AWS cloud services
for Research Computing and HPC

Dr. Kevin Jorissen
BD Manager, Research & Technical Computing

BIO

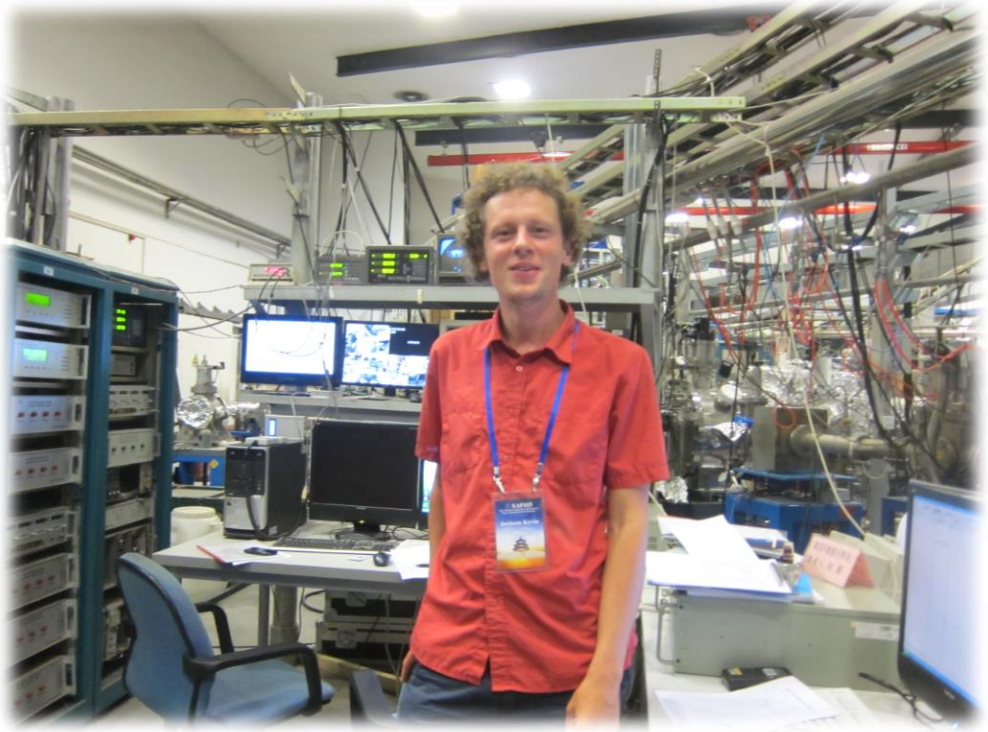

Kevin Jorissen

Seattle

Kevin has 10 years of experience in computational science, and holds a Ph.D. in Physics. He developed codes solving the quantum physics equations for light absorption by materials, taught workshops to scientists worldwide, and wrote about high performance computing in the cloud before it was fashionable. He worked as a postdoctoral researcher in Antwerp, Lausanne, Seattle, and Zurich. He contributed to the WIEN2k code (Density Functional Theory calculations of material properties, www.wien2k.at) and the FEFF code (X-ray and Electron absorption spectra, www.feffproject.org).

Kevin joined Amazon in 2015 to help accelerate the adoption of cloud computing in the scientific community globally.

Workflows for Computational Scattering Science 2013

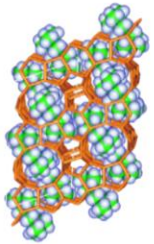

Report from the Community Workshop

January 31 – February 2, 2013
California Institute of Technology

(not so) Long ago…

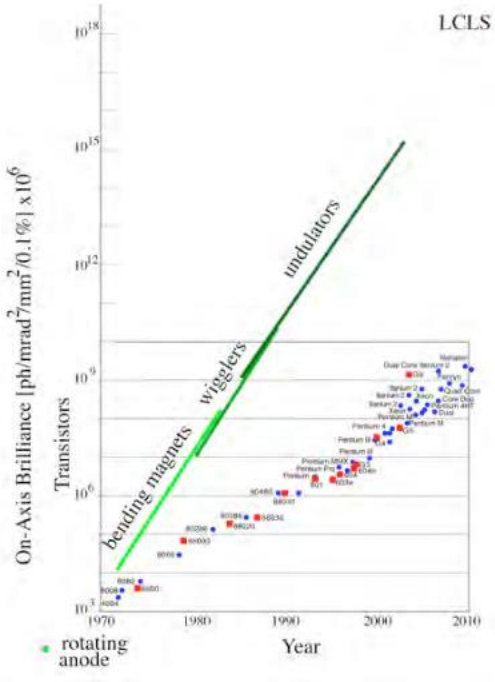

Fig. 1.1. X-ray light source performance versus year, superimposed on a typical plot demonstrating Moore's law (i.e., the exponential increase in the number of transistors on a device).

Workflows for Computational Scattering Science 2013

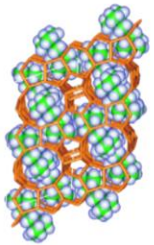

Report from the Community Workshop

January 31 – February 2, 2013
California Institute of Technology

“Where is all this data going to go? How is it going to be worked on?”

“Data is only as valuable as the decisions it enables.”

“Shouldn’t someone, maybe, budget some $ for all this?”

FF to 2017 – is the research world ready to solve challenges?

* Rapid development of technology/methods (ML, AI, IoT, …)
* Self-driving cars and personalized medicine look very real
* Cloud matured as a platform
* “Reproducibility” became “reusability”
* Expanding user base

* Research data mandates are everywhere
* We still don’t know who will pay for them
* New storage products needed for research data?

* Are policies up-to-date?
* Standards? (netcdf vs geotiff, light sources, metadata …)
* Are tools compatible with object stores (vs FS)?

* Legal challenges (NGOs, cities …)

* Privacy issues (compliance ; press ; math e.g. DP)

Public cloud as platform for research data

- Unlimited storage
- Platform for collaboration (bring everyone to the data)
- Security (job zero; sophisticated real-world controls; compliance e.g. HIPAA …)
- Toolset to extract value from data, rapidly expanding
- Services for ingesting data streams (IoT --- Kinesis --- SnowMobile)

AWS Snowmobile moves 100PB to the cloud per trip

1

Basics

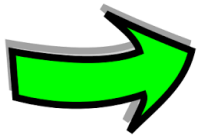

How is Cloud Computing different from on-premise?

- Request servers by the hour in the cloud, and dismiss them when your work is finished

The amount of processing power and the type of server always matches your needs. (1 GPU today, 1000 CPUs tomorrow)

- Pay for what you use (pay per server per hour, or per GB stored)

Each research budget $ is spent effectively – no paying for idle servers.

- AWS does all the undifferentiated work, and does it well (security, datacenter operations, cooling, power, hardware failures, patching/resizing/backing up databases, …)

Scientists can focus on science, instead of cluster maintenance.

Campus IT specialists can focus on helping the scientists design and operate HPC solutions.

AWS Storage Options for HPC Workloads

EFS	EC2+EBS	Amazon S3	Amazon Glacier
Highly available, multi-AZ, fully managed network-attached elastic file system.	Block storage device (SSD or HDD) for file system attached to EC2 instance. Can build parallel file system (e.g., using Intel Lustre).	Secure, durable, highly-scalable object storage. Fast access, low cost.	Secure, durable, long term, highly cost-effective object storage.
For near-line, highly-available storage of files in a traditional NFS format (NFSv4).	For near-line storage of files optimized for high I/O performance.	For long-term durable storage of data, in a readily accessible get/put access format.	For long-term storage and archival of data that is infrequently accessed.
Use for read-often, temporary working storage	**Use for high-IOPs, temporary working storage**	**Primary durable and scalable storage for research data**	**Use for long-term, lower-cost archival of research data**

Combining Compute and Storage

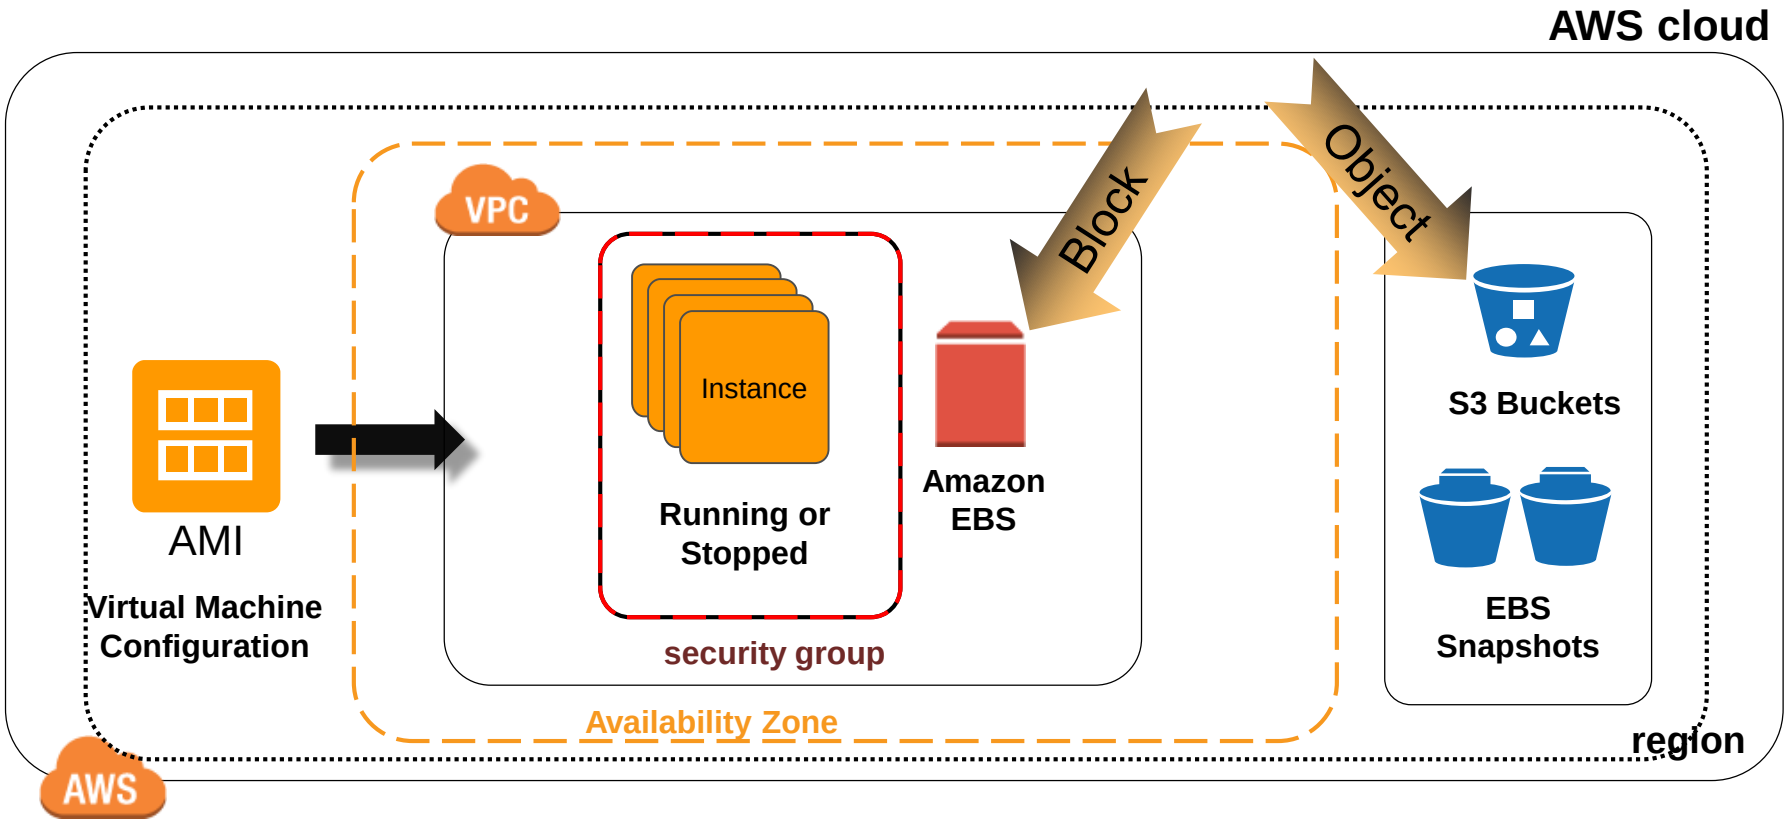

Clusters in the cloud are fit for purpose

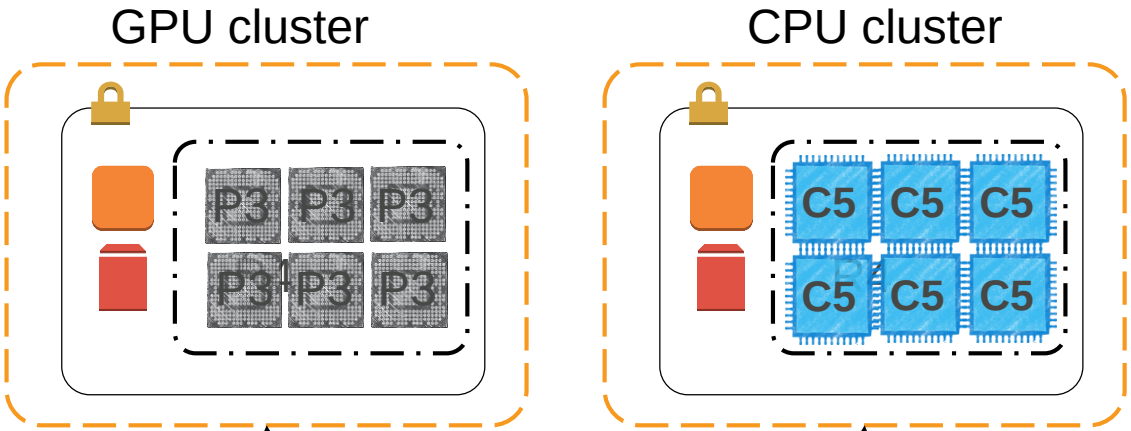

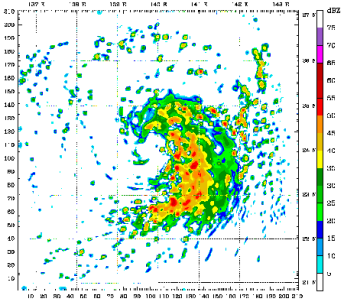

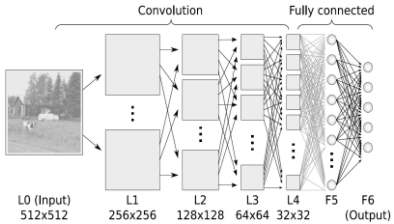

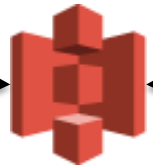

Amazon S3

Architected and Audited for Security

Certifications and accreditations for workloads that matter

AWS CloudTrail - AWS API call logging for governance & compliance

Log and review user activity

Stores data in S3, or archive to Glacier

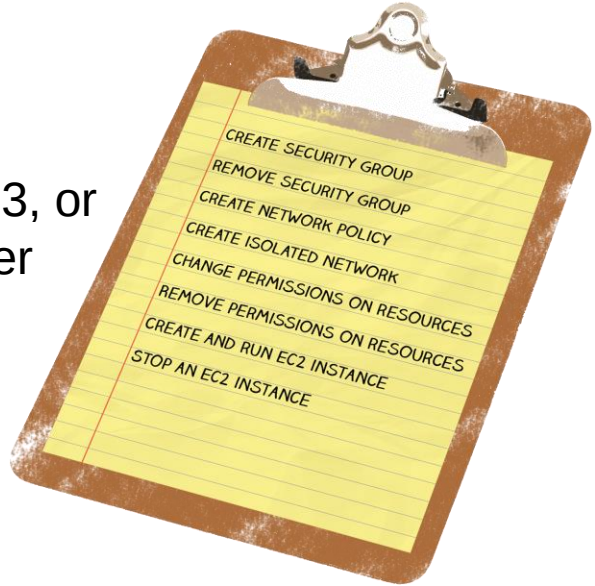

The CTO of NASA JPL, Tom Soderstrom, says that he thinks his scientific infrastructure in AWS is ***more secure* than that in his own data centers**.

S3 storage brings tons of capabilities to your data

AWS Machine Learning

Kinesis: ingest and process data streams

IoT: capture data from lots of small devices, e.g. sensor networks, smart cities, …

Athena: run SQL queries straight on S3 data (no cluster required)

Lambda: serverless compute actions on S3 objects

Rekognition: Image recognition

AppStream: build user-friendly platforms for users

EMR: Elastic MapReduce

RDS: Relational Database Service

… 90+ services to make the most of your data

Amazon Athena for Population-Scale Analyses

Amazon Athena allows you to quickly and cost-effectively query the genomic data of thousands of individuals to derive insights

Full Thousand Genomes Dataset (2.5K genomes = ~100K Exomes)

- Aggregate Allele/Genotype Frequencies + ClinVar Join
  - ~4 min
  - $0.50 for query
- Selecting data from single sample + ClinVar Join
  - ~10 sec
  - ~$0.0003 for query

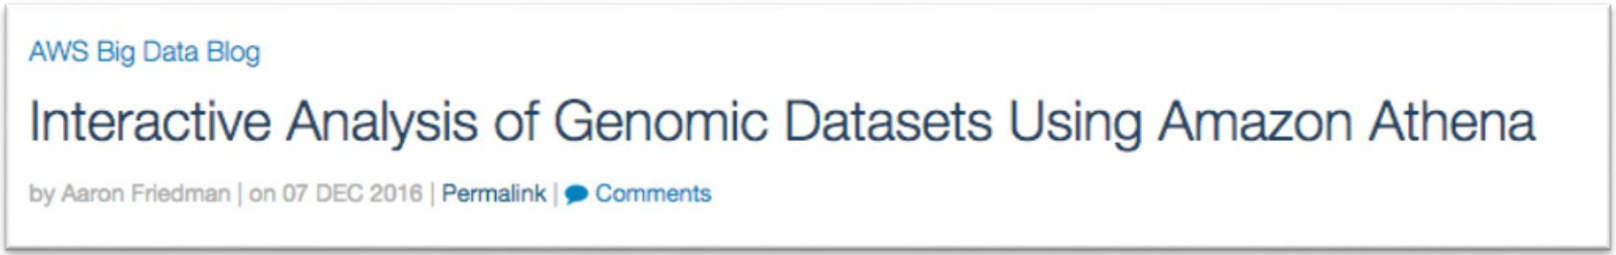

http://amzn.to/2m8zU0A

Scientific computing offerings:
AWS Marketplace – Technology Partners

2

AWS and collaboration

Increase scientific impact

Benefits of AWS Cloud for collaboration

- Built from the ground up with sophisticated, real-world security: share without giving up security.
- Use AWS worldwide network and data centers to reach your collaborators.
- Collaborators can analyze your shared data in their own account, and run your shared applications in their own account, using the “infinite” EC2 compute and data analytics capacity available, at their own expense.
- Not necessary for everyone to download a copy of the dataset: everyone can bring analytics to central copy.
- You retain full ownership. Data never leaves a country (“data sovereignty”) unless you explicitly move it.

It’s typically consuming and expensive to acquire, store, and analyze large data sets. Accessing data at scale is often a prohibitive challenge.
Sharing data on AWS makes it accessible to a large and growing community of researchers who use the AWS cloud.

Enabling Global Collaboration

Global Infrastructure

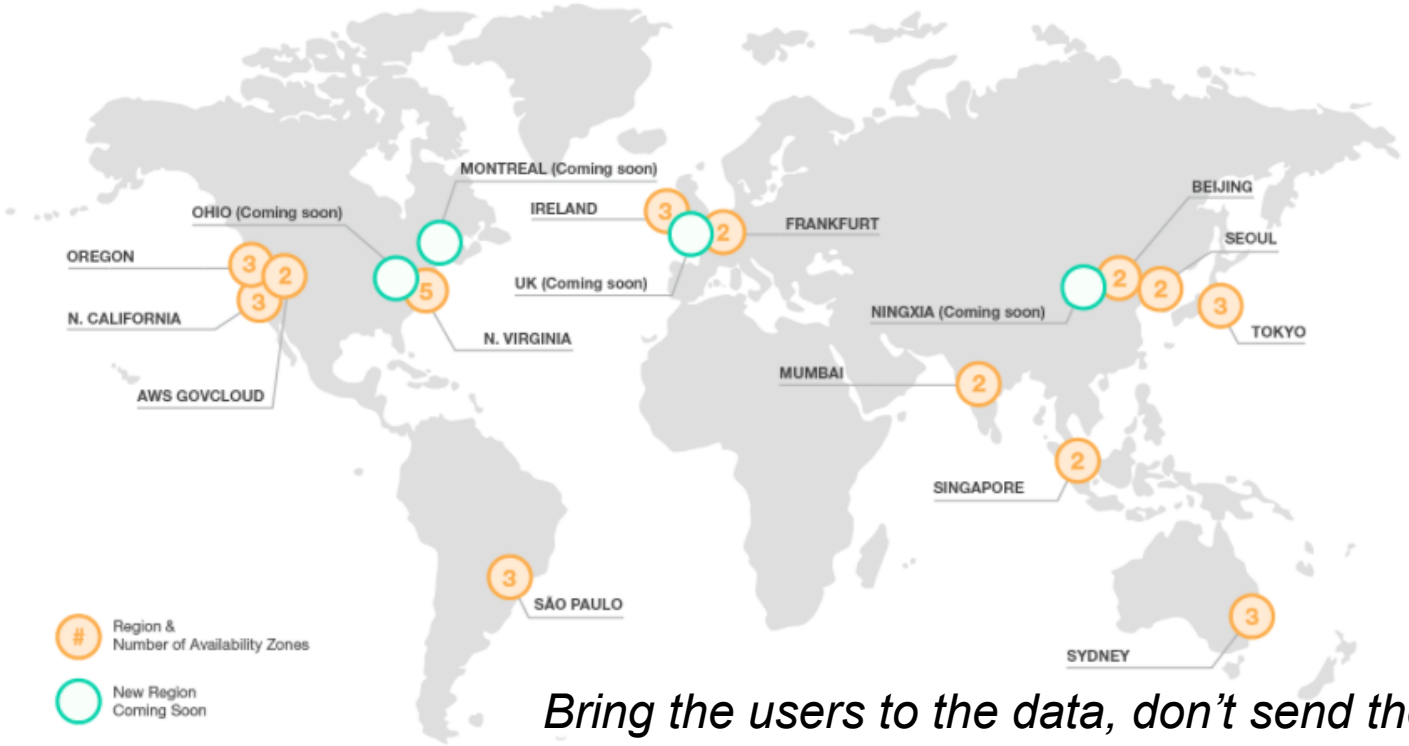

Bring the users to the data, don't send the data to the users

Public datasets on AWS

To enable more innovation, AWS hosts a selection of datasets that anyone can access for free. Data in our public datasets is available for rapid access to our flexible and low-cost computing resources.

Life Science

- **TCGA & ICGC**
- 1000 Genomes
- Genome in a Bottle
- Human Microbiome Project
- 3000 Rice Genome

Earth Science

- Landsat
- NEXRAD
- NASA NEX

Internet Science

- Common Crawl Corpus
- Google Books Ngrams
- Multimedia Commons

Earth on AWS

Build planetary-scale applications in the cloud with open geospatial data.

aws.amazon.com/earth

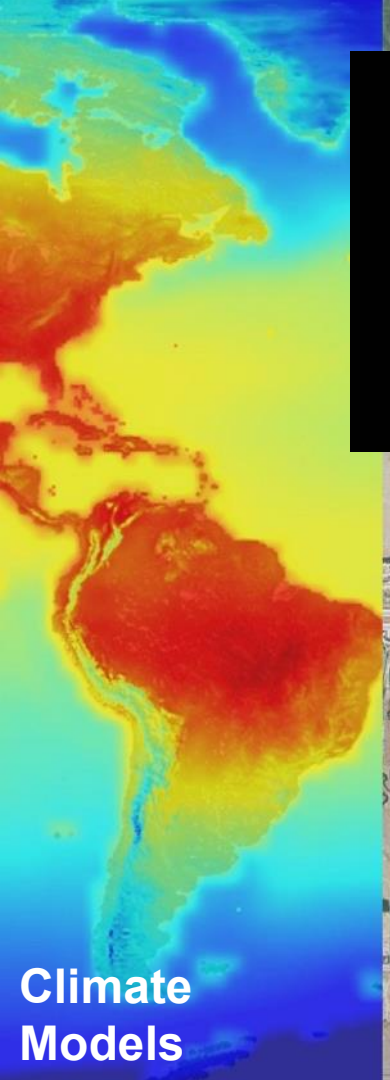

Climate Models

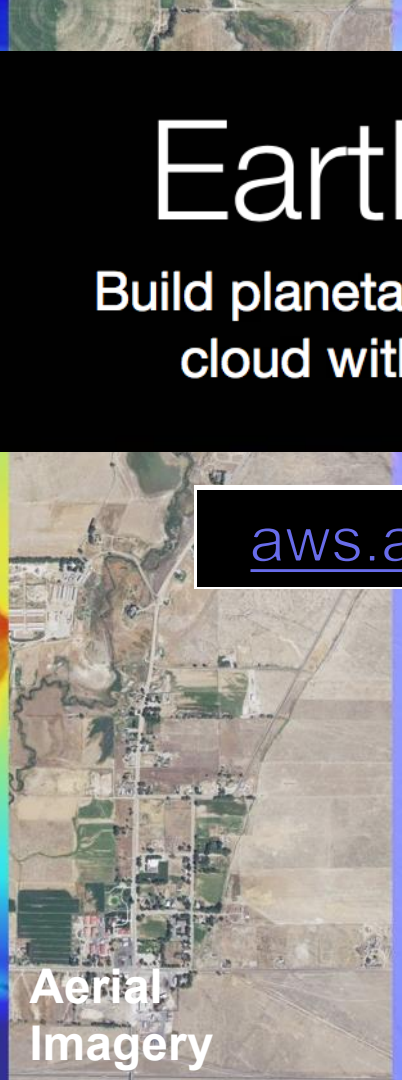

Aerial Imagery

Elevation Models

Satellite Imagery

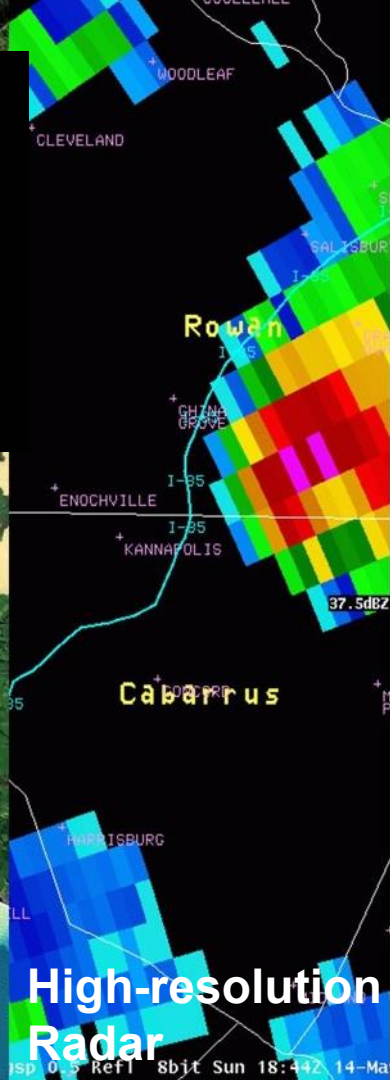

High-resolution Radar

NEXRAD on AWS

270TB of volume scan files and real-time chunks as objects on Amazon S3.

- Data can be accessed programmatically via a RESTful interface and quickly deployed to any of our products for analysis and processing.
- Amazon Simple Notification Service (SNS) allows subscription to notifications of new data.

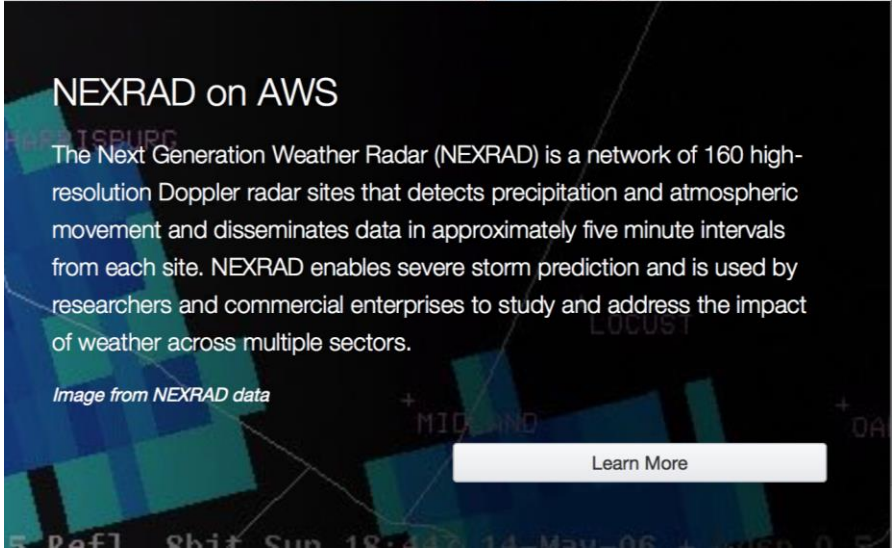

NEXRAD on AWS

Immediate usage:

- Climate Corporation **cut two weeks** out of an analysis pipeline
- Increased NEXRAD usage **2.3X**
- A weather data company stopped storing their own NEXRAD archive, **freeing up revenue** to build new products.

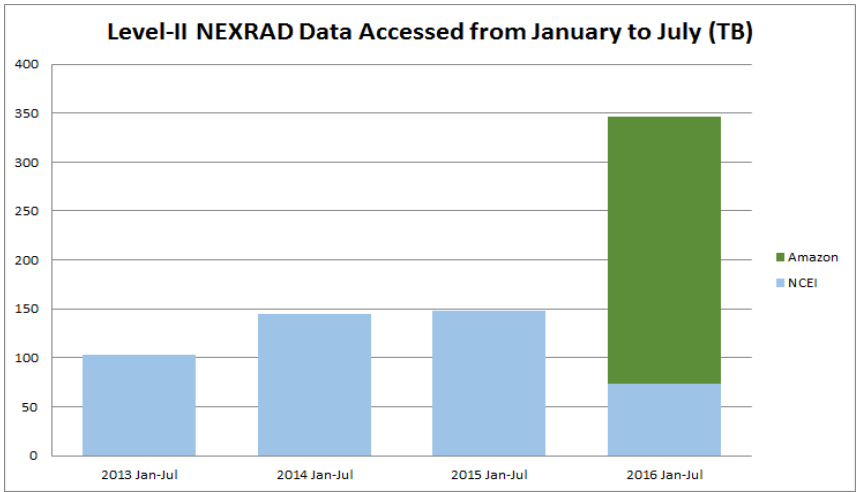

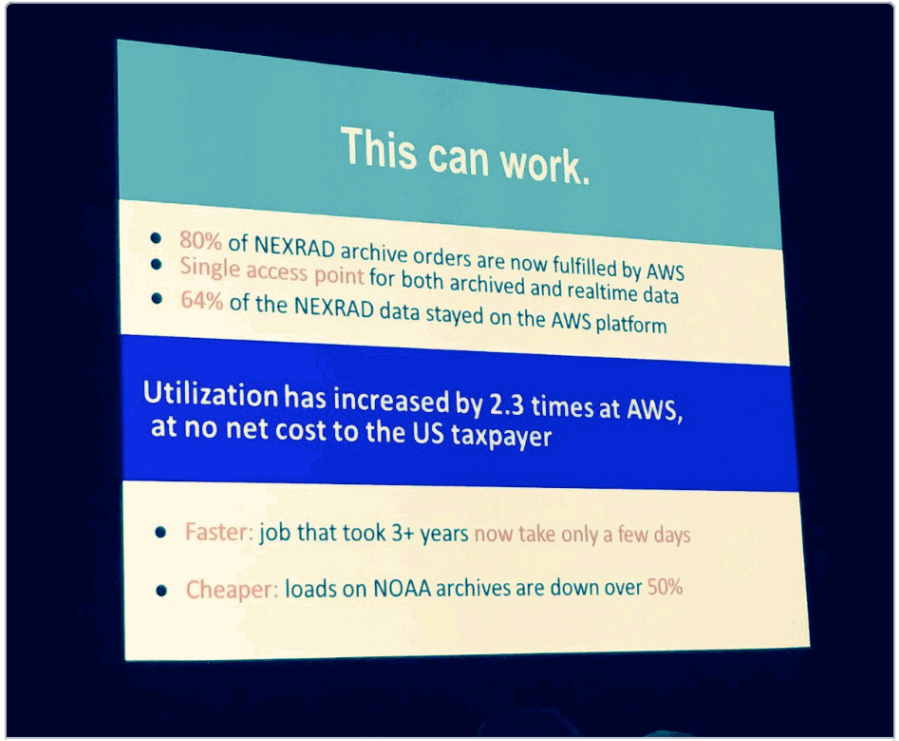

DJ Patil @DJ44

Follow

Wow. What happens when @NOAA puts their data on the cloud. #WHOpenData 👊👍🙌🇺🇸

4:04 PM - 28 Sep 2016

92 124

Landsat on AWS

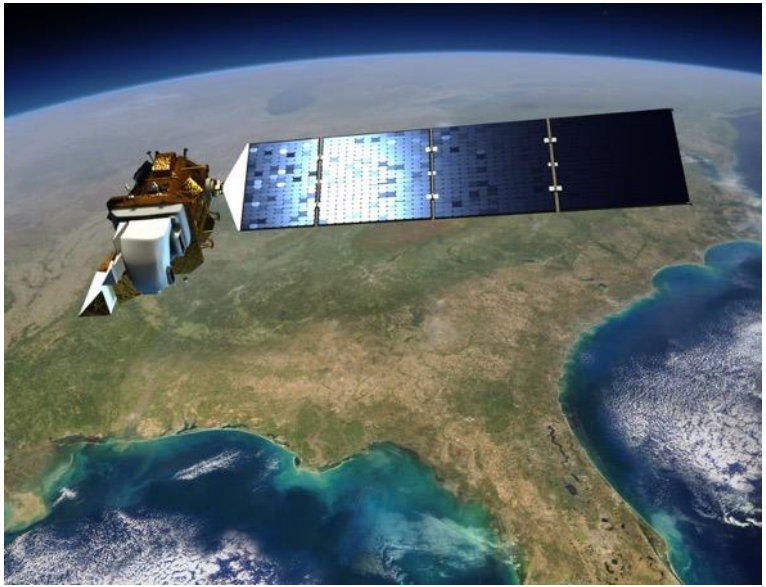

Landsat 8 satellite

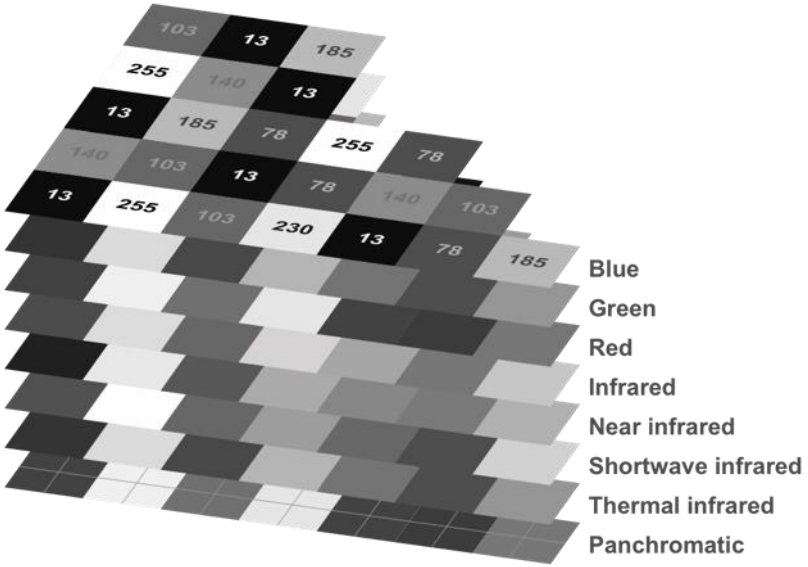

Raster data

Within the first year of Landsat on AWS, the data has been requested over 1 billion times, globally. Over 400,000 scenes are now available

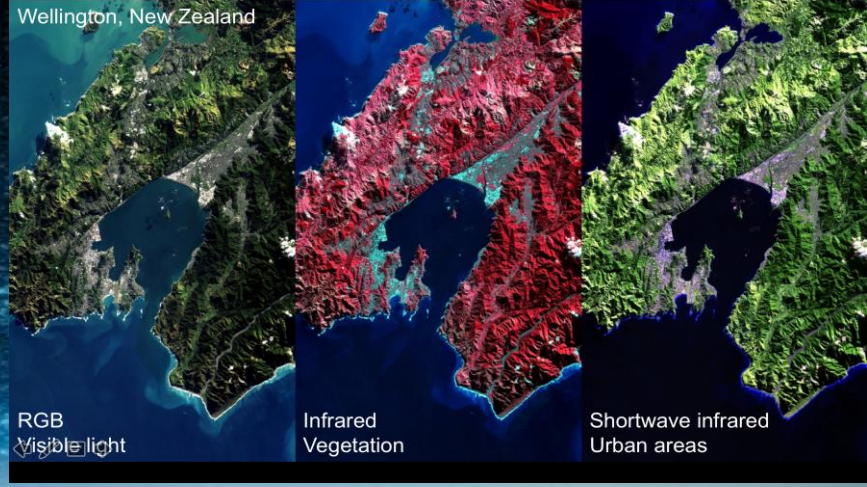

Science, not servers

New ways of working with research data

- Spot market for low-cost research computing
- Containers
- Accelerators
- Serverless computing
- Machine Learning

1. Spot: ICRAR/CHILES finding neutral hydrogen galaxies

- A global radio astronomy consortium (led by Columbia University in New York) needed to process observational data from the Very Large Array telescope in New Mexico. A 12-hour SLA meant they need ~$2 million of conventional HPC hardware. This was impossible because they had only $50k.
- Using the EC2 Spot market in AWS's northern Virginia region, they were able to deploy their HPC workload at a much larger scale -- so they always beat their SLA -- whilst averaging only $1,200 per month of EC2 compute resources, well within their 2-year budget of $50k.
- The project **produced a major discovery which smashed the previous record** for identifying a **neutral hydrogen galaxy** by nearly twice the redshift of its predecessor.

The “Spot Market” stretches your research budget $$

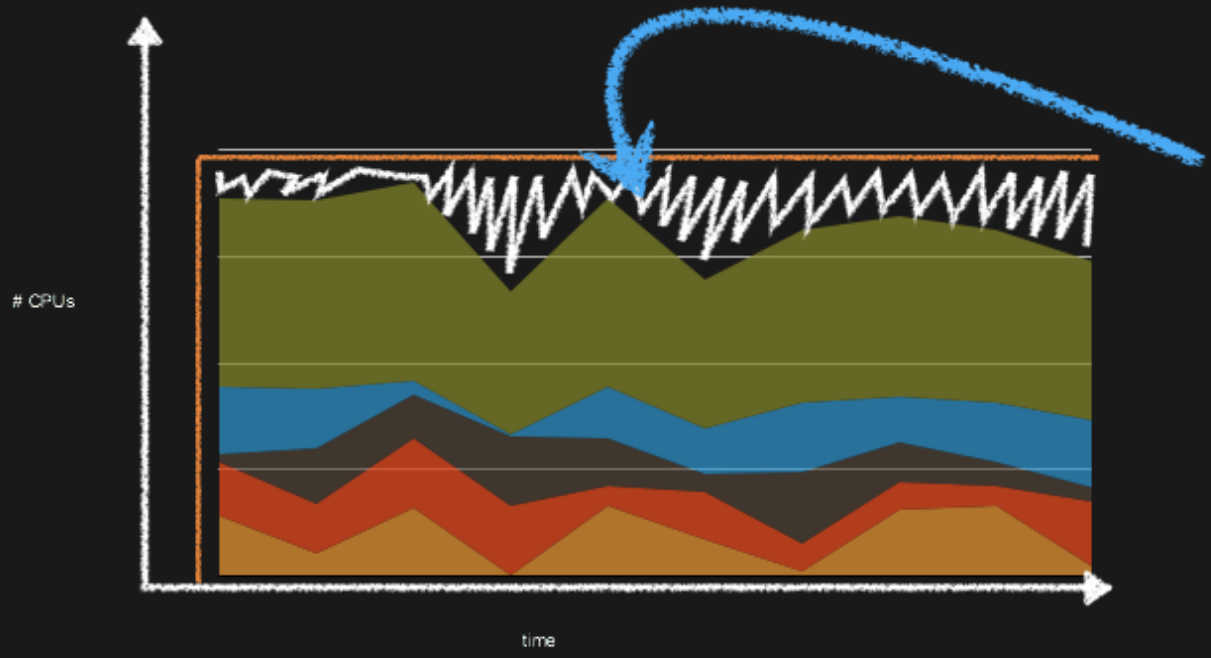

Spot Market

Our ultimate space filler.

Spot Instances allow you to name your own price for spare AWS computing capacity.

Great for workloads that aren’t time sensitive, and especially popular in research (hint: it’s really cheap).

Elasticity: Natural Language Processing at Clemson University – 550,000 cores & EC2 Spot Instances

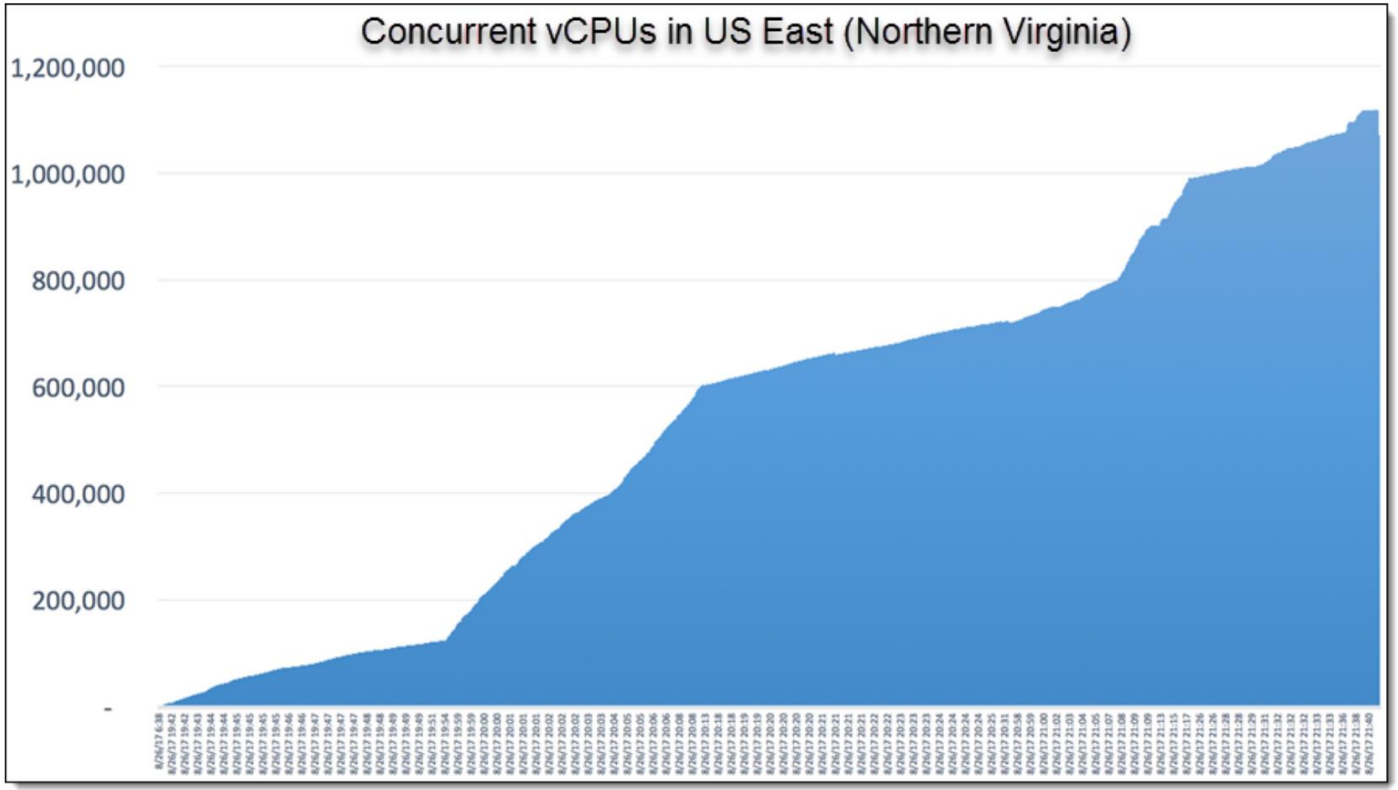

https://aws.amazon.com/blogs/aws/natural-language-processing-at-clemson-university-1-1-million-vcpus-ec2-spot-instances/

Elasticity: Natural Language Processing at Clemson University – 550,000 cores & EC2 Spot Instances

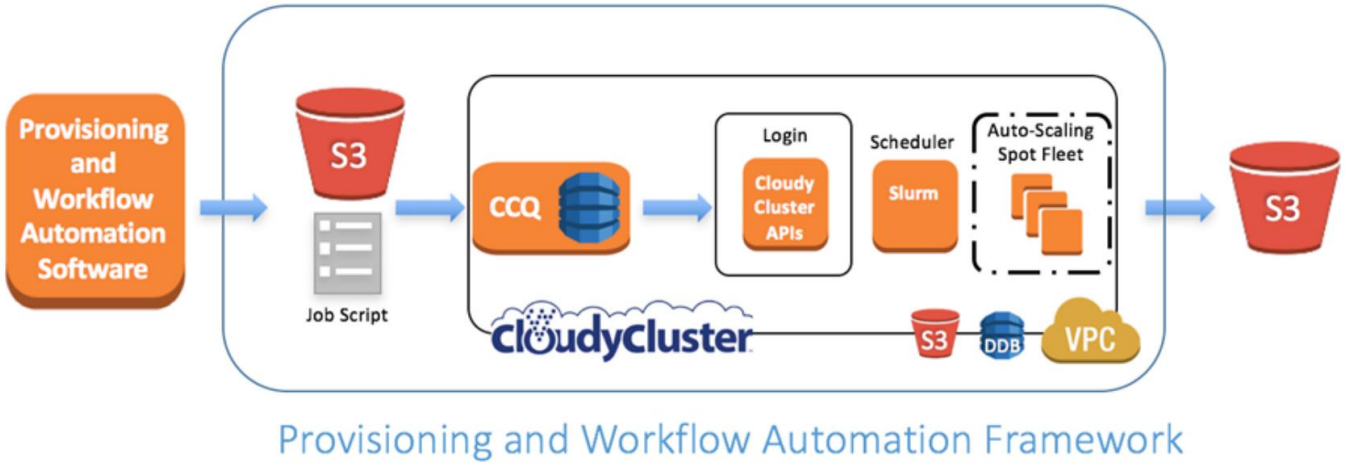

Provisioning and Workflow Automation Framework

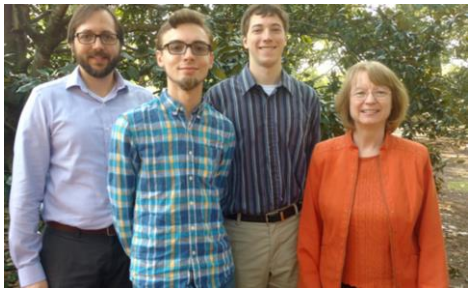

"I am absolutely thrilled with the outcome of this experiment. The graduate students on the project are amazing. They used resources from AWS and Omnibond and developed a new software infrastructure to perform research at a scale and time-to-completion not possible with only campus resources. Per-second billing was a key enabler of these experiments." – Prof. Amy Apon, Co-Director of the Complex Systems, Analytics and Visualization Institute

https://aws.amazon.com/blogs/aws/natural-language-processing-at-clemson-university-1-1-million-vcpus-ec2-spot-instances/

2. Evolving Compute Abstractions – Containers/Batch

Physical

On-premise server

Virtualisation

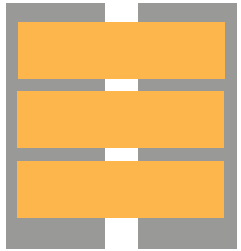

EC2 instance

Containerization

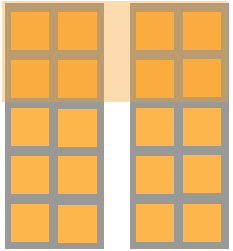

Docker

Serverless

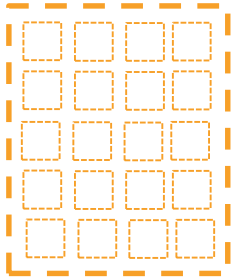

Lambda function

AWS Batch – a managed service for container based jobs

- Completely manages compute infrastructure for you
- Only define your applications and resource requirements
  - Shell scripts
  - Linux executables
  - Docker images
- Batch takes care of the rest
  - Jobs can be Lambda functions or applications
  - Array Jobs for Parametric Sweeps
  - Jobs are submitted to a Job Queue
  - Sophisticated job dependencies
  - Automated job retries
  - MPI next! (But, not right now.)

Example Genomics workflow on AWS

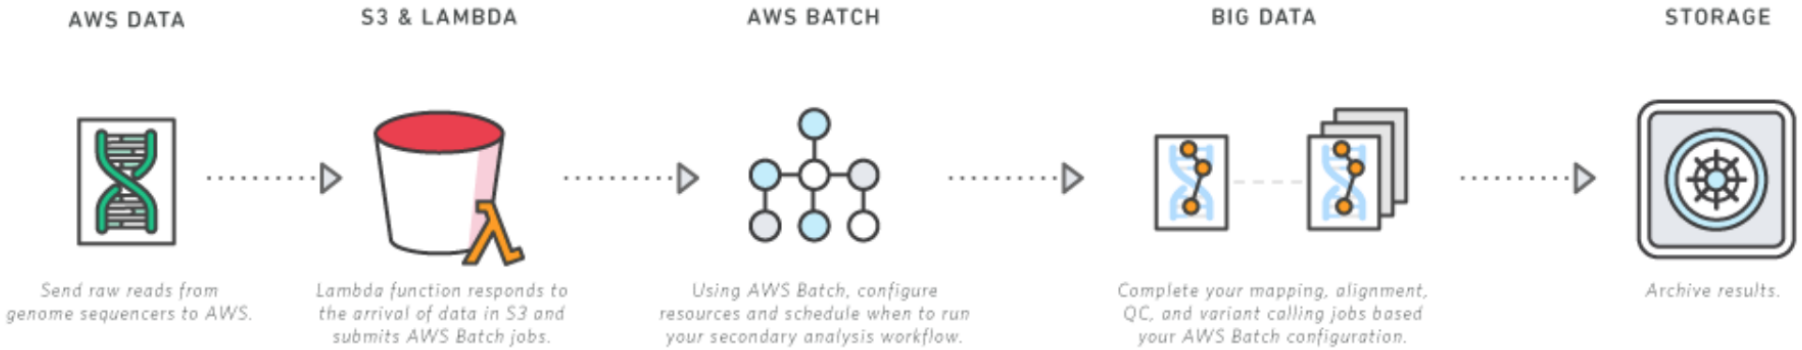

Example Genomics workflow on AWS

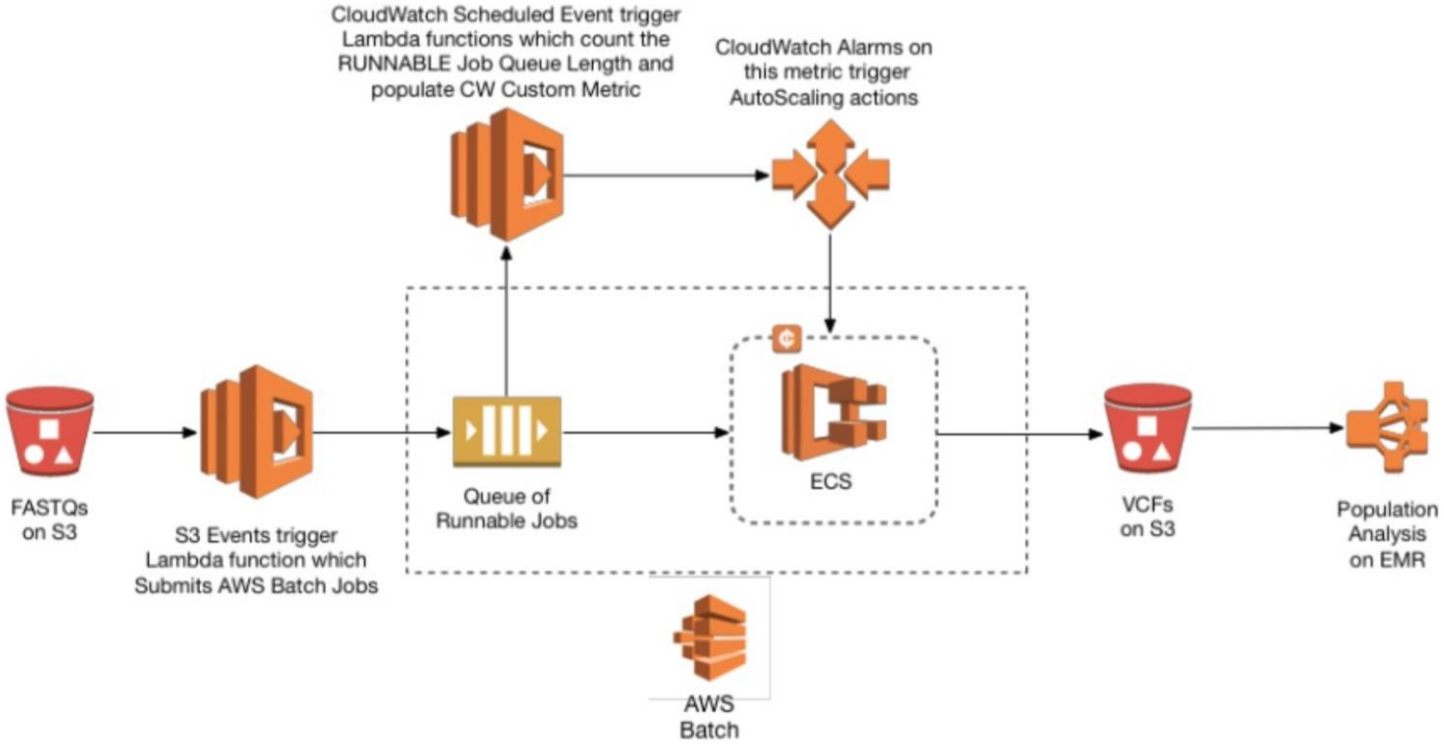

3. Serverless Computing: AWS Lambda

AWS Lambda is a service which allows for software functions in a variety of languages to be deployed into the cloud natively, and to be triggered directly or driven by events in the cloud. The infrastructure (hardware, operating system and software environment) for Lambda is managed by AWS and scales rapidly.

Bring your own code

Node.JS, Java, Python

Java = Any JVM based language such as Scala, Clojure, etc.

Bring your own libraries

Simple resource model

- Select memory from 128MB to 1.5GB in 64MB steps
- CPU & Network allocated proportionately to RAM
- Reports actual usage

Serverless Computing: AWS Lambda

AWS Lambda is a service which allows for software functions in a variety of languages to be deployed into the cloud natively, and to be triggered directly or driven by events in the cloud. The infrastructure (hardware, operating system and software environment) for Lambda is managed by AWS and scales rapidly.

No Servers to Manage

AWS Lambda automatically runs your code without requiring you to provision or manage servers. Just write the code and upload it to Lambda.

Continuous Scaling

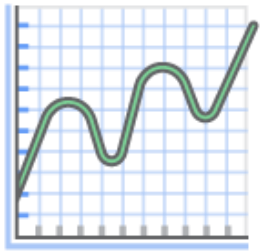

AWS Lambda automatically scales your application by running code in response to each trigger. Your code runs in parallel and processes each trigger individually, scaling precisely with the size of the workload.

Subsecond Metering

With AWS Lambda, you are charged for every 100ms your code executes and the number of times your code is triggered. You don't pay anything when your code isn't running.

Two examples of HPC on Lambda

CSIRO have built quickly scaling genomics analysis on AWS Lambda

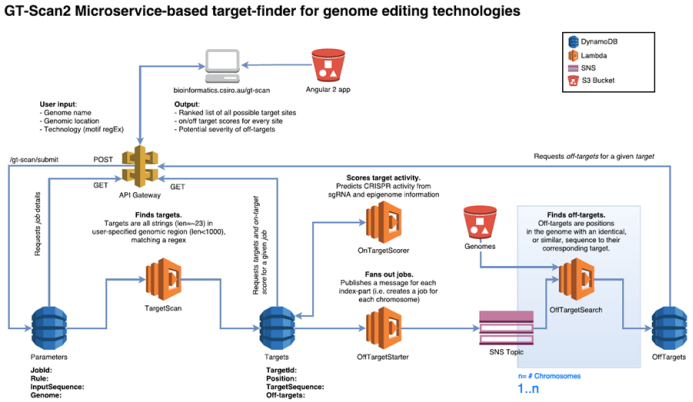

PyWren.io

riselab UC Berkeley

```
def my_function(b):
  x = np.random.normal(0, b, 1024)
  A = np.random.normal(0, b, (1024, 1024))
  return np.dot(A, x)

  pwex = pywren.default_executor()
  res = pwex.map(my_function, np.linspace(0.1, 100, 1000))
```

PyWren lets you run your existing python code at massive scale via AWS Lambda

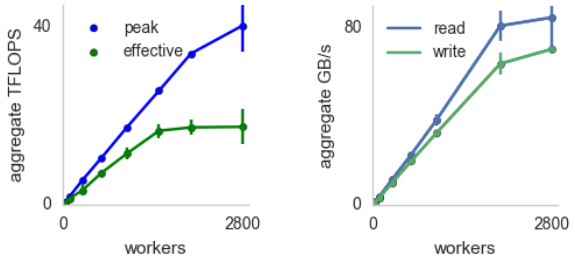

Pywren: Lambda in the context of Grid Computing

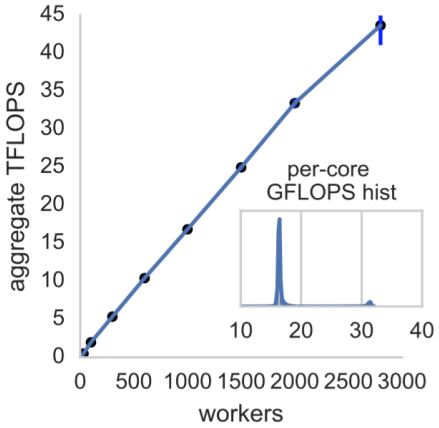

Figure 2: **Running a matrix multiplication benchmark inside each worker, we see a linear scalability of FLOPs across 3000 workers.**

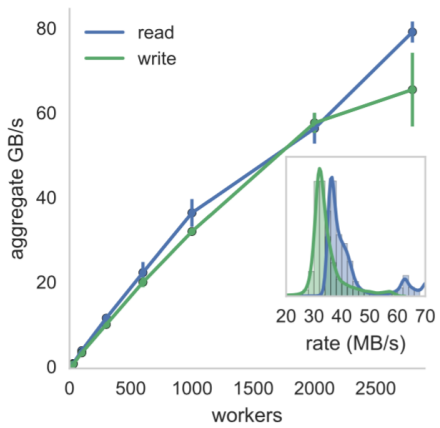

Figure 3: **Remote storage on S3 linearly scales with each worker getting around 30 MB/s bandwidth (inset histogram).**

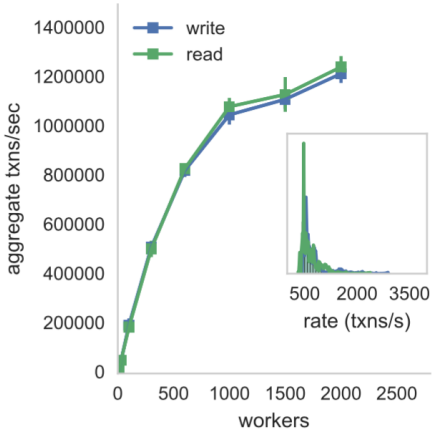

Figure 4: **Remote key-value operations to Redis scales up to 1000 workers. Each worker gets around 700 synchronous transactions/sec.**

Source: “Occupy the Cloud: Distributed Computing for the 99%”
https://arxiv.org/pdf/1702.04024.pdf

riselab
UC Berkeley

Pywren democratizes parallel scaling capabilities that used to be the sole preserve of large super-computing centers.

CSIRO – Cloud-based CRISPR prediction

CSIRO **used AWS Lambda functions to completely re-engineer a cluster HPC workload** to identify optimal gene editing sites for personalized treatment.

- "GTScan-2" job runtime varies from 1 second to 5 minutes, because the complexity of the targeted gene can vary dramatically.
- Rapid turn-around times are needed for real-time analysis.

Server-based solutions can't be provisioned efficiently to handle the variability and quick turn-around – either you have lots of servers sitting idle, or you have to wait minutes for new servers to spin up.

Deployed using AWS Lambda, the GTScan-2 runtime is stable at a few minutes **per complete job**, no matter how many jobs (i.e. genetic samples) are sent to it.

Re-casting of the code took **only a few weeks**.

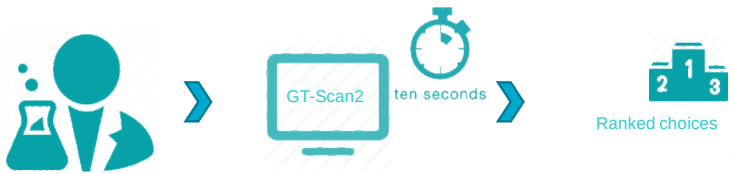

CSIRO – CRISPR search with AWS Lambda

GT-Scan2.0 is implemented as a microservices architecture using AWS Lambda

Serverless:

- Does not require users to have high-compute power

Scalable:

- Can be easily scaled to whole genome analysis

Also implement as a “stand-alone”

- Can be run on local servers
- Can incorporate your own ChIP-seq data rather than public data

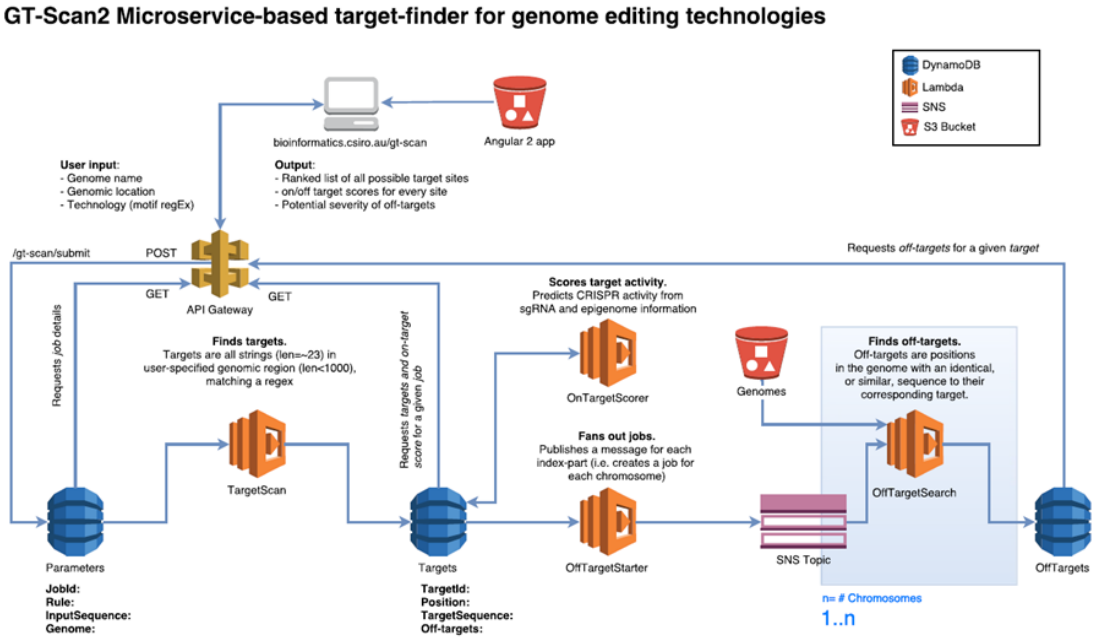

4. IoT Real-time Flood Mapping with PetaBencana.id

Critical Web Services for Emergency Management

Custom interface for Emergency Control Room

Real time flood data entered into system via web interface and sourced from Twitter

IoT water level sensing devices, to cheaply increase the monitoring across the waterway network in Jakarta

AWS Internet of Things

Prototype flood sensor

Build IoT water level sensing devices, to cheaply increase the monitoring across the waterway network in Jakarta.
Use ultrasound sonar devices to measure distance.
Transmit data securely across the Internet to the AWS cloud and store in the database.

5. Machine Learning – biggest trend today

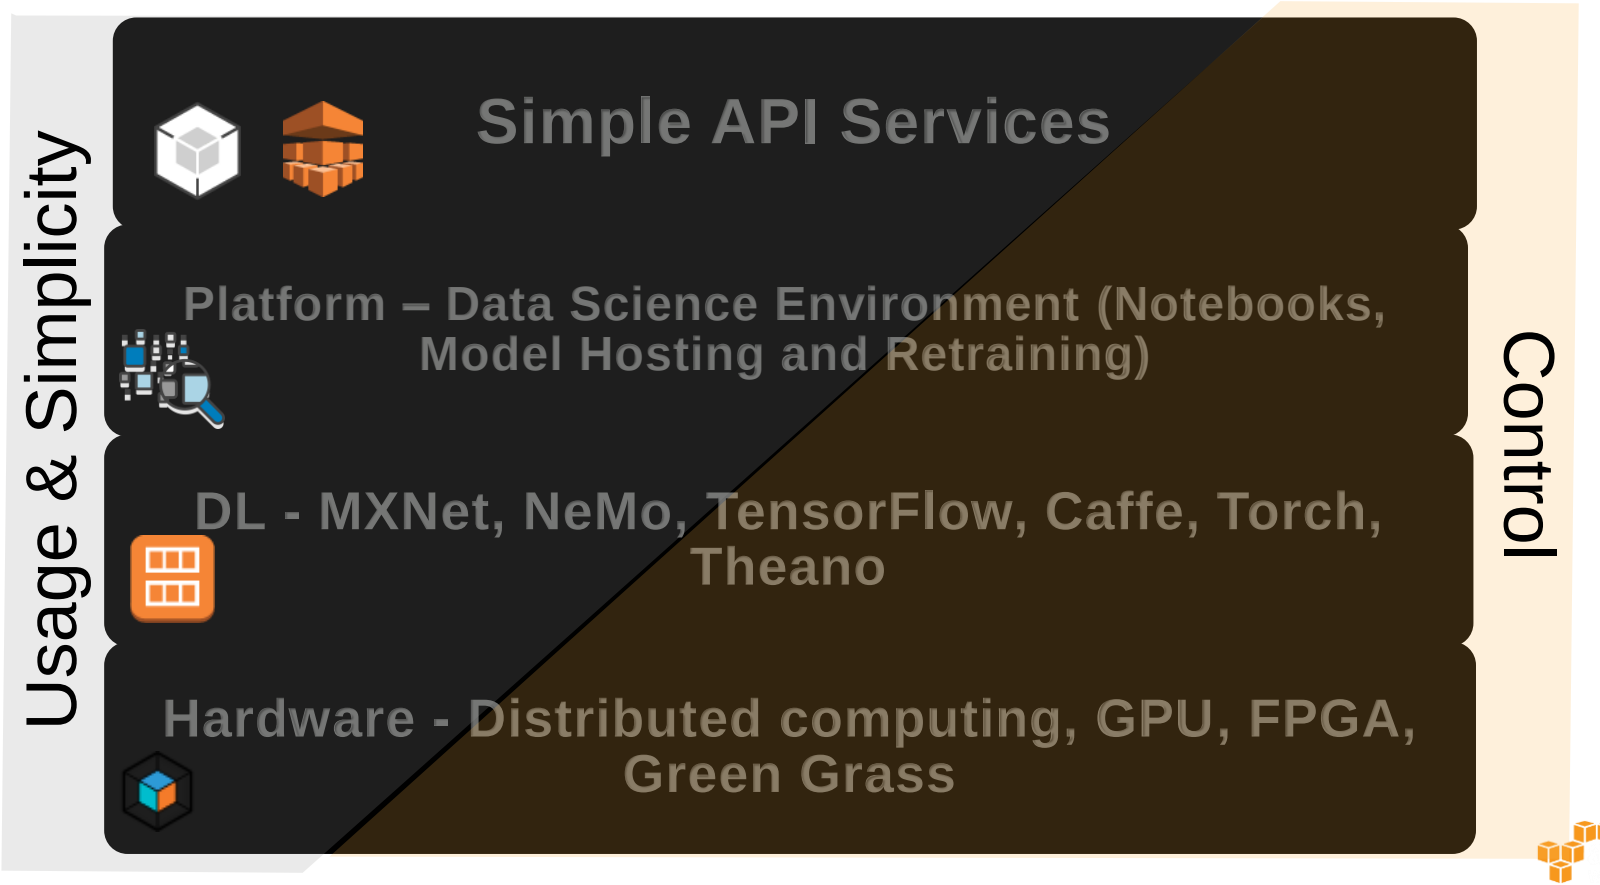

AMPLab & RISELab (Algorithms, Machines, People)

- **Collaborative 5-year research effort** between UC Berkeley, NSF, and industry partners (2012-2016) – AWS is founding partner.
- Students and researchers AMPLab used AWS to rapidly prototype and develop new systems at a scale and with a speed not possible before.
- Resulted in Apache Spark, developed on AWS, and integrated with AWS core services.

From batch data to advanced analytics

Algorithms
- Machine Learning, Statistical Methods
- Prediction, Business Intelligence

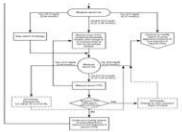

Machines
- Clusters and Clouds
- Warehouse Scale Computing

People
- Crowdsourcing, Human Computation
- Data Scientists, Analysts

Berkeley Data Analytics Stack

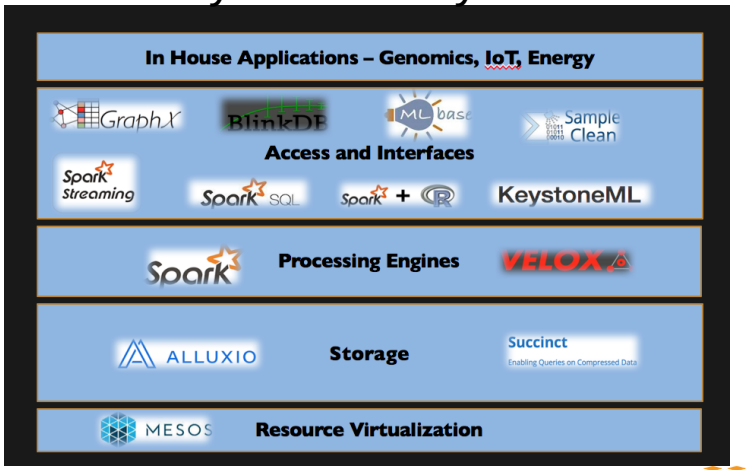

https://amplab.cs.berkeley.edu

AMPLab & RISELab (Real-time Intelligent Secure Execution)

- Collaborative 5-year effort between UC Berkeley, National Science Foundation, and industry partners. (2017-2021) – AWS is founding partner

Data only as valuable as the decisions it enables

Develop open source platforms, tools, and algorithms for intelligent real-time decisions on live-data

Typical decision system

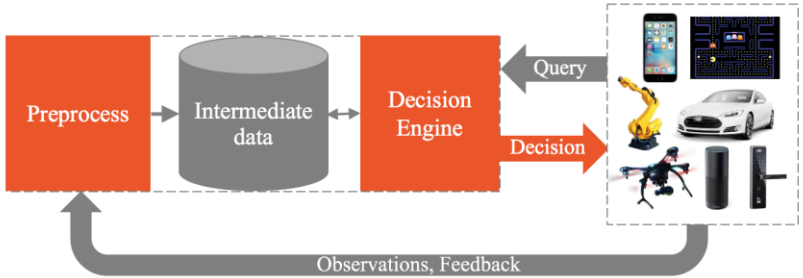

From live data to real-time decisions

7. Scale in the Cloud

Metric	Count
Compute Hours of Work	2,312,959 hours
Compute Days of Work	96,373 days
Compute Years of Work	264 years
Molecule Count	205,000 materials
Run Time	< 18 hours
Max Scale (cores)	156,314 cores across 8 regions
Max Scale (instances)	16,788 instances

CYCLECOMPUTING
LEADER IN CONDOR GRID COMPUTING SOLUTIONS

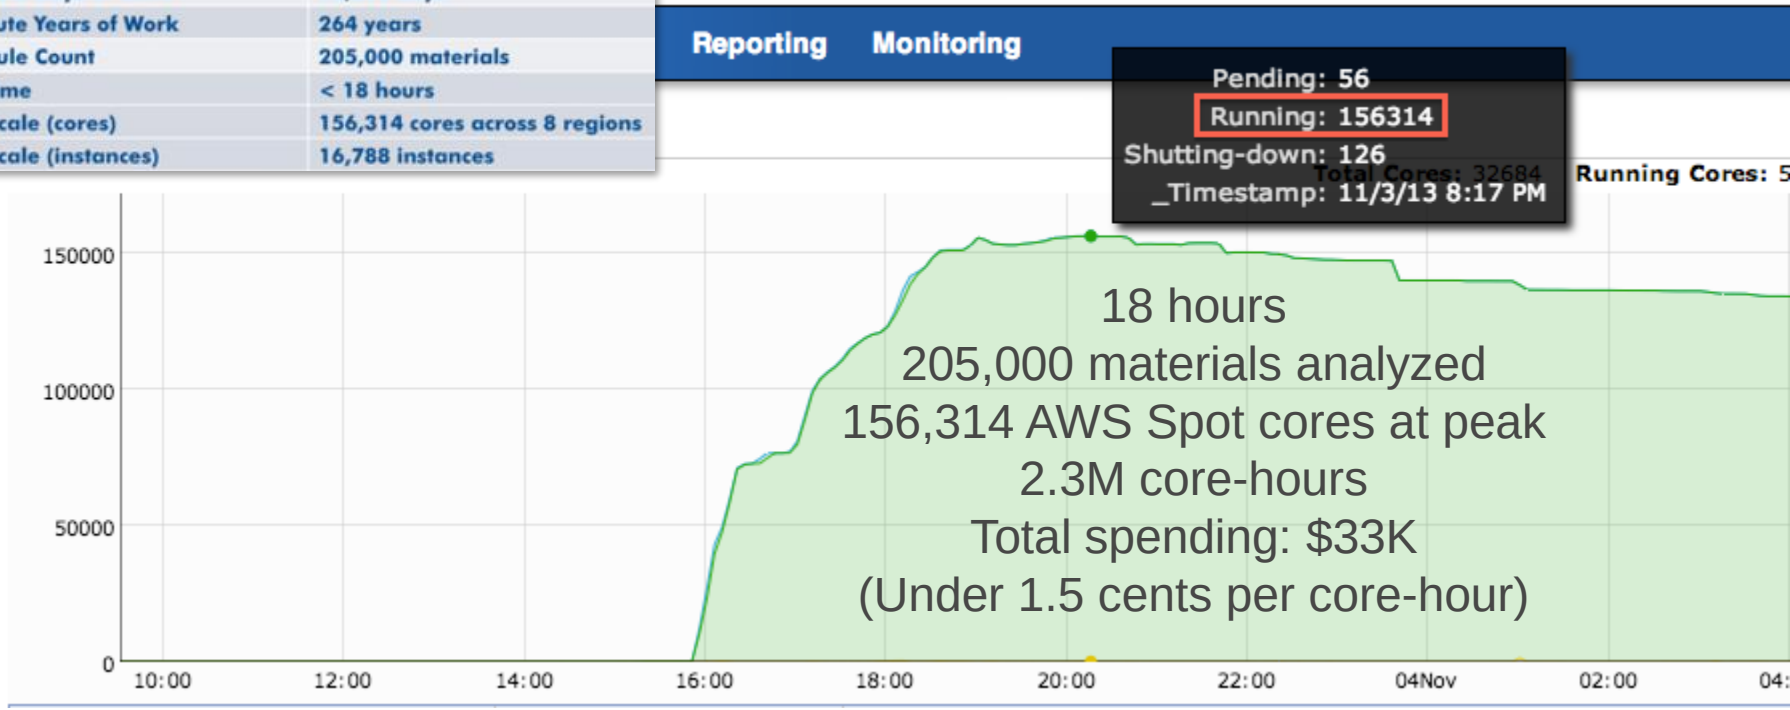

Fastest ever analysis of 1000 genomes

- 1,000 pediatric whole genomes
- Average 40X coverage
- Max 60X coverage
- Total runtime 2h 25min
- 1000 FPGA instances

edico genome

Children's Hospital of Philadelphia

… Available in “AWS App Store” for ~$24 / genome

Life Sciences

DNAnexus

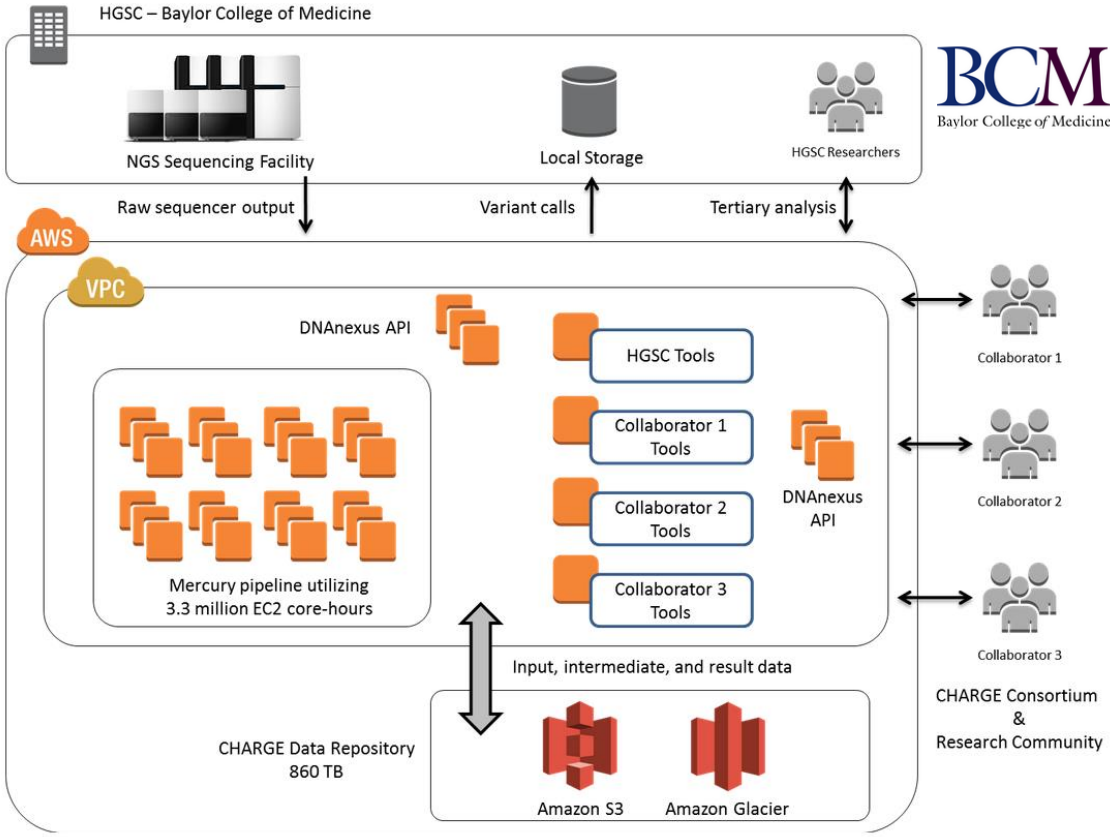

Baylor CHARGE project:

- Genomics analysis on 14,000 participants
- 24 terabases of sequencer content each month
- 1PB of raw data storage
- 21,000 AWS compute cores at peak
- Initial analysis completed in 10 days

ATLAS: High Throughput Computing at Scale

High Energy Physics

- Discovery of the Higgs Boson Particle
- Added 58,000 Spot Cores Elastically
- Monte Carlo Simulations Searching for Particles
- Reduced workload from 6 weeks to 10 days

BROOKHAVEN NATIONAL LABORATORY

CERN

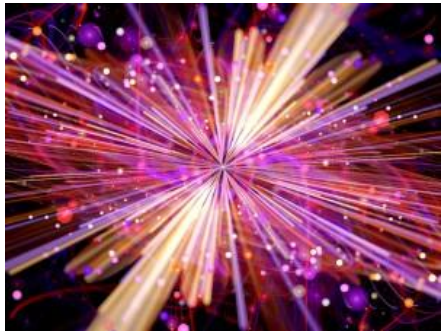

Source: https://aws.amazon.com/blogs/aws/experiment-that-discovered-the-higgs-boson-uses-aws-to-probe-nature/

ATLAS Architecture with AWS

ATLAS Architecture

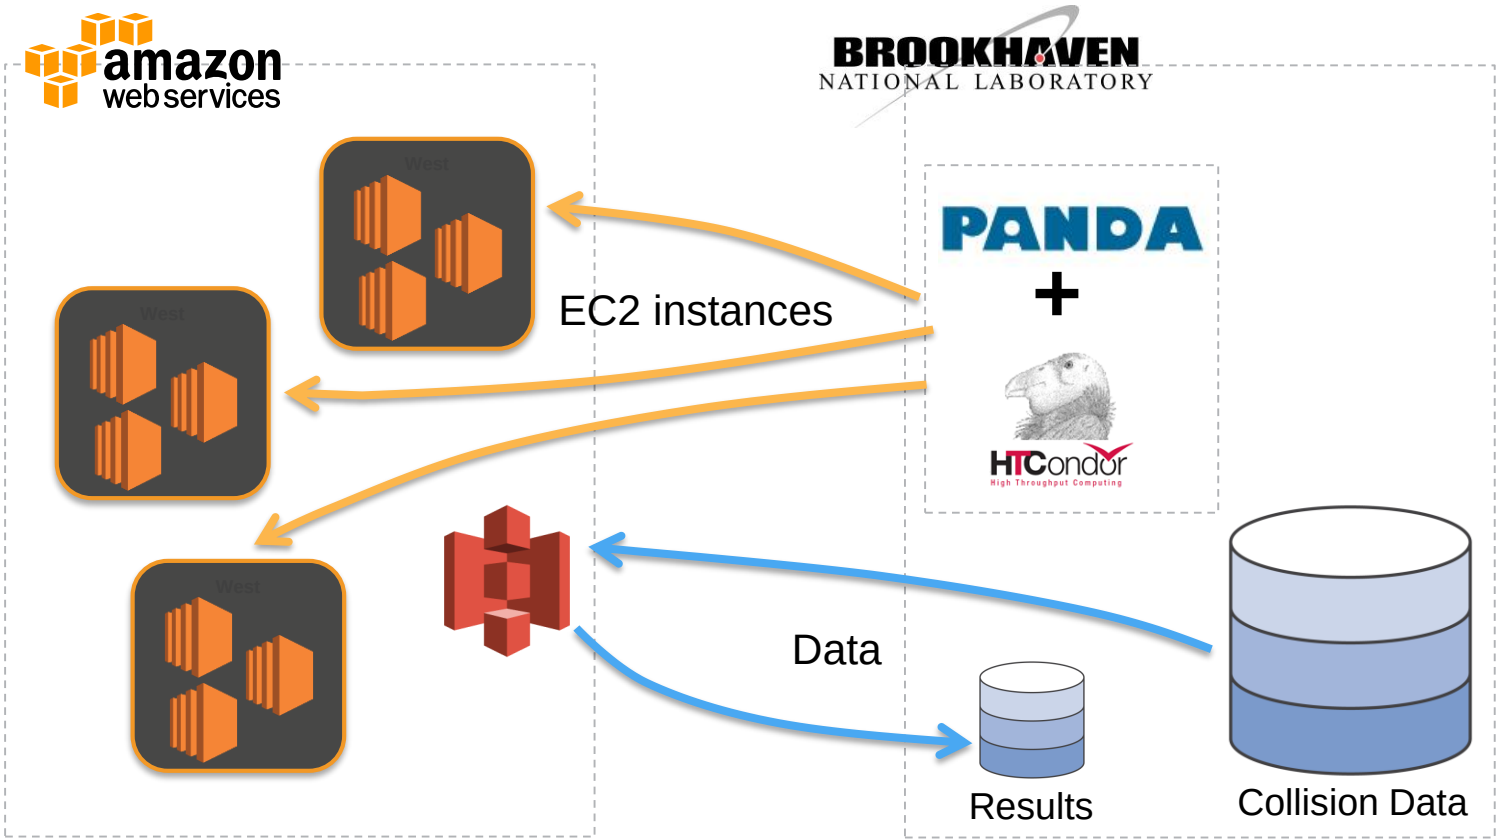

Measuring the Higgs with Machine Learning

Goal: identify particle spray (“jet”) originating from
Higgs Boson decaying into two bottom quarks (“double-b”)
Background: jets from other particles (quark/gluon)

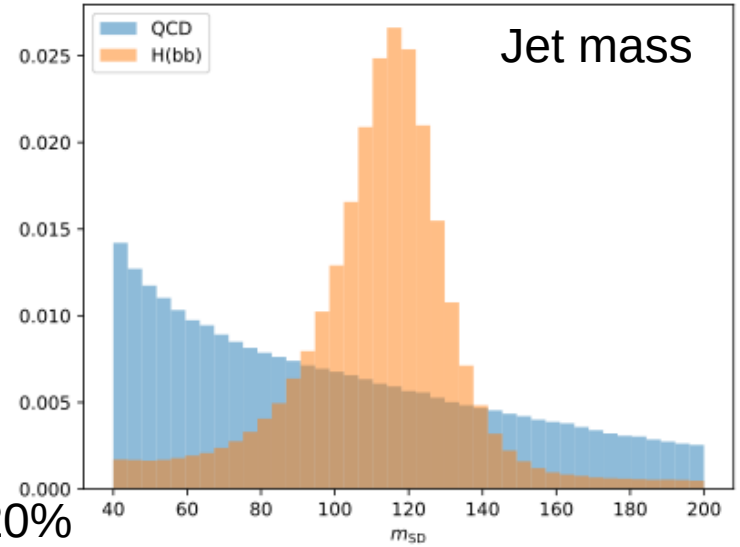

Data: 2.4 M q/g jets, 2.0 M H jets
Training/testing/validation split: 60% / 20% / 20%

Simply connected network: results

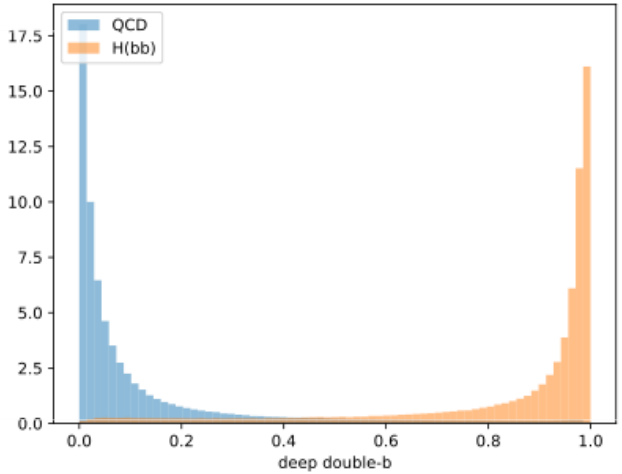

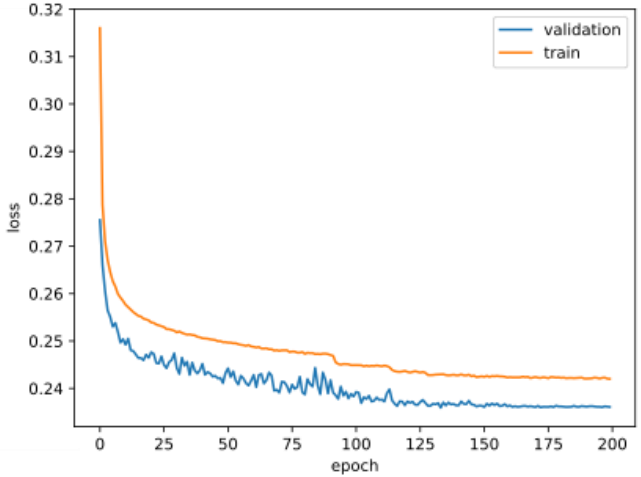

Training: Performance on AWS - big network

Significantly increase network complexity as an exercise

- Increased Conv1D filters from 32 to 1024
- Increased GRUs from 50 to 256
- Increased fully connected hidden layers from 100 to 512

Performance

- P2 instance (K80 GPU): 1735 min wallclock
- P3 instance (V100 GPU): 365 min wallclock (475% faster)

Caveat

- Have done very little performance tuning (using Keras as driver)
  - Unclear if we're making good use of Tensor Cores

4

AWS and the Research Community

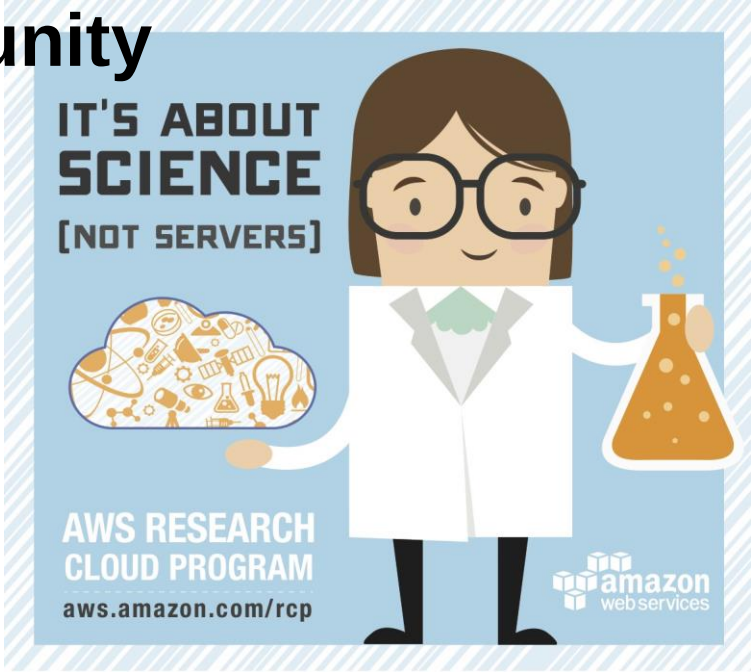

1. AWS Researcher's Handbook

The 200-page "missing manual" for science in the cloud.

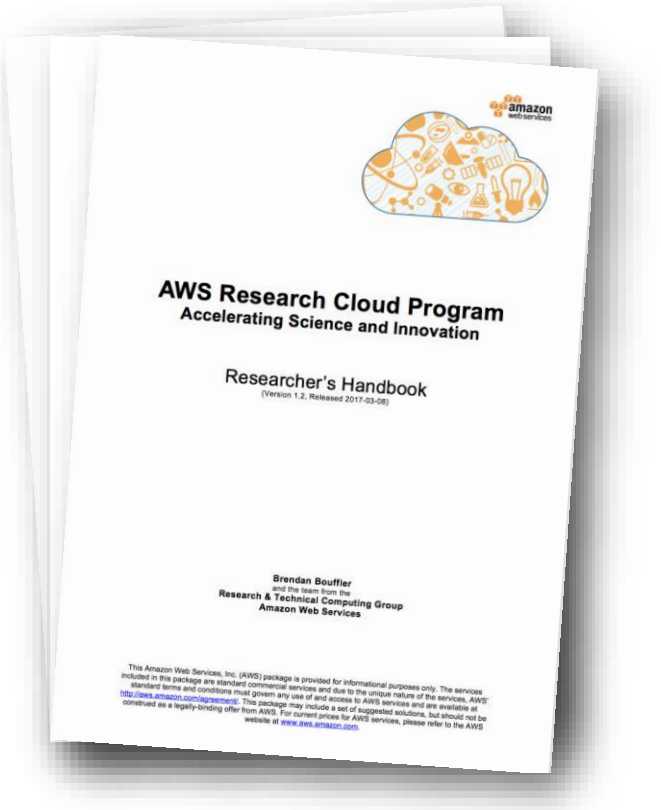

Written by Amazon's Research Computing community for scientists.

- Explains foundational concepts about how AWS can accelerate time-to-science in the cloud.
- Step-by-step best practices for securing your environment to ensure your research data is safe and your privacy is protected.
- Tools for budget management that will help you control your spending and limit costs (and preventing any over-runs).
- Catalogue of scientific solutions from partners chosen for their outstanding work with scientists.

aws.amazon.com/rcp

3. Serverless Computing: AWS Lambda

AWS Lambda is a service which allows for software functions in a variety of languages to be deployed into the cloud natively, and to be triggered directly or driven by events in the cloud. The infrastructure (hardware, operating system and software environment) for Lambda is managed by AWS and scales rapidly.

Bring your own code

Node.JS, Java, Python

Java = Any JVM based language such as Scala, Clojure, etc.

Bring your own libraries

Simple resource model

- Select memory from 128MB to 1.5GB in 64MB steps
- CPU & Network allocated proportionately to RAM
- Reports actual usage

1. AWS Researcher's Handbook

The 200-page "missing manual" for science in the cloud.

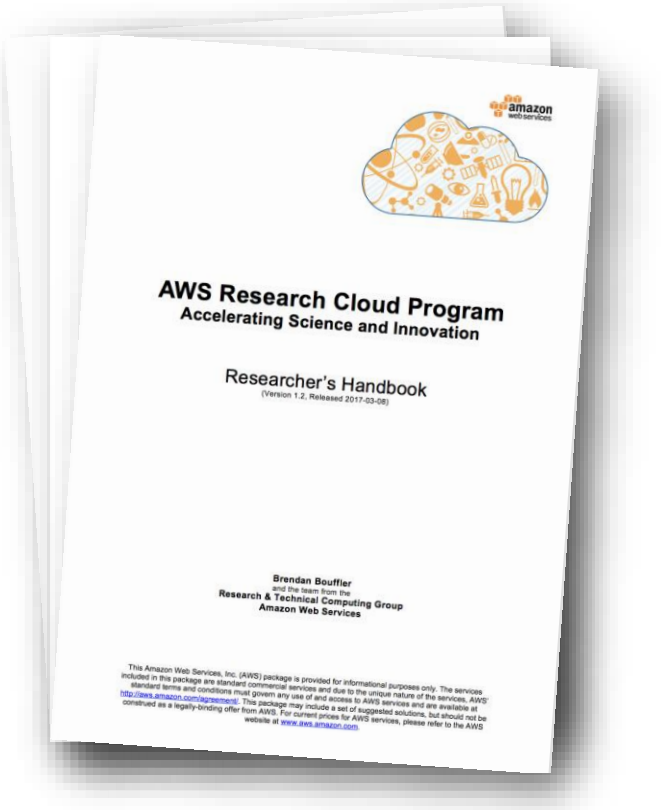

Written by Amazon's Research Computing community for scientists.

- Explains foundational concepts about how AWS can accelerate time-to-science in the cloud.
- Step-by-step best practices for securing your environment to ensure your research data is safe and your privacy is protected.
- Tools for budget management that will help you control your spending and limit costs (and preventing any over-runs).
- Catalogue of scientific solutions from partners chosen for their outstanding work with scientists.

aws.amazon.com/rcp

2. Global Data Egress Waiver reduces data download cost

All qualifying research customers should use this!

What?

Waives data egress charges from qualified accounts (capped at 15% of Total Spend)

Who?

Available to Degree-granting / Research Institutions

Why?

Researchers strongly need Predictable Budgets

How?

Join Research Cloud Program,

Or talk to your Account Team.

AWS peers with National Research and Education Networks for reliable, high-performance connection to/from AWS.
e.g. Esnet, Internet2, Geant, Jisc, SInet, AARnet, …

3. AWS Cloud Credits for Research

provide promotional AWS cloud credits for **anyone** to conduct research on AWS.

aws.amazon.com/research-credits

We especially like to support you in:

- building community tools
- proof of concept

4. Partnership with research funding agencies

- AWS initiated collaborative program with the National Science Foundation (NSF)
- The program provides NSF funds up to $26.5 million in addition to $3 million in AWS Cloud Credits to researchers to perform **cutting edge Big Data research on cloud for a period of 3-4 years**
- Precedent for similar collaborative programs with other agencies, and international research entities.
- RFP awardees selected by NSF per usual review process, i.e. cloud part of grant funding process

"In today's era of data-driven science and engineering, we are pleased to work with the AWS Research Initiative via the NSF BIGDATA program to provide cloud resources for our Nation's researchers to foster and accelerate discovery and innovation." -- **Dr. Jim Kurose,** Assistant Director of the National Science Foundation (NSF) for Computer and Information Science and Engineering Directorate (CISE)

5. AWS Educate & Academy

aws educate

- Self-service membership
- AWS usage credits
- Access to AWS Training content
- Curated content from AWS and educators
- Self-study learning paths and digital badges for students
- Job board for students

amazon web services | Academy

- Authorized ~60-hour curriculum developed & maintained by AWS
- Aligned to industry-recognized AWS Certifications
- Educator training and “instructor accreditation”
- Educator & Student discount for AWS Certification exam
- Free Practice Exams

Institutions, educators, and students benefit from both.

6. Billing: Budgets and Organizations

AWS Budgets:

- Track which project each expense belongs to
- Get notified when bill reaches a threshold
- Automatically shut down resources when limit exceeded
- See up-to-date spend anytime in web browser

AWS Organizations:

- Central management of multiple AWS accounts
- Control access policies and compliance
- Track costs & Control bill payment across accounts (& get volume discounts)
- Free

7. Academic Overhead – U of Washington empowers researchers to choose the best solution

Indirect Cost (F&A) Waiver for UW-IT Research Storage, Compute and Cloud Services

Effective April 1, 2015, UW-IT research storage, compute and contracted cloud services are no longer subject to Indirect Costs, also known as Facilities & Administrative (F&A) charges for sponsored research expenditures.

This waiver applies to the following UW-IT services:

Contracted Cloud Services

- Microsoft Azure
- Amazon Web Services

Benefits

Extending the F&A waiver to these UW-IT services allows Principal Investigators (PIs) in a sponsored research program to choose the most appropriate solution to meet their computing needs, whether it be research storage, compute or contracted cloud services through UW-IT, or

U. Washington also created a very active “eScience Institute” that supports campus researchers/educators with Cloud adoption and other IT needs.

Thank You

jorissen@amazon.com

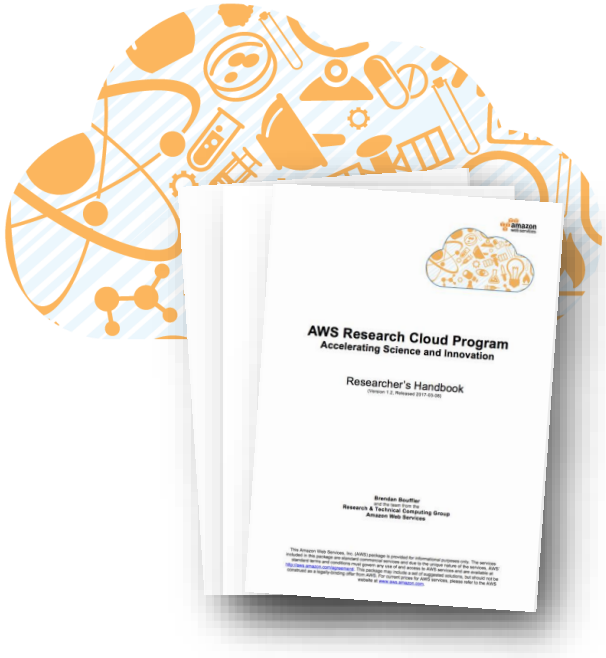

Register and enroll in the
AWS Research Cloud Program

https://aws.amazon.com/rcp