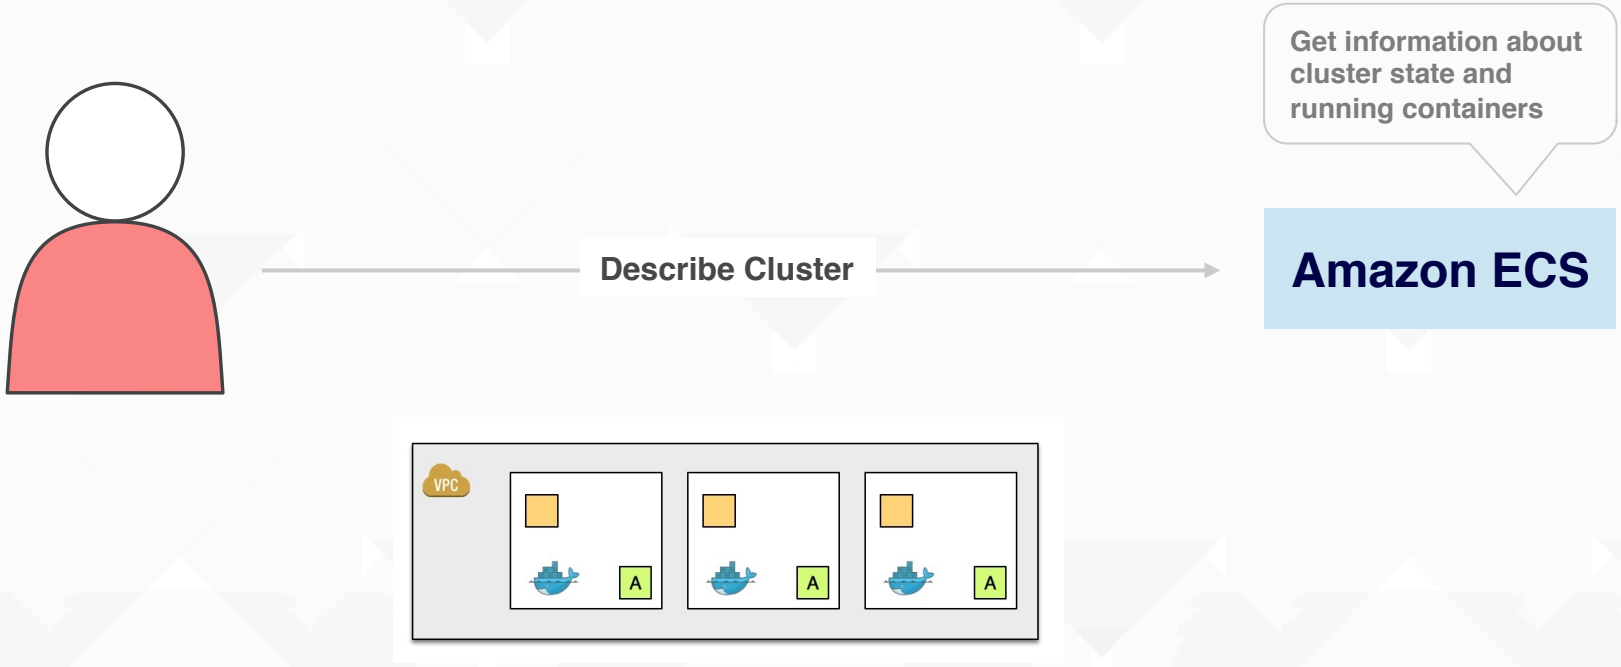

Thank you!

Jamie Kinney
jkinney@amazon.com
@jamiekinney

amazon web services

Additional resources…

- aws.amazon.com/big-data
- aws.amazon.com/compliance
- aws.amazon.com/datasets
- aws.amazon.com/grants
- aws.amazon.com/genomics
- aws.amazon.com/hpc
- aws.amazon.com/security