Amazon Web Services with HTCondor

HTCondor Week 2017, Madison, Wisconsin

Sanjay Padhi, Ph.D

AWS Research and Technical Computing

sanpadhi@amazon.com

Today

- Amazon Web Services (AWS)
- Scientific Computing using AWS and HTCondor
- Machine Learning

AWS Global Infrastructure

16 Regions – **42** Availability Zones – **74** Edge Locations

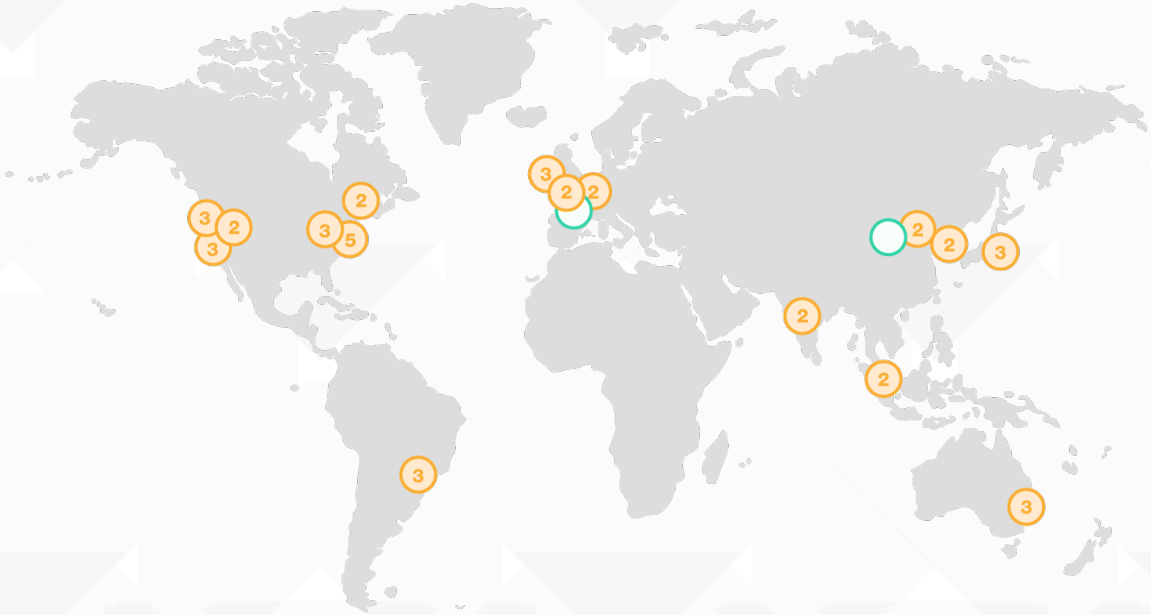

Region & Number of Availability Zones

AWS GovCloud (2)

US West

Oregon (3)

Northern California (3)

US East

N. Virginia (5), Ohio (3)

Canada

Central (2)

South America

São Paulo (3)

EU

Ireland (3)

Frankfurt (2)

London (2)

Asia Pacific

Singapore (2)

Sydney (2), Tokyo (3),

Seoul (2), Mumbai (2)

China

Beijing (2)

Announced Regions

Paris, Ningxia

amazon web services

Zoom In: AWS Region

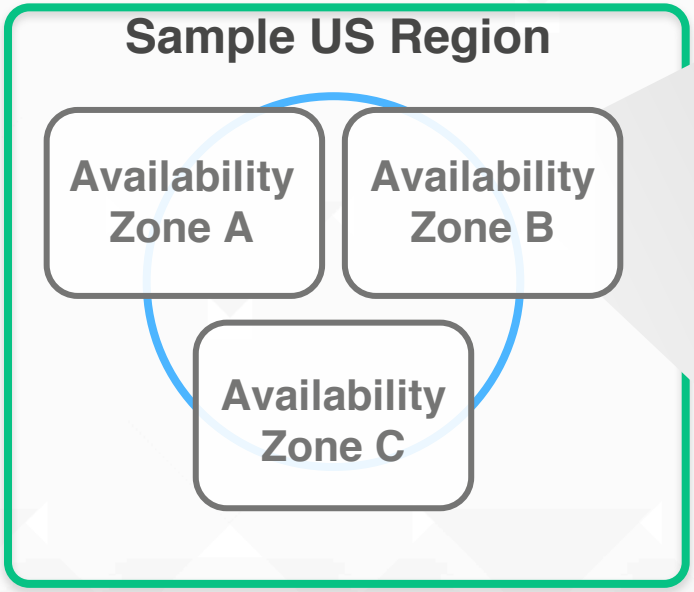

Zoom In: AWS AZ

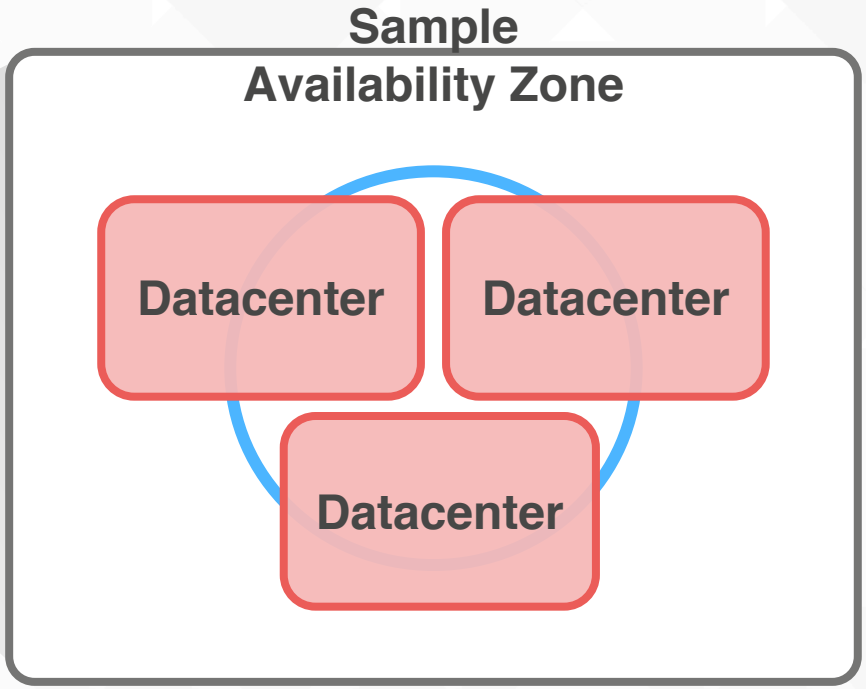

AWS: Comprehensive Infrastructure Services Platform with 90+ services

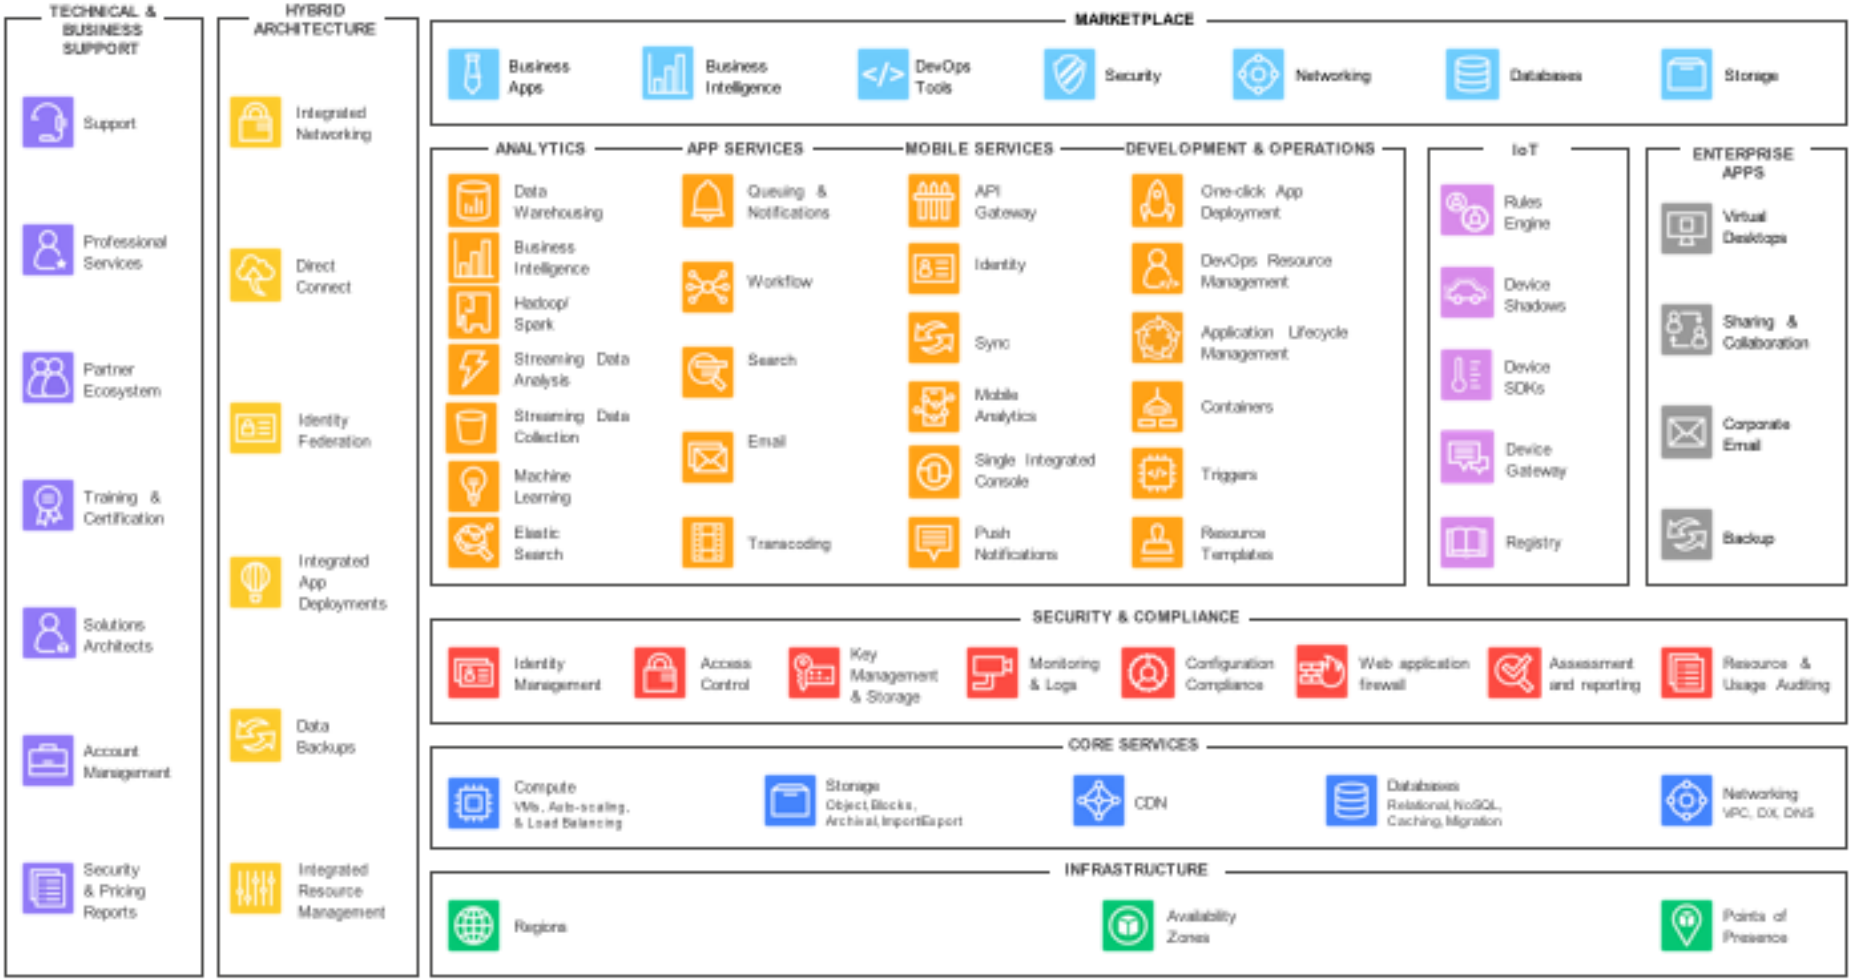

AWS GovCloud (US)

Designed to host sensitive data and regulated workloads in the cloud

- Supports U.S. government compliance requirements, including ITAR and FedRAMP
- Operated by employees who are vetted "U.S. Persons"
- Root account holders are confirmed U.S. Persons
- Available to U.S. government agencies and organizations in government-regulated industries, that meet GovCloud (US) requirements for access

amazon web services

Architected for Government Security Requirements

ISO 9001
Global Quality Standard

ISO 27001
Security Management Standard

ISO 27017
Cloud Specific Controls

ISO 27018
Personal Data Protection

PCI DSS Level 1
Payment Card Standards

SOC 1
Audit Controls Report

SOC 2
Compliance Controls Report

SOC 3
General Controls Report

CJIS
Criminal Justice Information Services

DoD SRG
DoD Data Processing

FDA
Food and Drug Administration

FedRAMP
Government Data Standards

FedRAMP TIC
FedRAMP-Trusted Internet Connection

FERPA
Educational Privacy Act

FIPS
Government Security Standards

FISMA
Federal Information Security Management

GXP
Quality Guidelines and Regulations

HIPAA
Protected Health Information

SEC Rule 17a-4(f)
Financial Data Standards

ITAR
International Arms Regulations

FISC [Japan]
Financial Industry Information Systems

IRAP [Australia]
Australian Security Standards

MLPS Level 3 [China]
Multi-Level Protection

MTCS Tier 3 [Singapore]
Multi-Tier Cloud Security Standard

My Number Act [Japan]
Personal Information Protection

DNB [Netherlands]
Dutch Financial Regulations

EU Data Protection
Data Protection Framework

G-Cloud [UK]
UK Government Standards

IT-Grundschutz [Germany]
Baseline Protection Methodology

Privacy Shield
EU-US Data Transfer

UK Cyber Essentials Plus
Cyber Threat Protection

https://aws.amazon.com/compliance/

amazon web services

AWS: Pace of Innovation

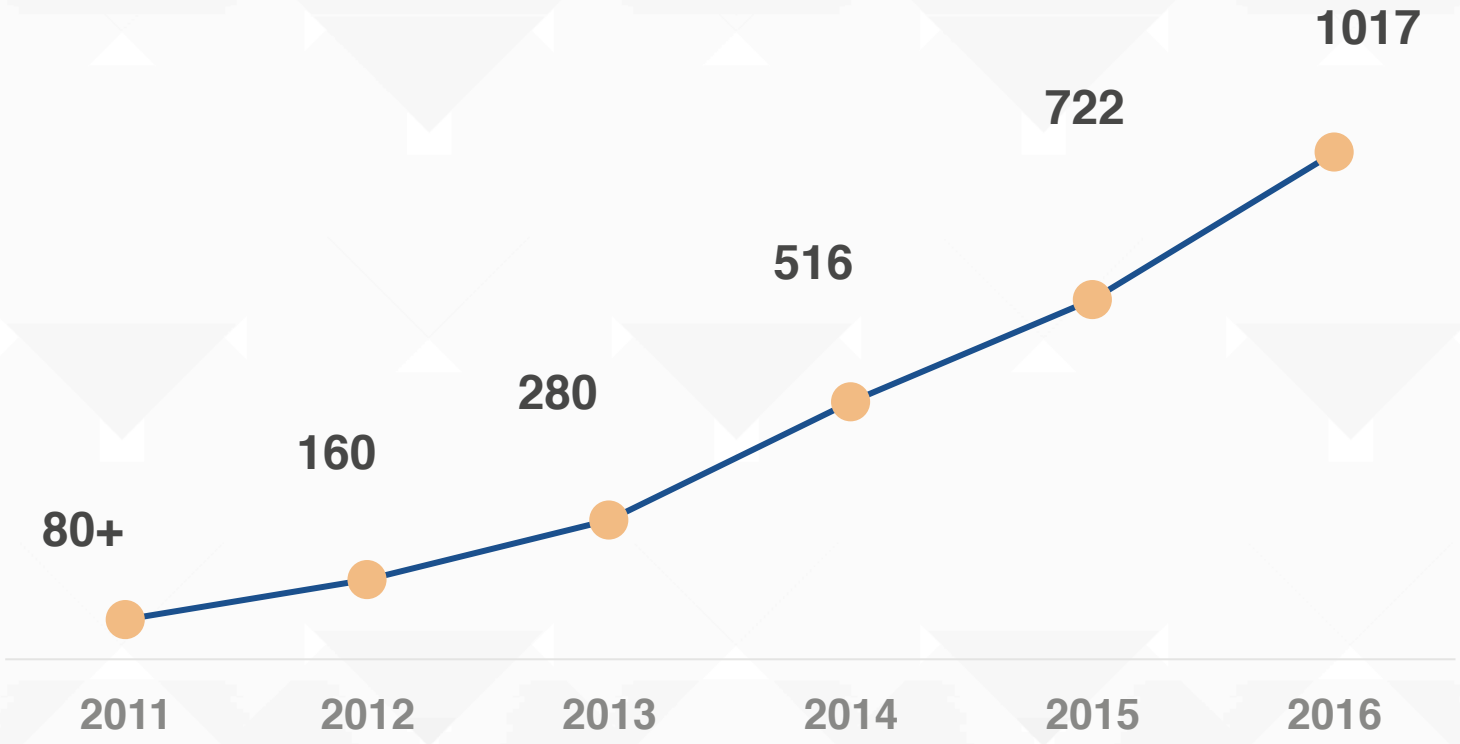

amazon web services

AWS in the Public Sector

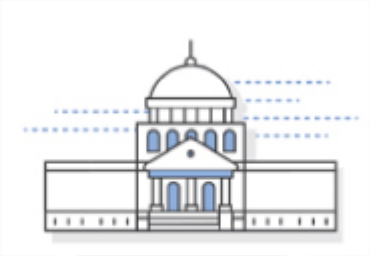

2,300+

government agencies

7,000+

educational institutions

22,000+

nonprofit organizations

amazon web services

Compute Services

Virtual Server Hosting, Container management, and Serverless Computing

Amazon EC2

Provides resizable cloud-based compute capacity in the form of EC2 instances, which are equivalent to virtual servers

Amazon EC2 Container Service

A highly scalable, high performance container management service

AWS Lambda

Run code without thinking about servers.

bankinter

airbnb NETFLIX

NASA Jet Propulsion Laboratory
California Institute of Technology

Time Inc. LAMBORGHINI

amazon web services

Scientific Computing using AWS and HTCondor

Large Hadron Collider

The Large Hadron Collider @ CERN includes 6,000+ researchers from over 40 countries and produces approximately 25PB of data each year.

The ATLAS and CMS experiments are using AWS for Monte Carlo simulations, processing, and analysis of LHC data.

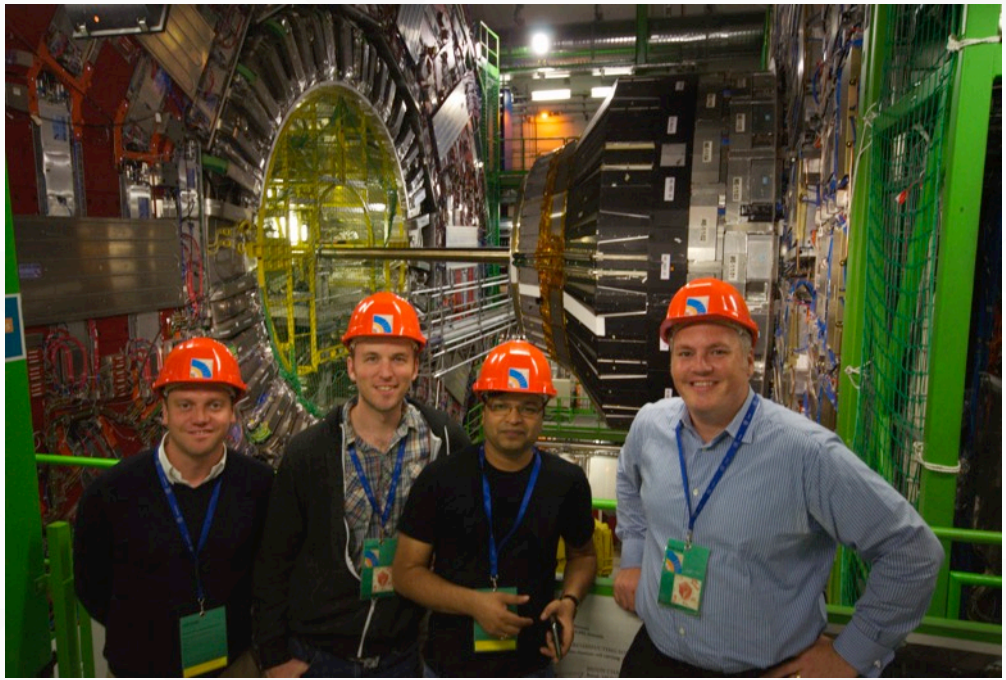

amazon web services

CMS Detector

- **80 Million electronic channels**
  - x 4 bytes
  - x 40MHz
  - ---
  - ~ **10 Petabytes/sec** of information
  - x 1/1000 zero-suppression
  - x 1/100,000 online event filtering
  - ---
  - ~ 100-1000 Megabytes/sec raw data to tape
  - **1 to 10 Petabytes of raw data per year written to tape, not counting simulations.**

- 2000 Scientists (1200 Ph.D. in physics)
  - ~ 180 Institutions
  - ~ 40 countries

- 12,500 tons, 21m long, 16m diameter

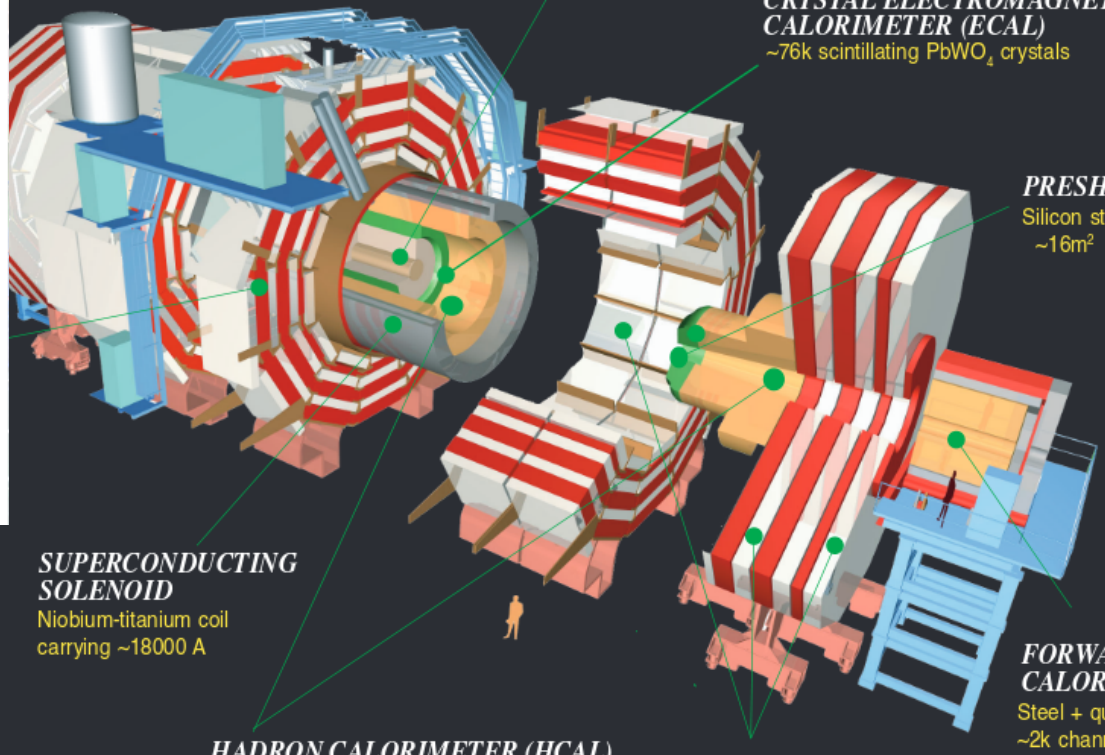

Total weight	: 14000 tonnes
Overall diameter	: 15.0 m
Overall length	: 28.7 m
Magnetic field	: 3.8 T

Clouds provided elasticity in computing

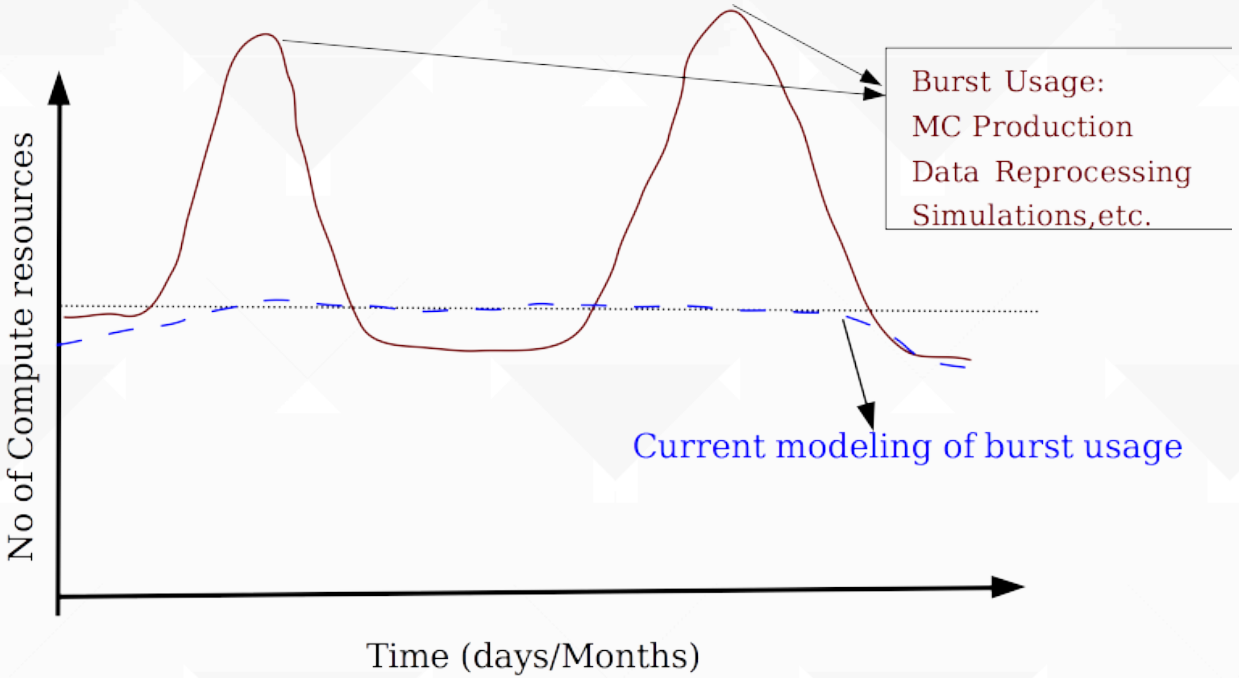

Finite number of resources by the experiments (Compute as well as Storage)

"Burst" usage is modeled using delays (~months) due to (re)processing capabilities

Elasticity in the system is really essential

amazon web services

Tutorial for Scientific Computing using Amazon

Monday 2 Feb 2015, 08:00 → 18:00 Europe/Zurich

30-7-018 - Kjell Johnsen Auditorium (CERN)

Maria Girone (CERN) , Sanjay Padhi (Univ. of California San Diego (US))

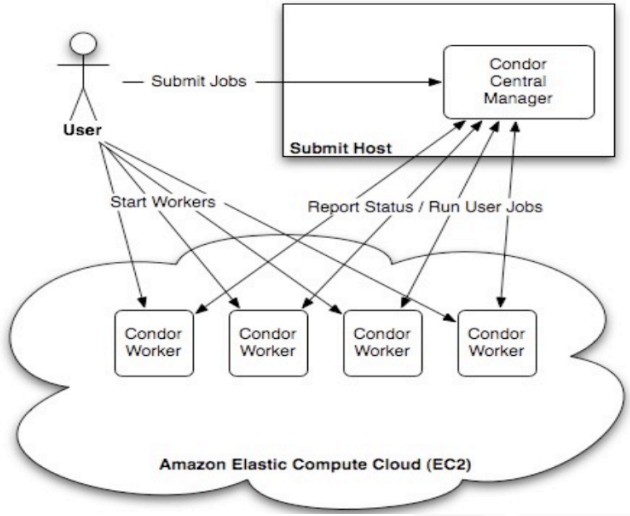


```
Task name:                    150131_105339_crab3test99:mmascher_crab_tutorial_MC_analysis_test1
Task status:                  SUBMITTED
Glidemon monitoring URL:      http://glidemon.web.cern.ch/glidemon/jobs.php?taskname=150131_105339_crab
Dashboard monitoring URL:     http://dashb-cms-job.cern.ch/dashboard/templates/task-analysis/#user=mmas
50131_105339_crab3test99%3Ammascher_crab_tutorial_MC_analysis_test1
Details:                      finished      5.6% ( 1/18)
                              idle         88.9% (16/18)
                              running       5.6% ( 1/18)

Publication status:           finished      5.6% ( 1/18)
                              unsubmitted  94.4% (17/18)

Output dataset:               /GenericTTbar/mmascher-CRAB3_tutorial_MC_analysis_test1-37773c17ce2994cf16892d5f04945e41/USER
Output dataset url:           https://cmsweb.cern.ch/das/request?input=%2FGenericTTbar%2Fmmascher-CRAB3_tutorial_MC_analysis_tes
nce=prod%2Fphys03
```

Start » [mmascher] » Tasks » Jobs

Data | Charts | Show 25 entries | Task: 150131_105339_crab3test99:mmascher_crab_tutorial_MC_analysis_test1 NJobTotal: 18 Pending: 16 Running: 1 Unknown: 0 Cancelled: 0 Success: 1 Failed: 0 WNPostProc: 0 ToRetry: 0

Id	Status	AppExitCode	Site	Retries	Submitted	Started	Finished	Wall Time	Job Log	File Access	FTS File Status
1	finished	0	unknown	1	2015-01-31T10:54:26	2015-01-31T10:55:41	2015-01-31T13:02:04	02:06:23	Job Log,Job Log JSON,Post Job Log	File Info	N/A
Attempt No.	Restarts No.	Error Code/ Details	Job Status	Site	Submitted	Started	Finished	Wall Time	Job Log		
1	1	0 / Application finished properly Postprocessing step finished properly	finished	unknown	2015-01-31T10:54:26	2015-01-31T10:55:41	2015-01-31T13:02:04	02:06:23	Job Log,Job Log JSON,Post Job Log		
2	running	N/A	unknown	1	2015-01-31T10:54:26	2015-01-31T13:06:38	1970-01-01T00:00:00	00:00:00	Not available	File Info	N/A
3	pending	N/A	unknown	1	2015-01-31T10:54:26	1970-01-01T00:00:00	1970-01-01T00:00:00	00:00:00	Not available	File Info	N/A
4	pending	N/A	unknown	1	2015-01-31T10:54:26	1970-01-01T00:00:00	1970-01-01T00:00:00	00:00:00	Not available	File Info	N/A
5	pending	N/A	unknown	1	2015-01-31T10:54:26	1970-01-01T00:00:00	1970-01-01T00:00:00	00:00:00	Not available	File Info	N/A
6	pending	N/A	unknown	1	2015-01-31T10:54:26	1970-01-01T00:00:00	1970-01-01T00:00:00	00:00:00	Not available	File Info	N/A
7	pending	N/A	unknown	1	2015-01-31T10:54:26	1970-01-01T00:00:00	1970-01-01T00:00:00	00:00:00	Not available	File Info	N/A

amazon web services

Streaming data interaction every 25 nano sec – Occupancy (Finding patterns)

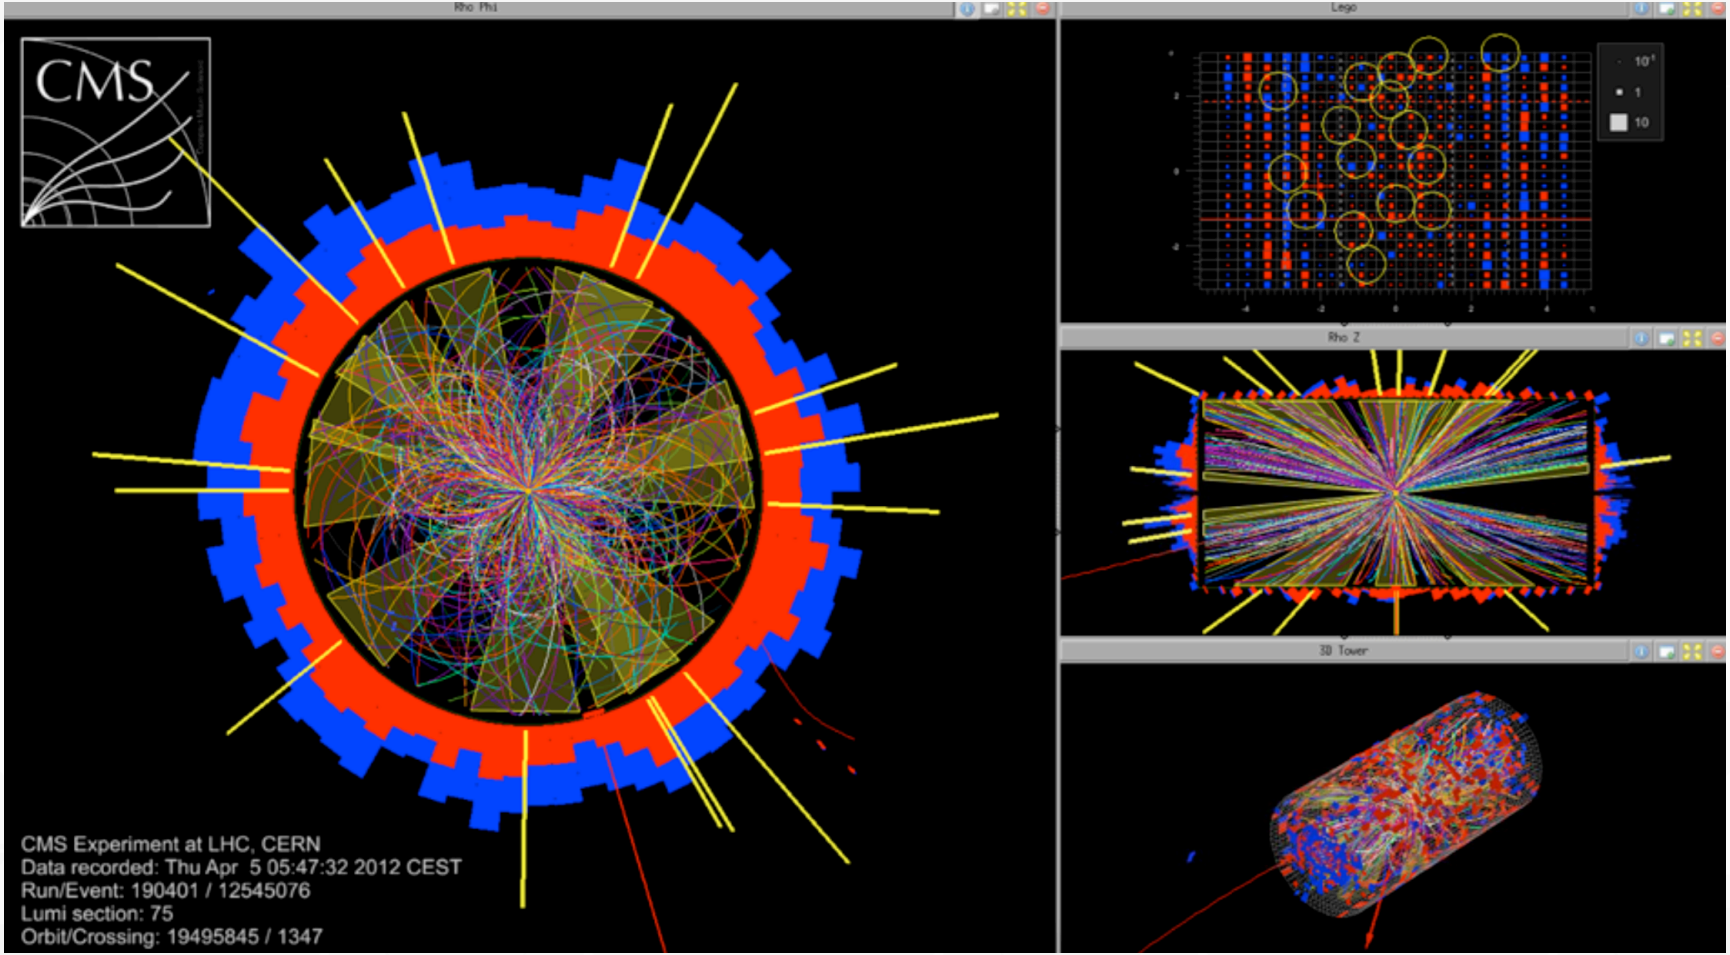

HepCloud : CMS Workflow at Fermilab (Auto-expansion to AWS using HTCondor)

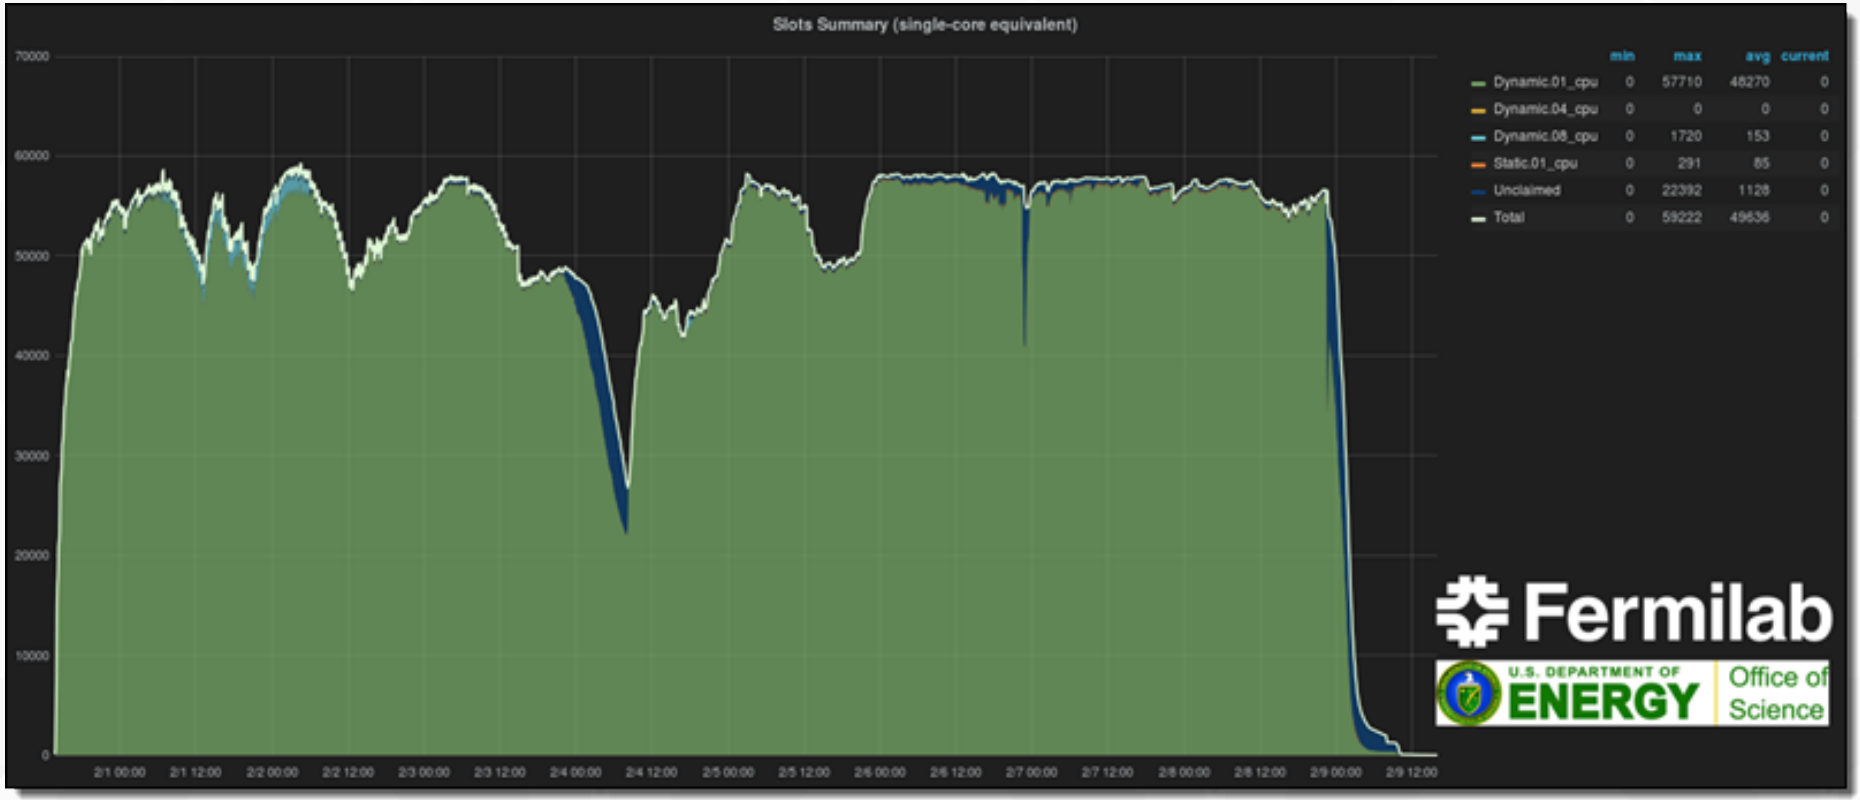

~60,000 slots using AWS spot instances. A factor of 5 larger than Fermilab capacity!

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

amazon web services

Results from the CMS Use Case

- **All CMS requests fulfilled for the "Moriond" conference**
  - **2.9 million jobs, 15.1 million wall hours**
    - 9.5% badput – includes preemption from spot pricing
    - 87% CPU efficiency
  - **518 million events generated**

/DYJetsToLL_M-50_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext4-v1/AODSIM
/DYJetsToLL_M-10to50_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext3-v1/AODSIM
/TTJets_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext1-v1/AODSIM
/WJetsToLNu_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext4-v1/AODSIM

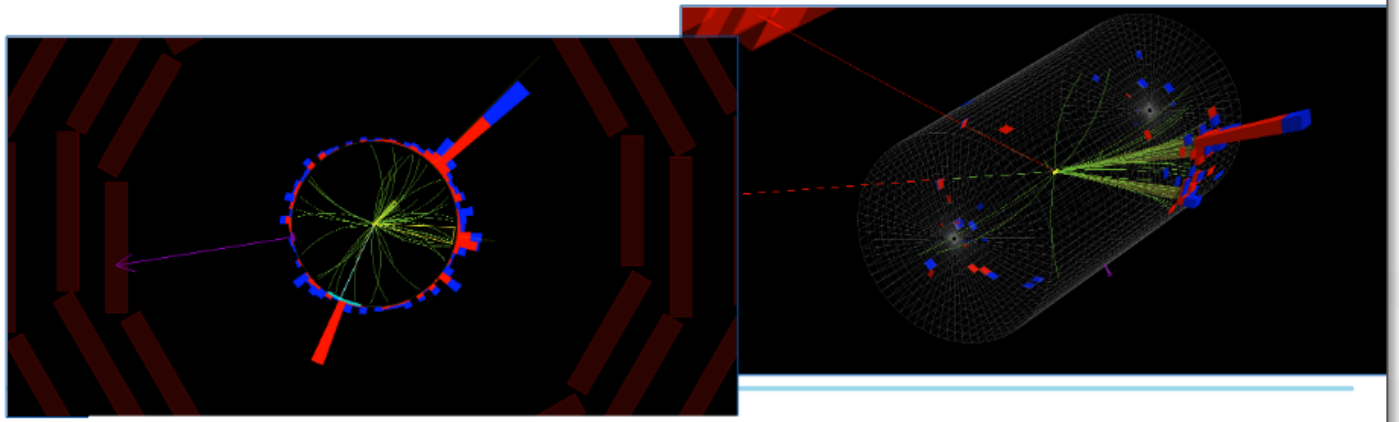

ATLAS Workflow in Cloud – At Scale and Low Cost

- Joint project between AWS, BNL and ESNET
- Investigate the technical and financial feasibility of large-scale usage of Cloud
- AWS: Provided expertise & guidance
- BNL: ATLAS compatible VMs, provisioning infrastructure, VM life management
- ESNET: High performance connectivity between AWS and US site

M. Ernst: Director of the RHIC and ATLAS Computing Facility, Brookhaven National Laboratory

"ATLAS has met the challenge of data intensive computing at a scale not seen before"

"The joint project with the AWS Scientific Computing Team and ESnet has been crucial to the successful implementation"

"The cost of AWS/EC2 spot is slightly lower than dedicated farm resources at BNL"

www.informationweek.com/cloud/infrastructure-as-a-service/brookhaven-lab-finds-aws-spot-instances-hit-sweet-spot/d/d-id/1324145

amazon web services

Enabling Global Collaboration

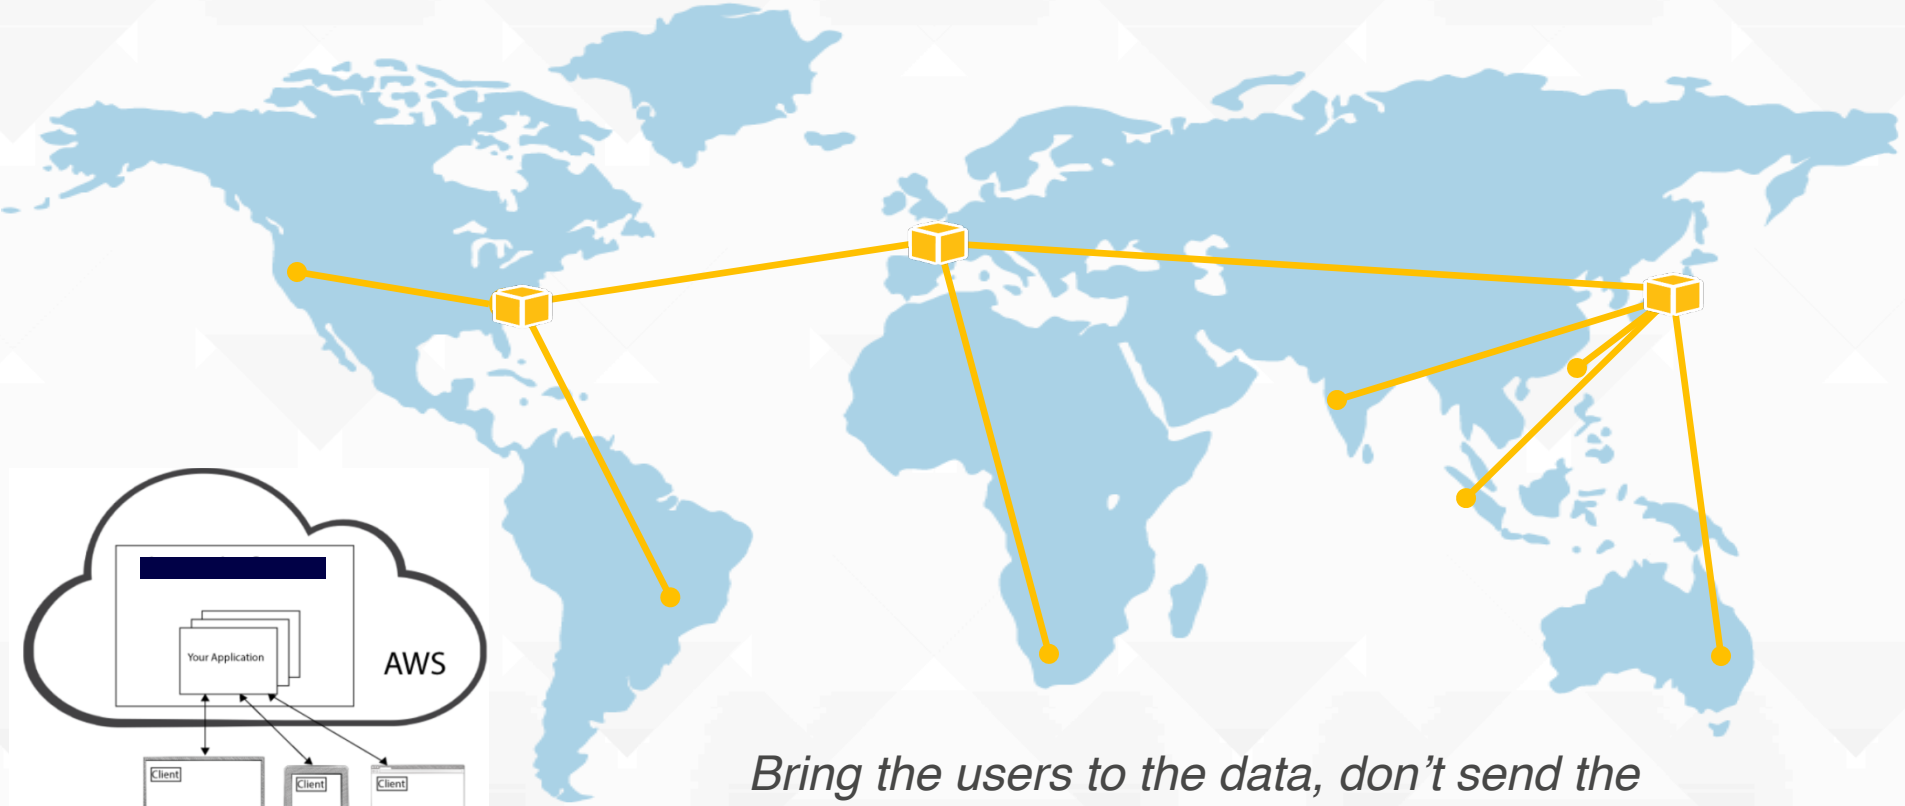

Bring the users to the data, don't send the data to the users

amazon web services

NOvA uses AWS to Shed Light on Neutrino Mysteries

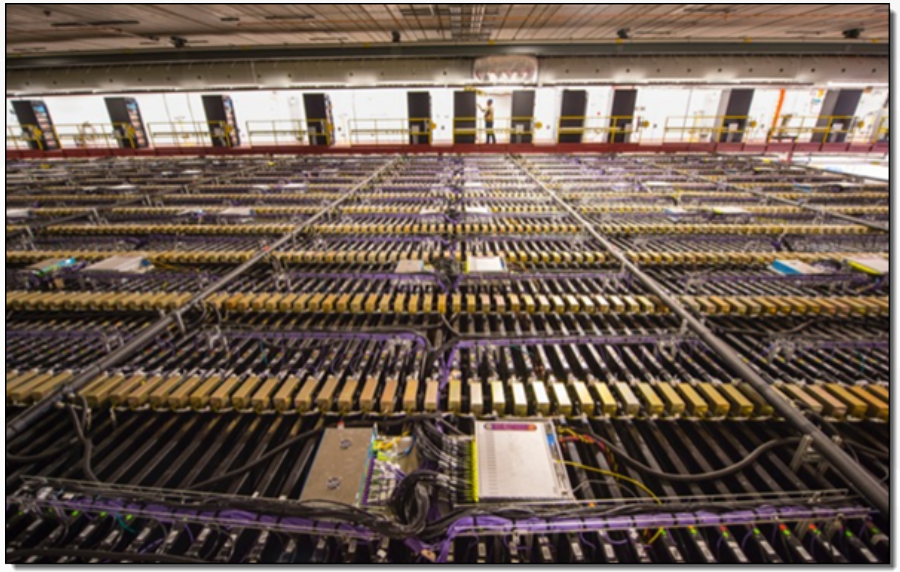

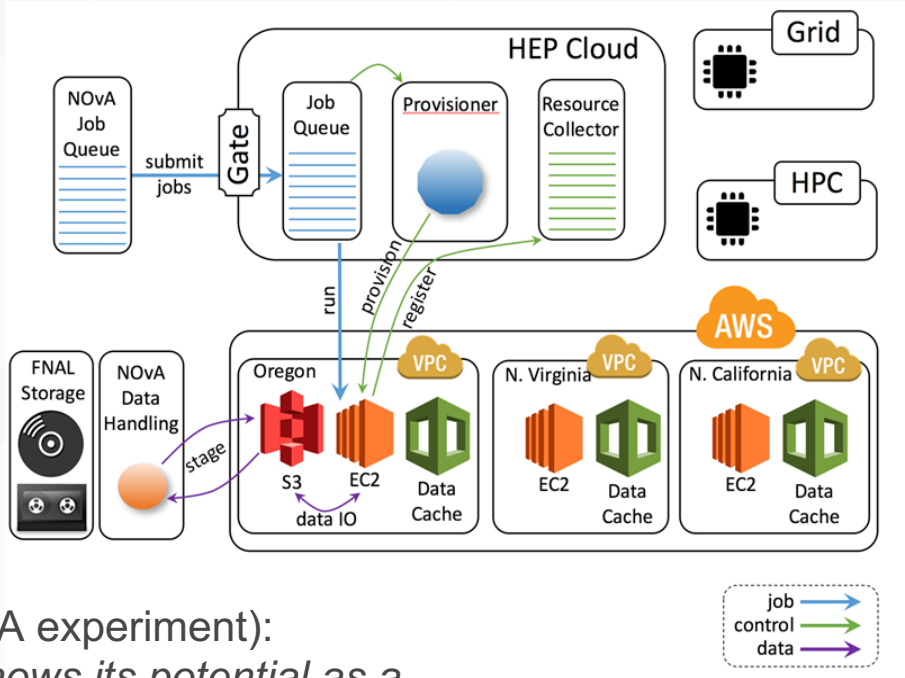

Peter Shanahan (Co-spokesperson of the NOvA experiment):
"Our experience with Amazon Web Services shows its potential as a reliable way to meet our peak data processing needs at times of high demand"

https://aws.amazon.com/blogs/aws/nova-uses-aws-to-shed-light-on-neutrino-mysteries/

Neutrinos are ghost like particles → Needed advanced ML analytics to detect

amazon web services

Scalability using AWS

CYCLECOMPUTING
LEADER IN CONDOR GRID COMPUTING SOLUTIONS

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

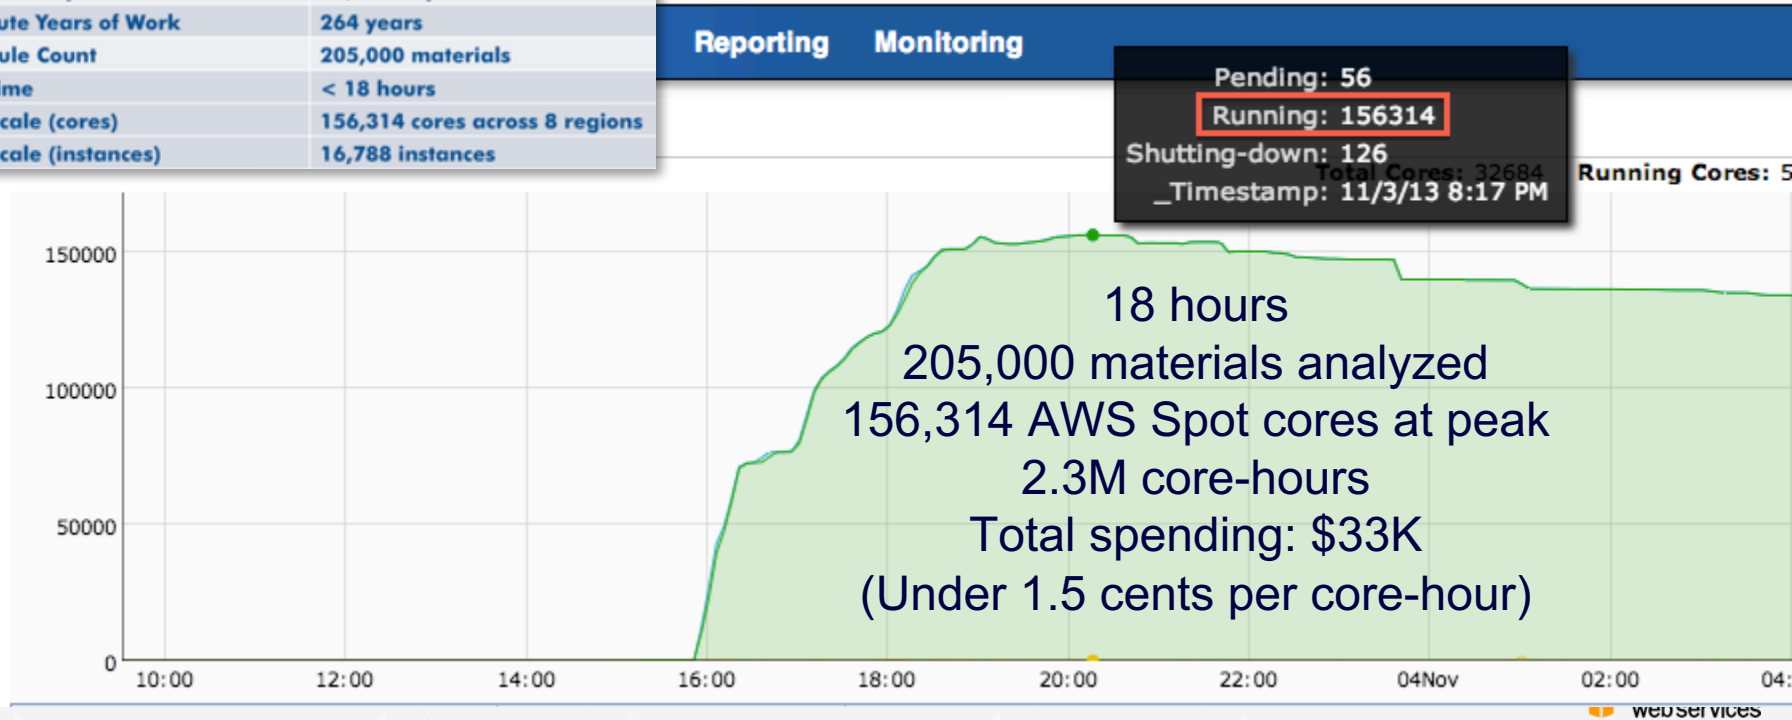

Development of HTCondor Annex

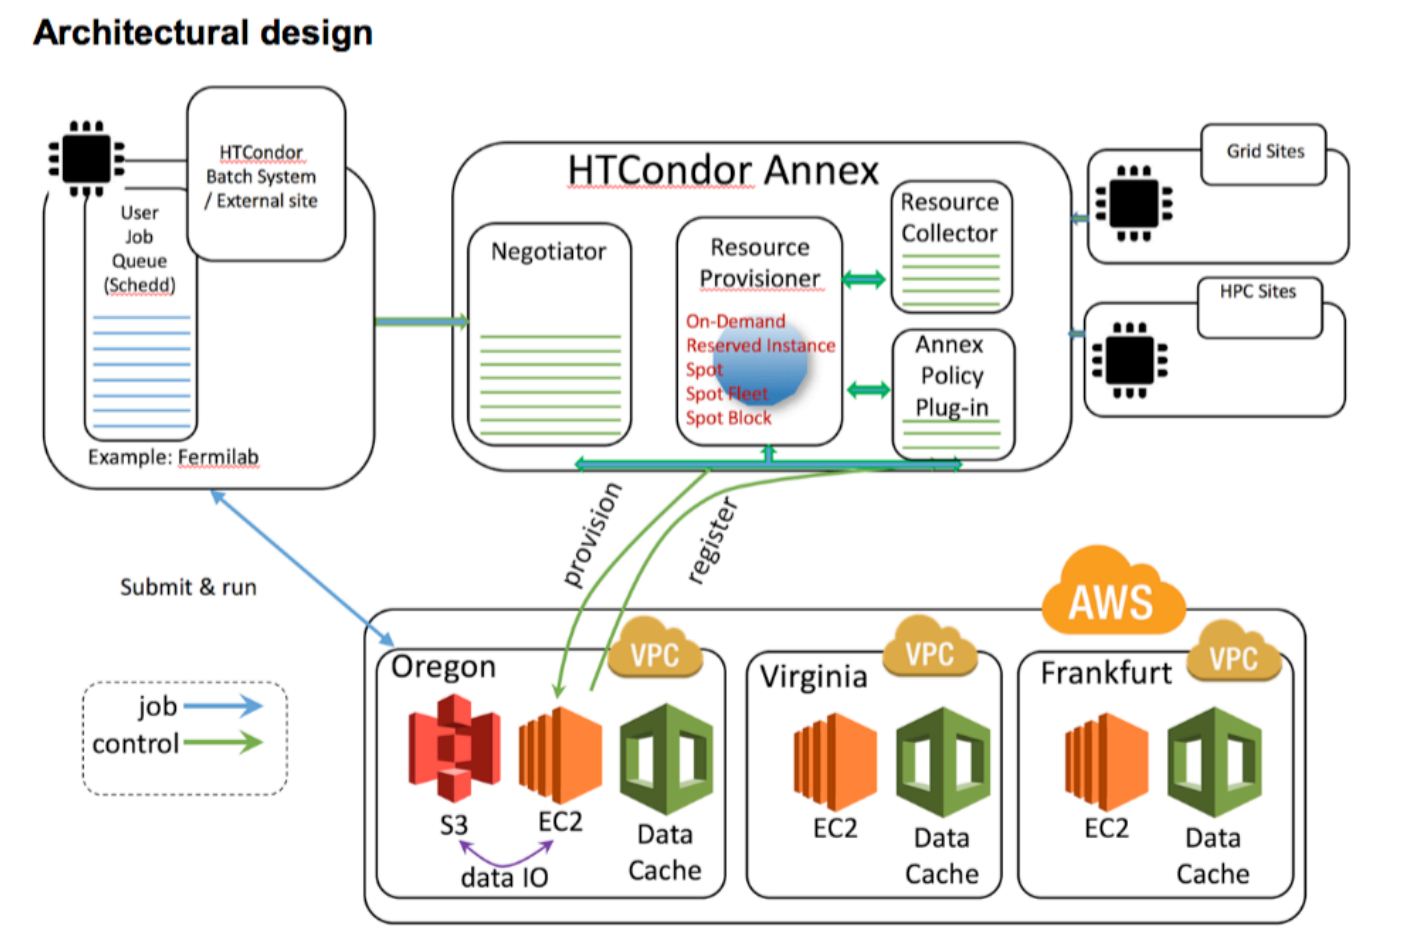

Machine Learning

Amazon AI: https://aws.amazon.com/amazon-ai/

Amazon Machine Learning: https://aws.amazon.com/machine-learning/

AWS Deep Learning AMI: https://aws.amazon.com/blogs/ai/the-aws-deep-learning-ami-now-with-ubuntu/

ARTIFICIAL INTELLIGENCE

Early artificial intelligence stirs excitement.

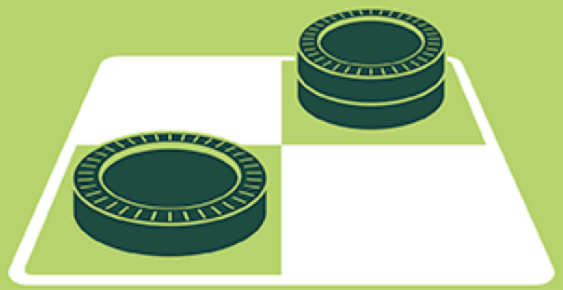

MACHINE LEARNING

Machine learning begins to flourish.

DEEP LEARNING

Deep learning breakthroughs drive AI boom.

1950's 1960's 1970's 1980's 1990's 2000's 2010's

Machine Learning

	Supervised Learning	Unsupervised Learning
Discrete	classification or categorization	clustering
Continuous	regression	dimensionality reduction

Supervised Learning:

- Learning from “labelled data”
- Classification, Regression, Prediction, Function Approx

Unsupervised Learning:

- Method to find similar groups in the data clusters
- Groups that are similar to near clusters
- Groups different far away from each other

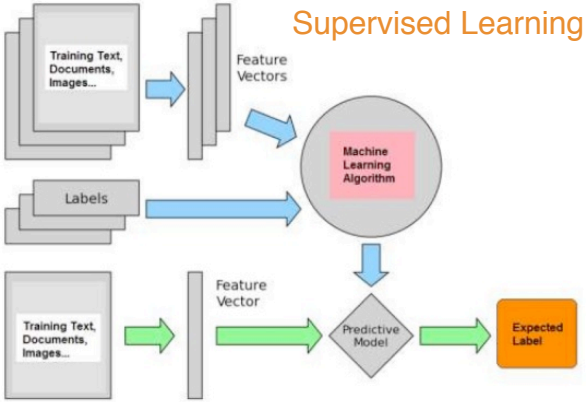

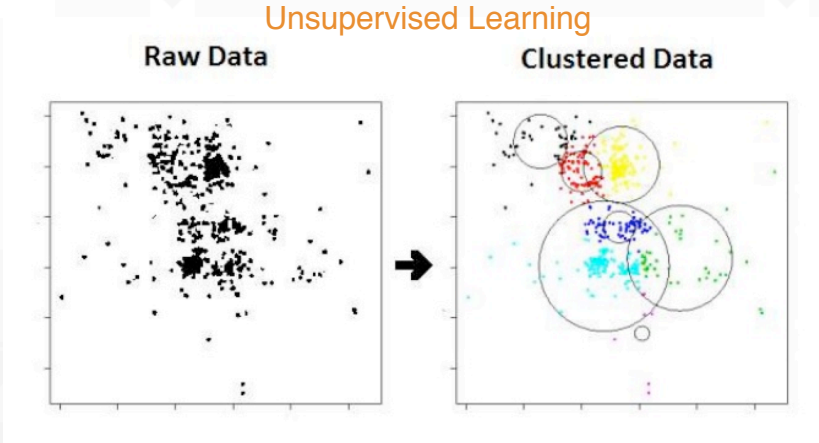

amazon web services

Machine Learning (Classification, Regression and Ranking) :

Euclidean Distance Score

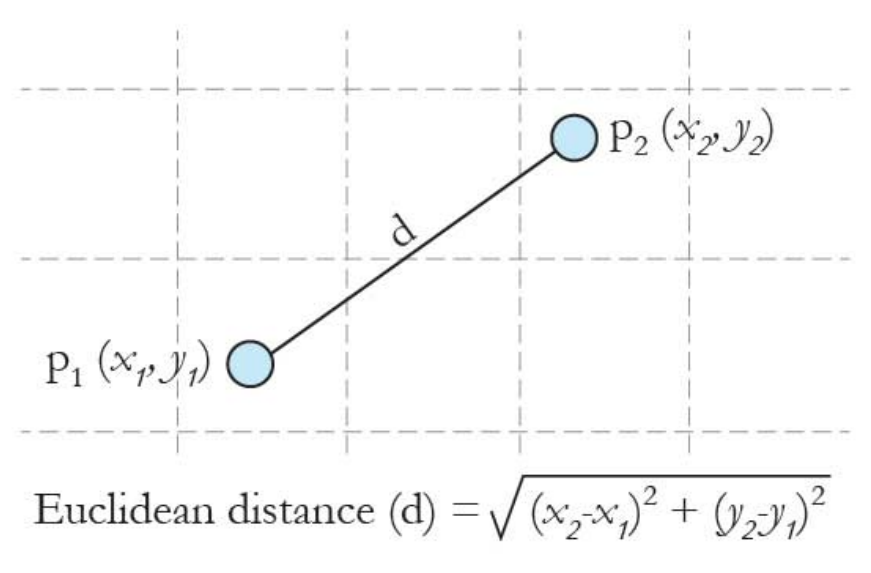

Pearson Correlation Score

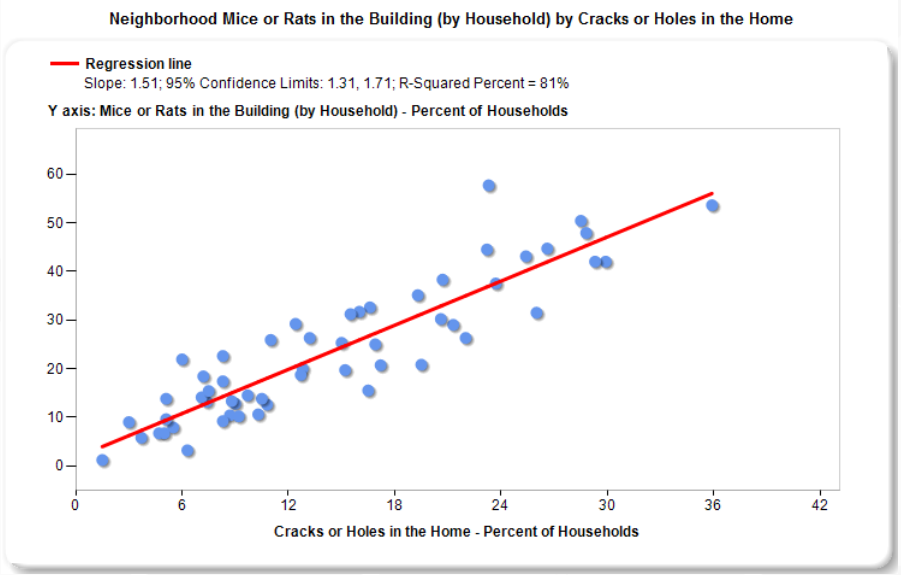

amazon web services

Convolutional Neural Net

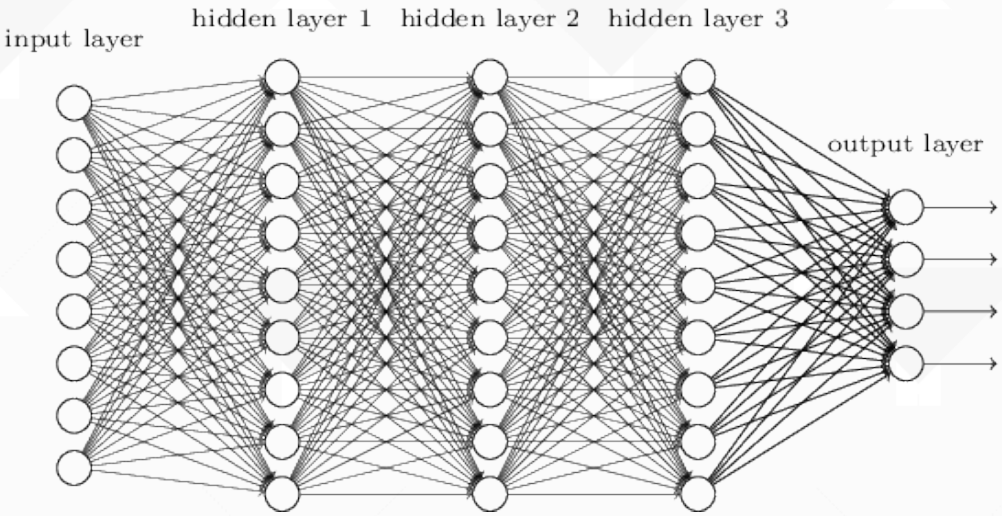

- Convolutional neural nets are a very successful deep learning method.
- Inspired by research showing that the cells in the visual cortex are only responsive to small portions of the visual field - “receptive field”.
- Some cells collect information from small patches – sensitive to edge-like features.
- Other cells collect information from large patches.
- Effectively, these cells are applying convolutional kernels across the visual field.

Amazon Rekognition - Image Detection and Recognition Powered by Deep Learning

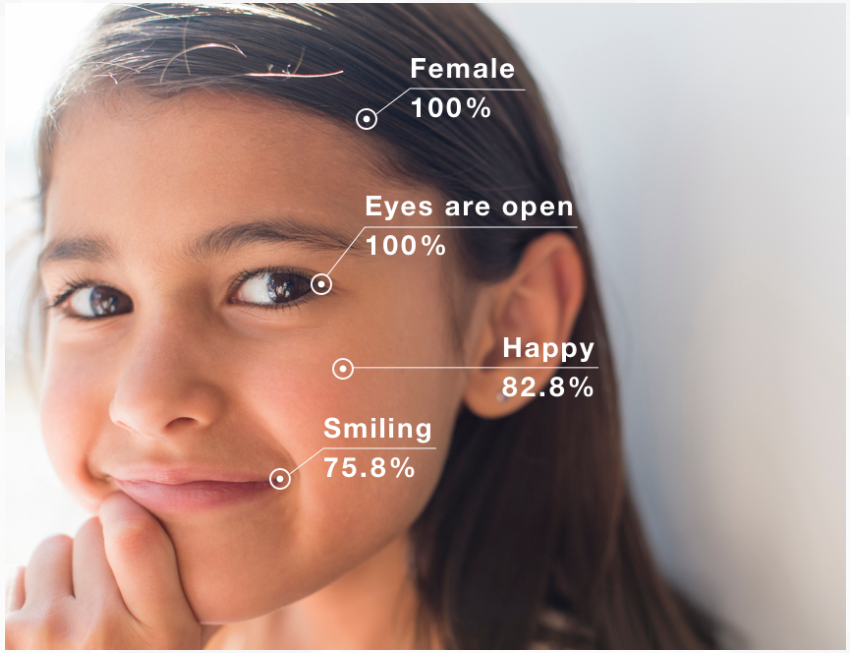

https://aws.amazon.com/rekognition/

amazon web services

Amazon Rekognition - Image Detection and Recognition Powered by Deep Learning

DOE Workshop on Future Online Analysis Platform. April 2017

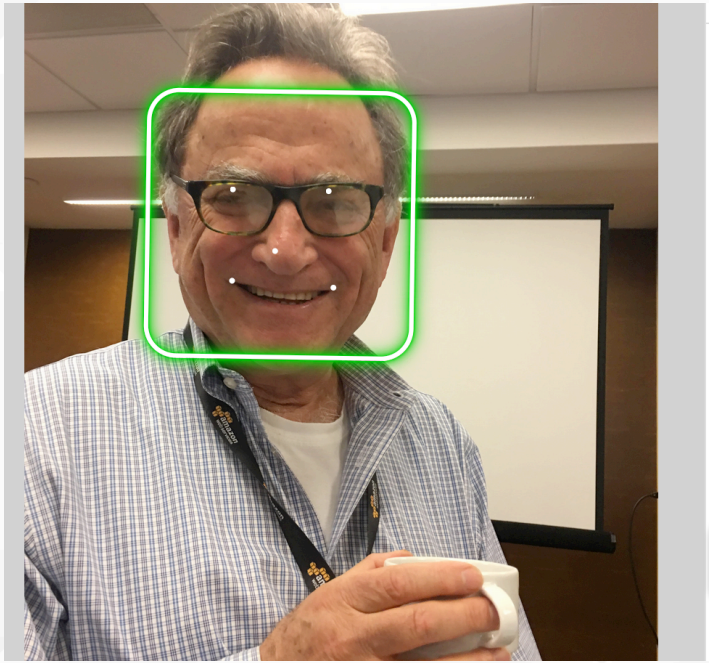

Results

looks like a face	99.9%
appears to be male	99.9%
age range	60 - 80 years old
smiling	99.7%
appears to be happy	80%
wearing eyeglasses	99.9%

Results

looks like a face	100%
appears to be male	99.9%
age range	45 - 63 years old
smiling	99.1%
appears to be happy	97.1%
wearing eyeglasses	99.9%

https://aws.amazon.com/rekognition/

amazon web services

Automatic Grading of Diabetic Retinopathy through Deep Learning using AWS

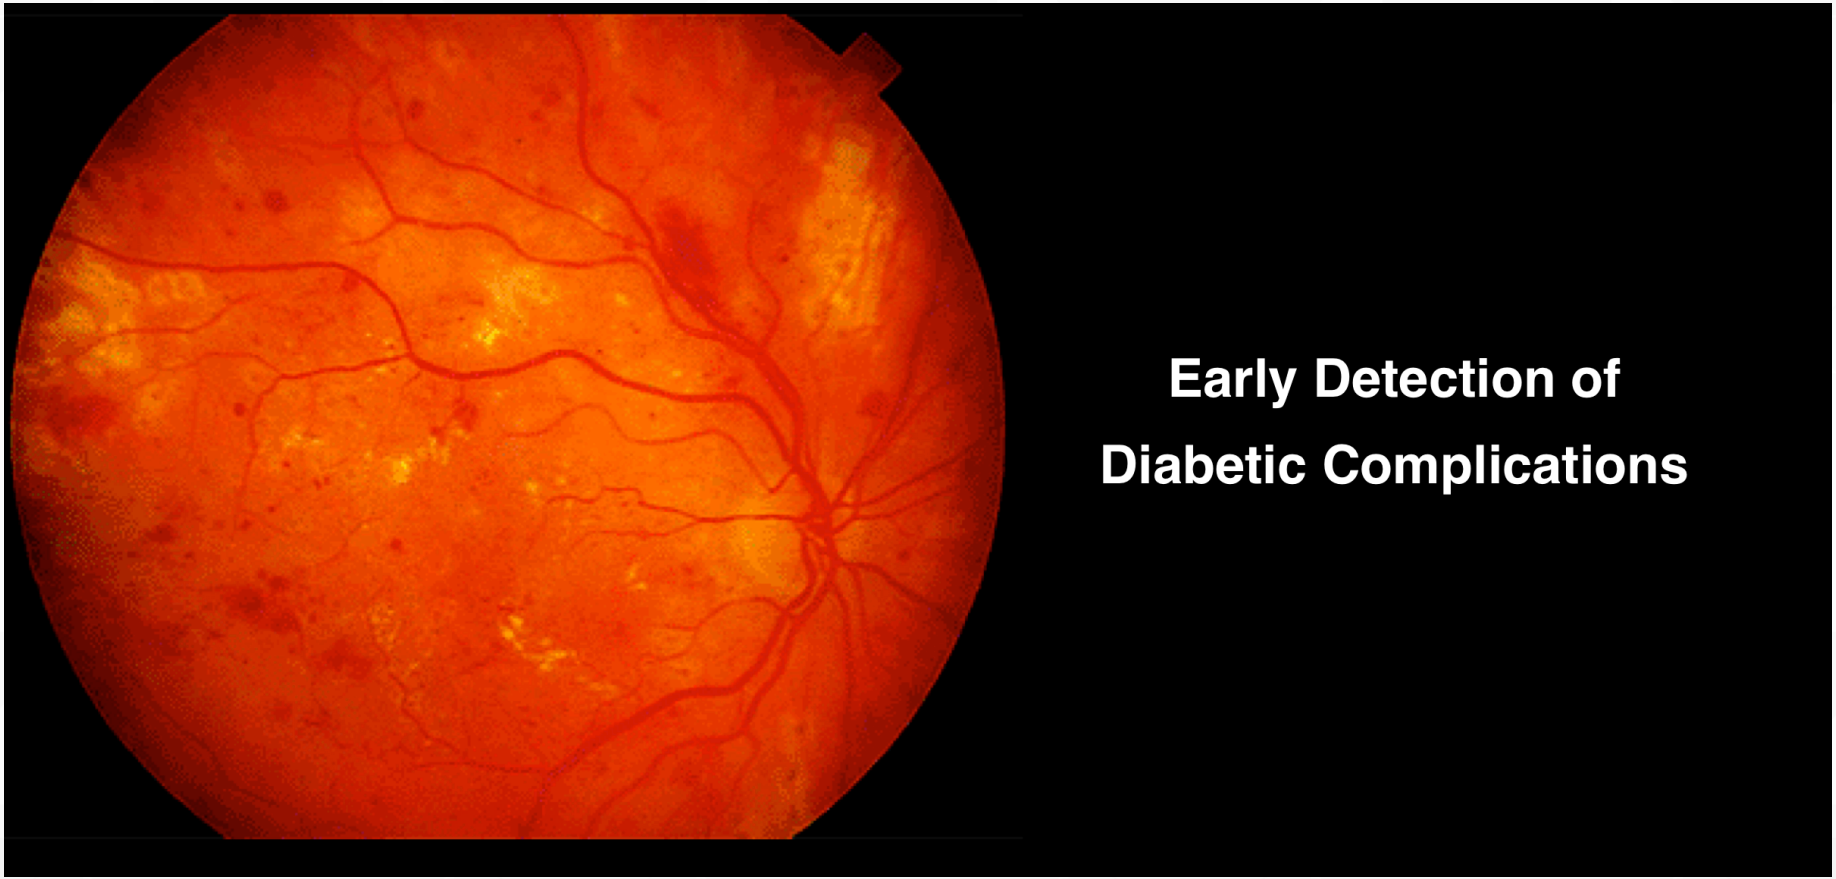

Skin Cancer Detection
At Physician-Levels (or better)

Lung Cancer Detection With Deep Learning & Medical Imaging

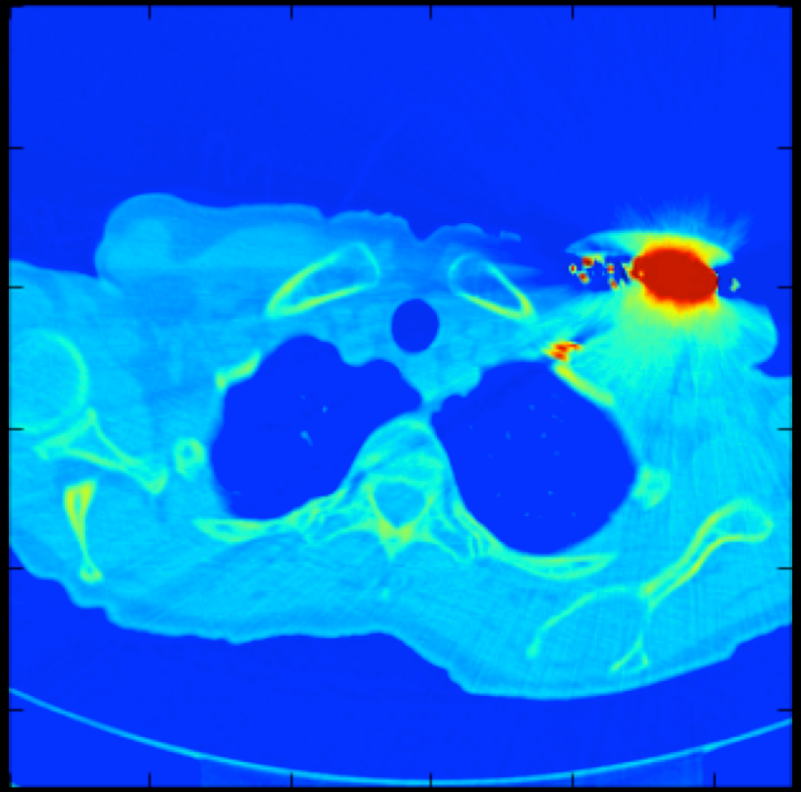

Amazon Lex

Conversational interfaces for your applications

Powered by the same deep learning technologies as Alexa

OhioHealth

"We are excited about utilizing evolving speech recognition and natural language processing technology to enhance the lives of our customers. Amazon Lex represents a great opportunity for us to deliver a better experience to our patients. Everything we do at OhioHealth is ultimately about providing the right care to our patients at the right time and in the right place. Amazon Lex's next generation technology and the innovative applications we are developing using it will help provide an improved customer experience. We are just scratching the surface of what is possible."

– Michael Krouse, Senior Vice President Operational Support and Chief Information officer, OhioHealth

Natural Language Understanding (NLU) & Automatic Speech Recognition (ASR)

as in Amazon ALEXA - Powered by Deep Learning

https://aws.amazon.com/lex/

amazon web services

Amazon Polly

Turn text into lifelike speech using deep learning

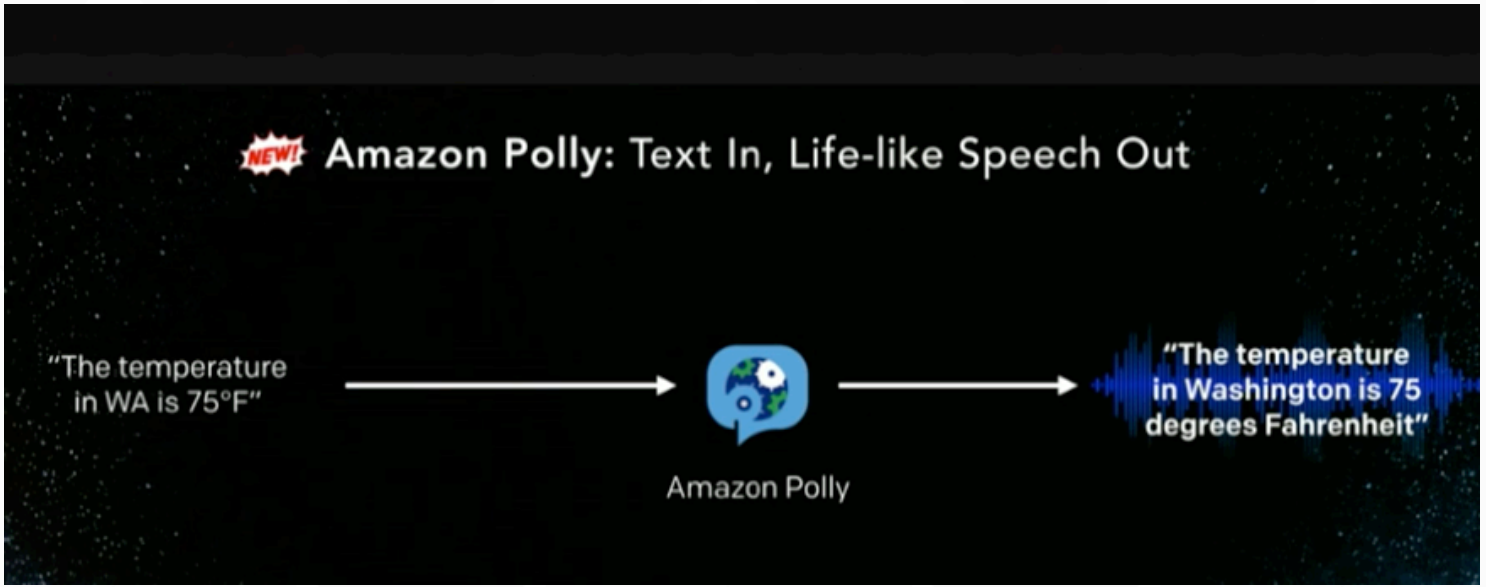

Incorporates ~47 different voices and fully managed services

https://aws.amazon.com/polly/

amazon web services

Deep Learning

Significantly improve many applications on multiple domains

image understanding

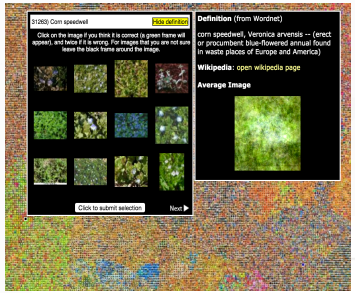

speech recognition

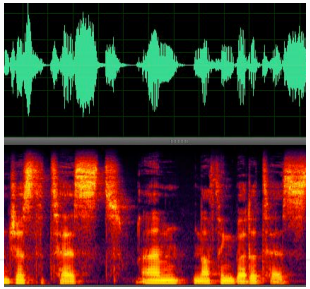

natural language processing

autonomy

. . .

amazon web services

Autonomous Driving Systems

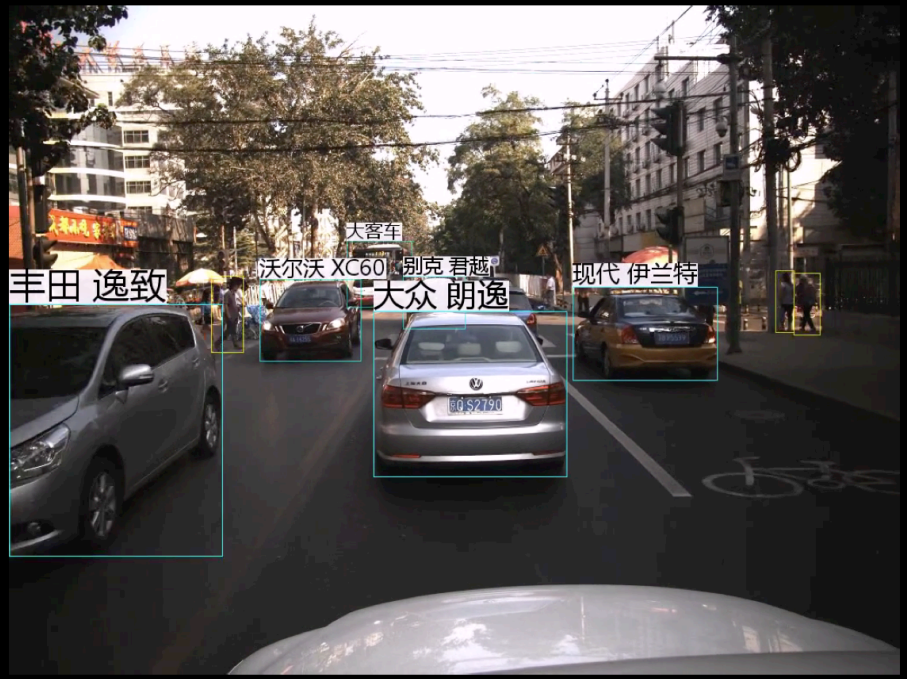

“The future is here,

It’s just not evenly distributed yet”

William Gibson

Amazon AI: Building Intelligent Systems

Model
Training

Inference
in the Cloud

Inference
at the Edge

Amazon AI: Democratized Artificial Intelligence

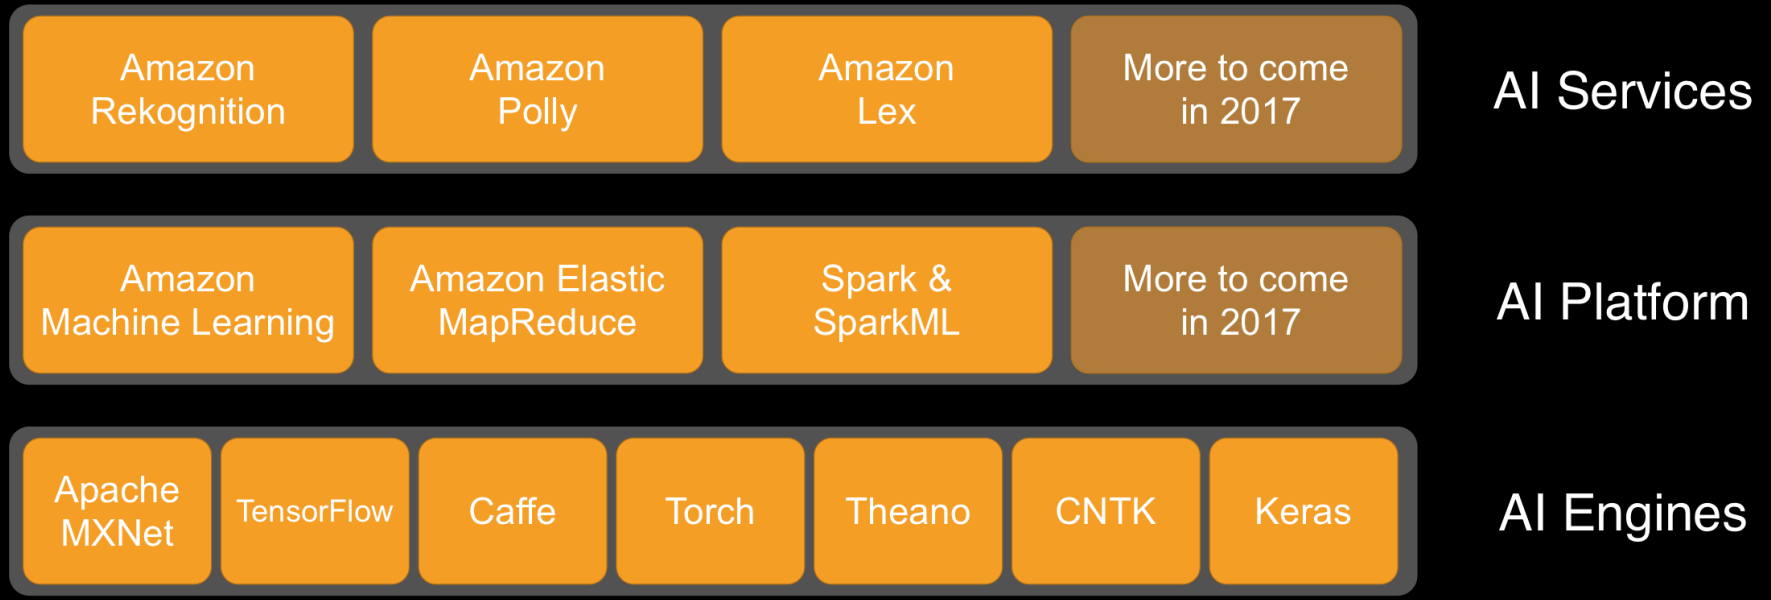

aws research initiative

https://aws.amazon.com/government-education/research-and-technical-computing/nsf-aribd/

- AWS initiated collaborative program with the National Science Foundation (NSF)
- The program by multiple directorates at NSF, provides funds up to $26.5 million in addition to $3 million in AWS Cloud Credits to perform **cutting edge Big Data research on cloud for a period of 3-4 years** (up to 2021)
- **This opens up a venue for collaborative programs with national, federal, and state agencies.**

https://www.nsf.gov/pubs/2017/nsf17534/nsf17534.htm

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)

amazon web services

Thank you!

Sanjay Padhi

sanpadhi@amazon.com