Amazon Innovation, Culture & AWS Resources for education

Louis Ho 何盧穎

Head of Business Development, Public Sector

holouis@amazon.com

A **quick update** on the AWS business…

$26.7b

revenue run rate

(ANNUALIZED FROM Q3 2018)

45.7%

y/y growth

(Q3 2018 VS Q3 2017)

Millions of active customers

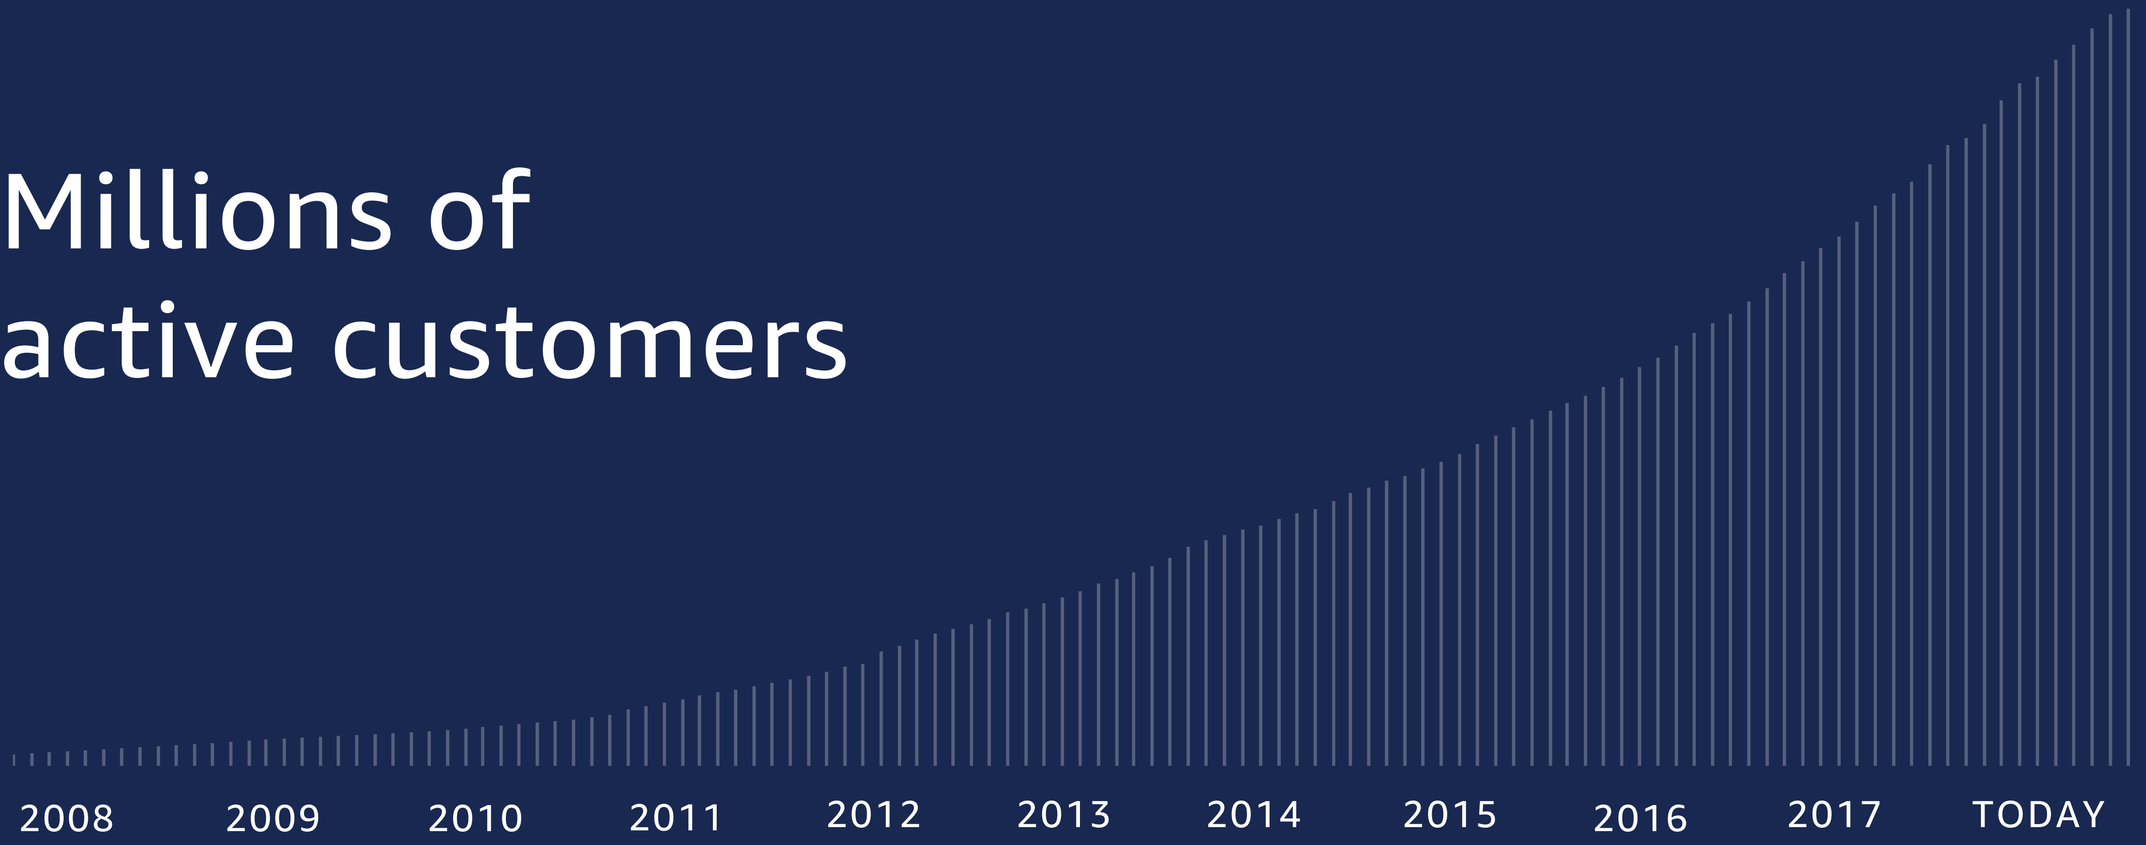

MOST ROBUST, FULLY FEATURED TECHNOLOGY INFRASTRUCTURE PLATFORM

TECHNICAL & BUSINESS SUPPORT

- Support
- Professional Services
- Optimization Guidance
- Partner Ecosystem
- Training & Certification
- Solutions Management
- Account Management
- Security & Billing Reports
- Personalized Dashboard

MARKETPLACE

- Business Apps
- Business Intelligence
- DevOps Tools
- Security
- Networking
- Databases
- Storage

ANALYTICS

- Data Warehousing
- Business Intelligence
- Hadoop/Spark
- Streaming Data Analysis
- Streaming Data Collection
- Elasticsearch
- Data Pipelines
- Interactive SQL Queries
- ETL

APP SERVICES

- Queuing & Notifications
- Workflow
- Search
- Email
- Transcoding

DEV OPS

- One-click App Deployment
- Resource Templates
- Build & Test
- Application Lifecycle Management
- DevOps Resource Management
- Triggers
- Containers
- Analyze & Debug
- Patching

MOBILE SERVICES

- API Gateway
- Single Integrated Console
- Identity
- Sync
- Mobile Analytics
- Mobile App Testing
- Targeted Push Notifications

IoT

- Rules Engine
- Device Shadows
- Device SDKs
- Device Gateway
- Registry
- Local Compute

MACHINE LEARNING

- Custom Model Training & Hosting
- Image & Scene Recognition
- Facial Recognition & Analysis
- Facial Search
- Text to Speech
- Conversational Chatbots
- Deep Learning (Apache MXNet, TensorFlow, & others)

ENTERPRISE APPS

- Virtual Desktops
- Sharing & Collaboration
- Corporate Email
- App Streaming
- Communications
- Contact Center

HYBRID ARCHITECTURE

- Data Integration
- Integrated Networking
- Integrated Identity & Access
- Integrated Resource & Deployment Management
- Integrated Devices & Edge Systems

MIGRATION

- Schema Conversion
- Exabyte-Scale Data Migration
- Application Migration
- Database Migration
- Server Migration

INFRASTRUCTURE

- Regions
- Availability Zones
- Points of Presence

CORE SERVICES

- **Compute** — VMs, Auto-scaling, Load Balancing, Containers, Virtual Private Servers, Batch Computing, Cloud Functions, Elastic GPUs, Edge Computing
- **Storage** — Object, Blocks, File, Archivals, Import/Export, Exabyte-scale data transfer
- **Databases** — Relational, NoSQL, Caching, Migration, PostgreSQL compatible
- **Networking** — VPC, DX, DNS
- **CDN**

SECURITY & COMPLIANCE

- Identity Management
- Configuration Compliance
- Access Control
- Key Management & Storage
- Monitoring & Logs
- Account Grouping
- Assessment & Reporting
- Resource & Usage Auditing
- Web Application Firewall
- DDOS Protection

MANAGEMENT TOOLS

- Manage Resources
- Monitoring
- Service Catalogue
- Server Management
- Configuration Tracking
- Resource Templates

AI SERVICES

Vision			Speech		Language		Chatbots	Forecasting	Recommendations
REKOGNITION IMAGE	REKOGNITION VIDEO	TEXTRACT	POLLY	TRANSCRIBE	TRANSLATE	COMPREHEND	LEX	FORECAST	PERSONALIZE

ML SERVICES

AMAZON SAGEMAKER

AWS DeepLens

AWS DEEP RACER

BUILD	TRAIN	DEPLOY
Pre-built algorithms & notebooks	One-click model training & tuning	One-click deployment & hosting
Data labeling (GROUND TRUTH)	Reinforcement learning	
Algorithms & models (AWS MARKETPLACE FOR MACHINE LEARNING)	Optimization (NEO)	

ML FRAMEWORKS & INFRASTRUCTURE

Frameworks	Interfaces	Infrastructure				
TensorFlow	GLUON	EC2 P3 & P3dn	EC2 C5	FPGAs	GREENGRASS	ELASTIC INFERENCE
mxnet	Keras					
PYTORCH						

aws

“AIoT” (AI and IoT) – Industrial Trend

ML: Train in the Cloud, Infer at the Edge

- ✓ Process State
- ✓ Product Quality
- ✓ Equipment Health
- ✓ Output Forecast
- ✓ Equipment Efficiency
- ✓ Leak Detection
- ✓ Asset Planning
- ✓ Ops Optimization
- ✓ Tool Productivity

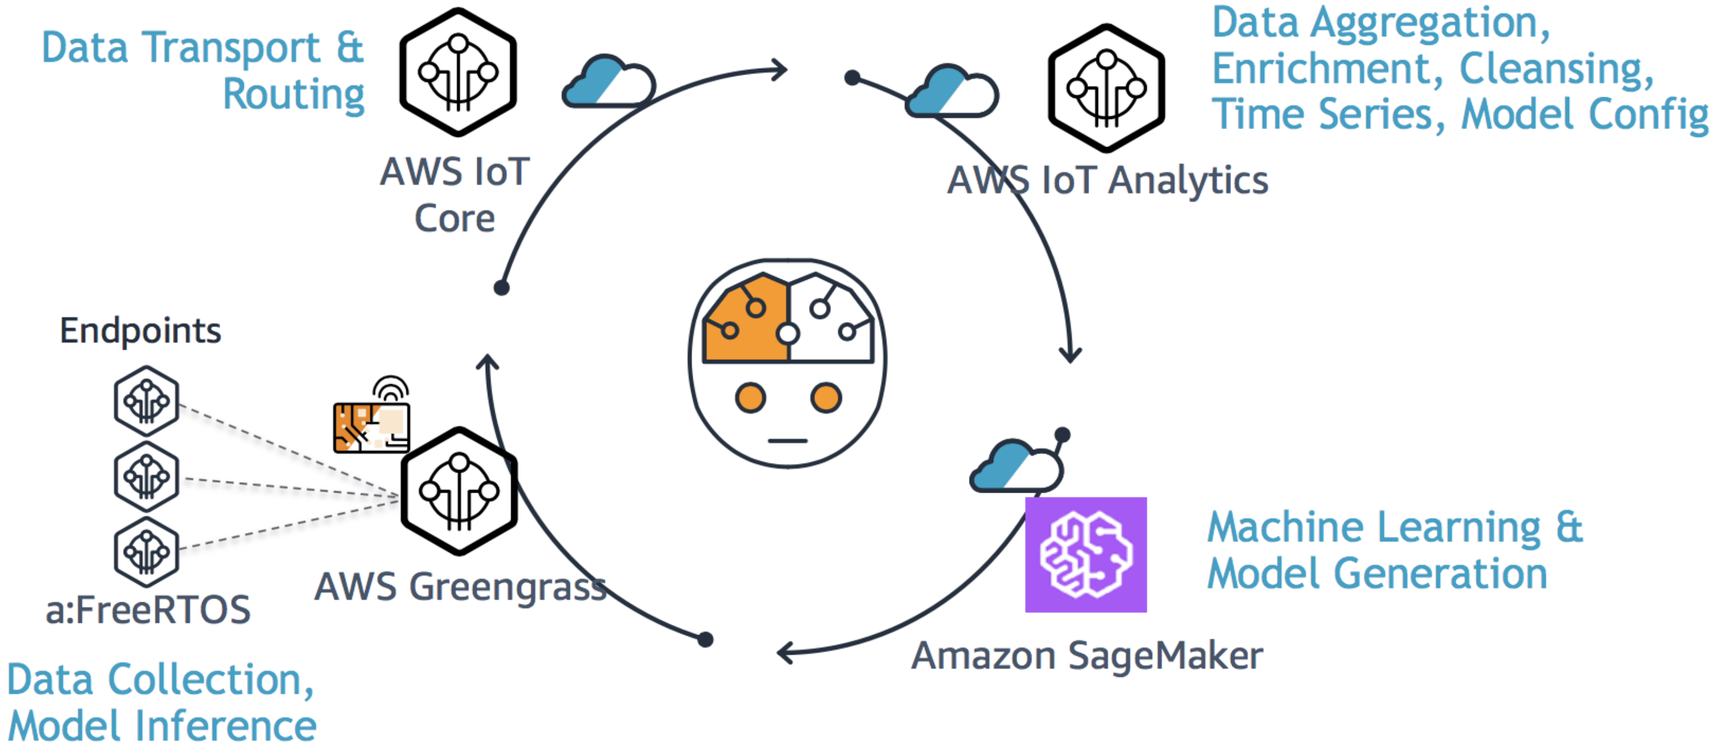

PACE OF INNOVATION

New capabilities daily

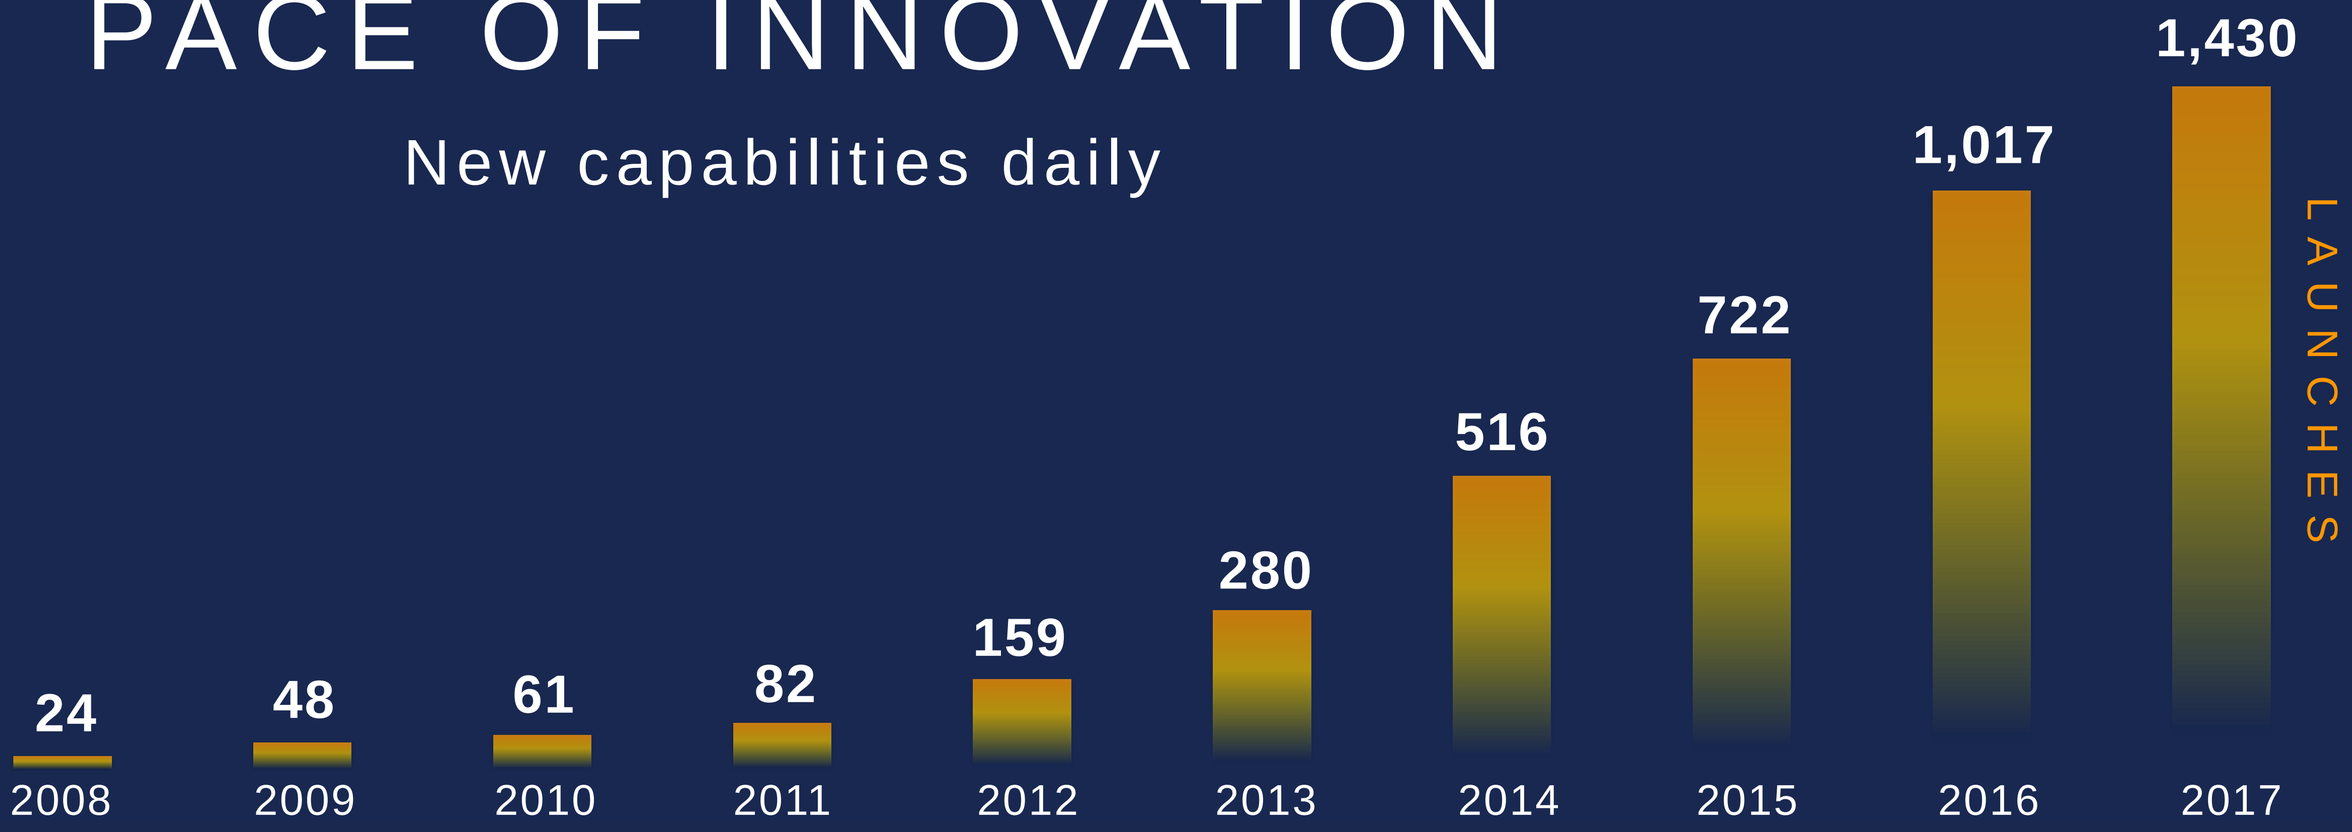

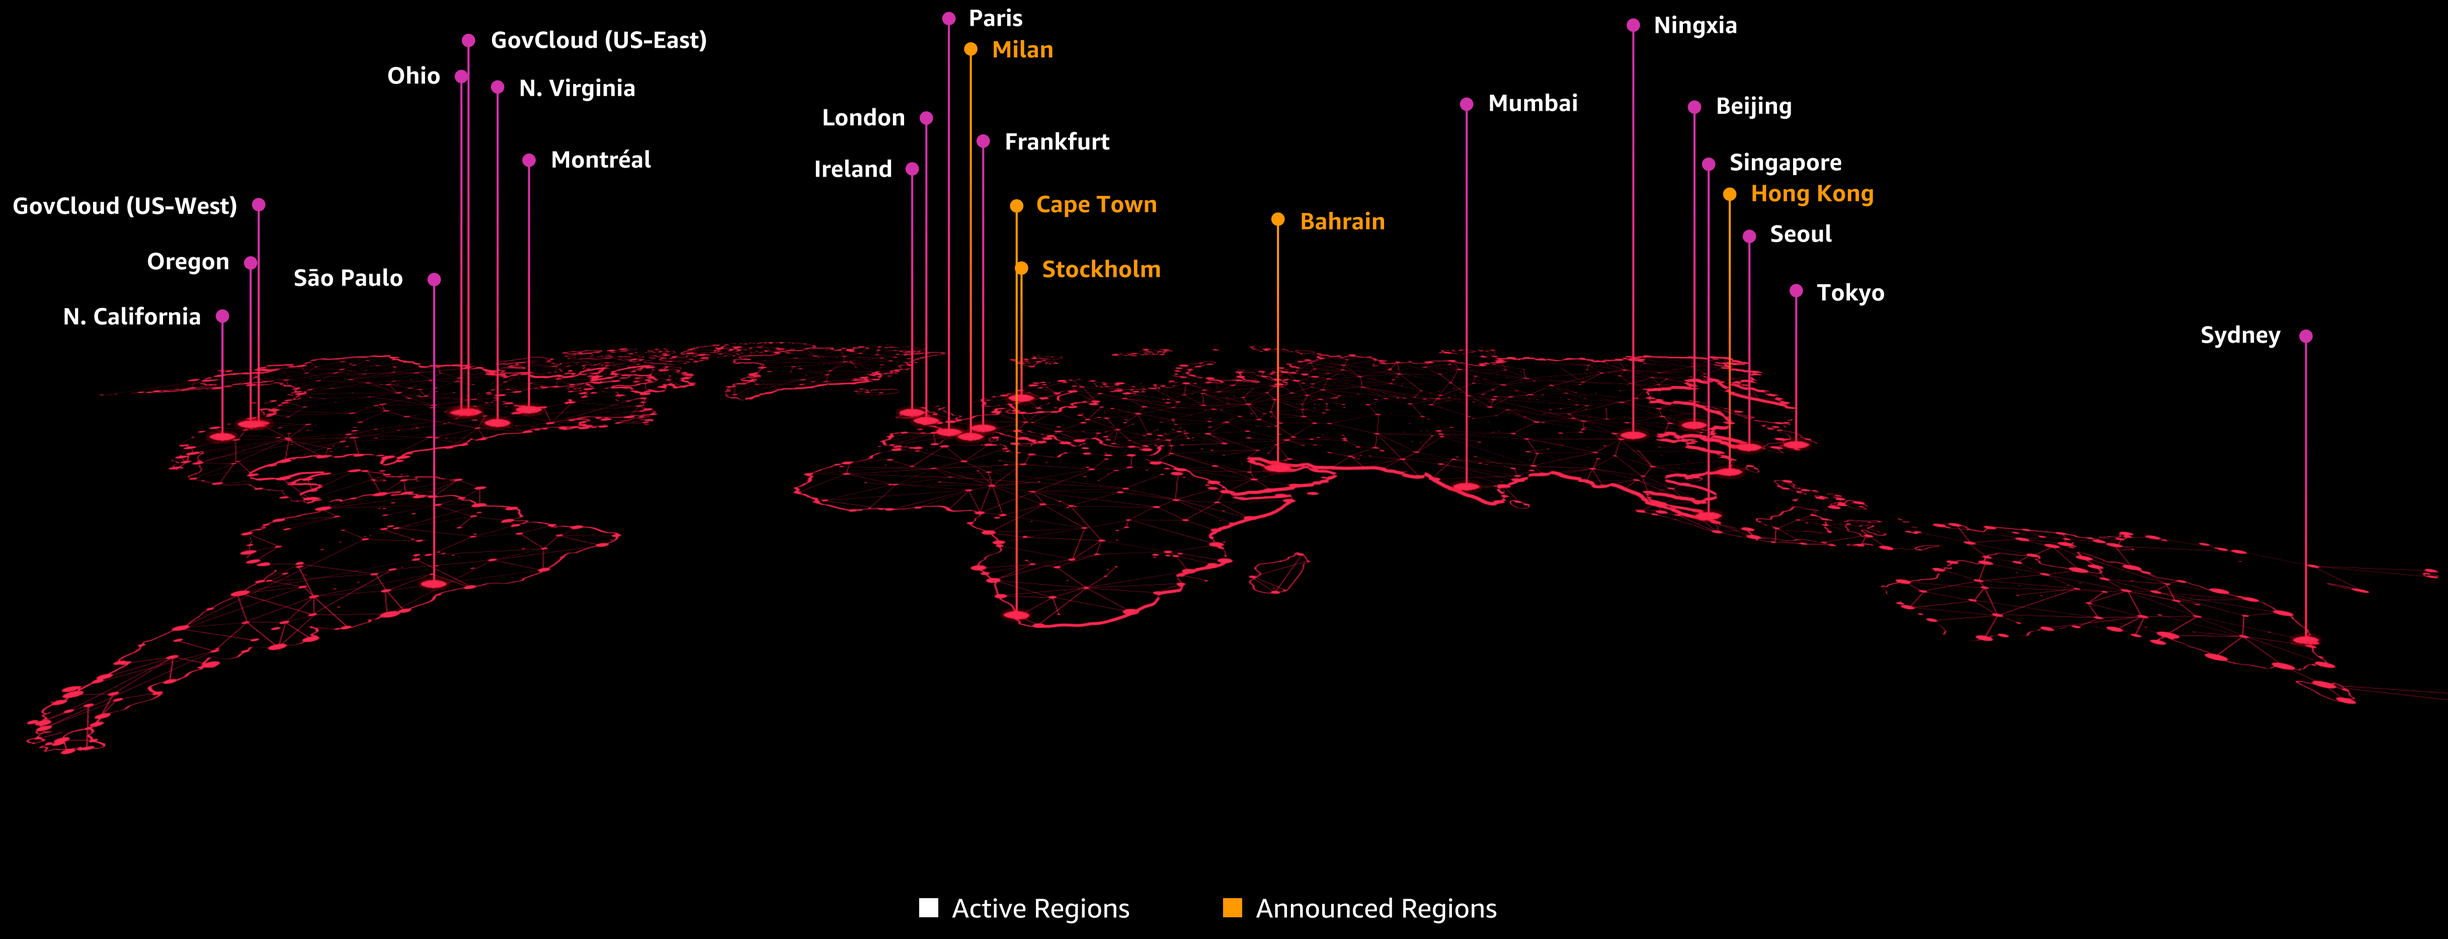
GovCloud (US-East)
Ohio
N. Virginia
Montréal
GovCloud (US-West)
Oregon
São Paulo
N. California
Paris
Milan
London
Frankfurt
Ireland
Cape Town
Stockholm
Bahrain
Mumbai
Ningxia
Beijing
Singapore
Hong Kong
Seoul
Tokyo
Sydney
Active Regions
Announced Regions

Fully-scaled Availability Zone
Highly redundant regional network with **massive** capacity

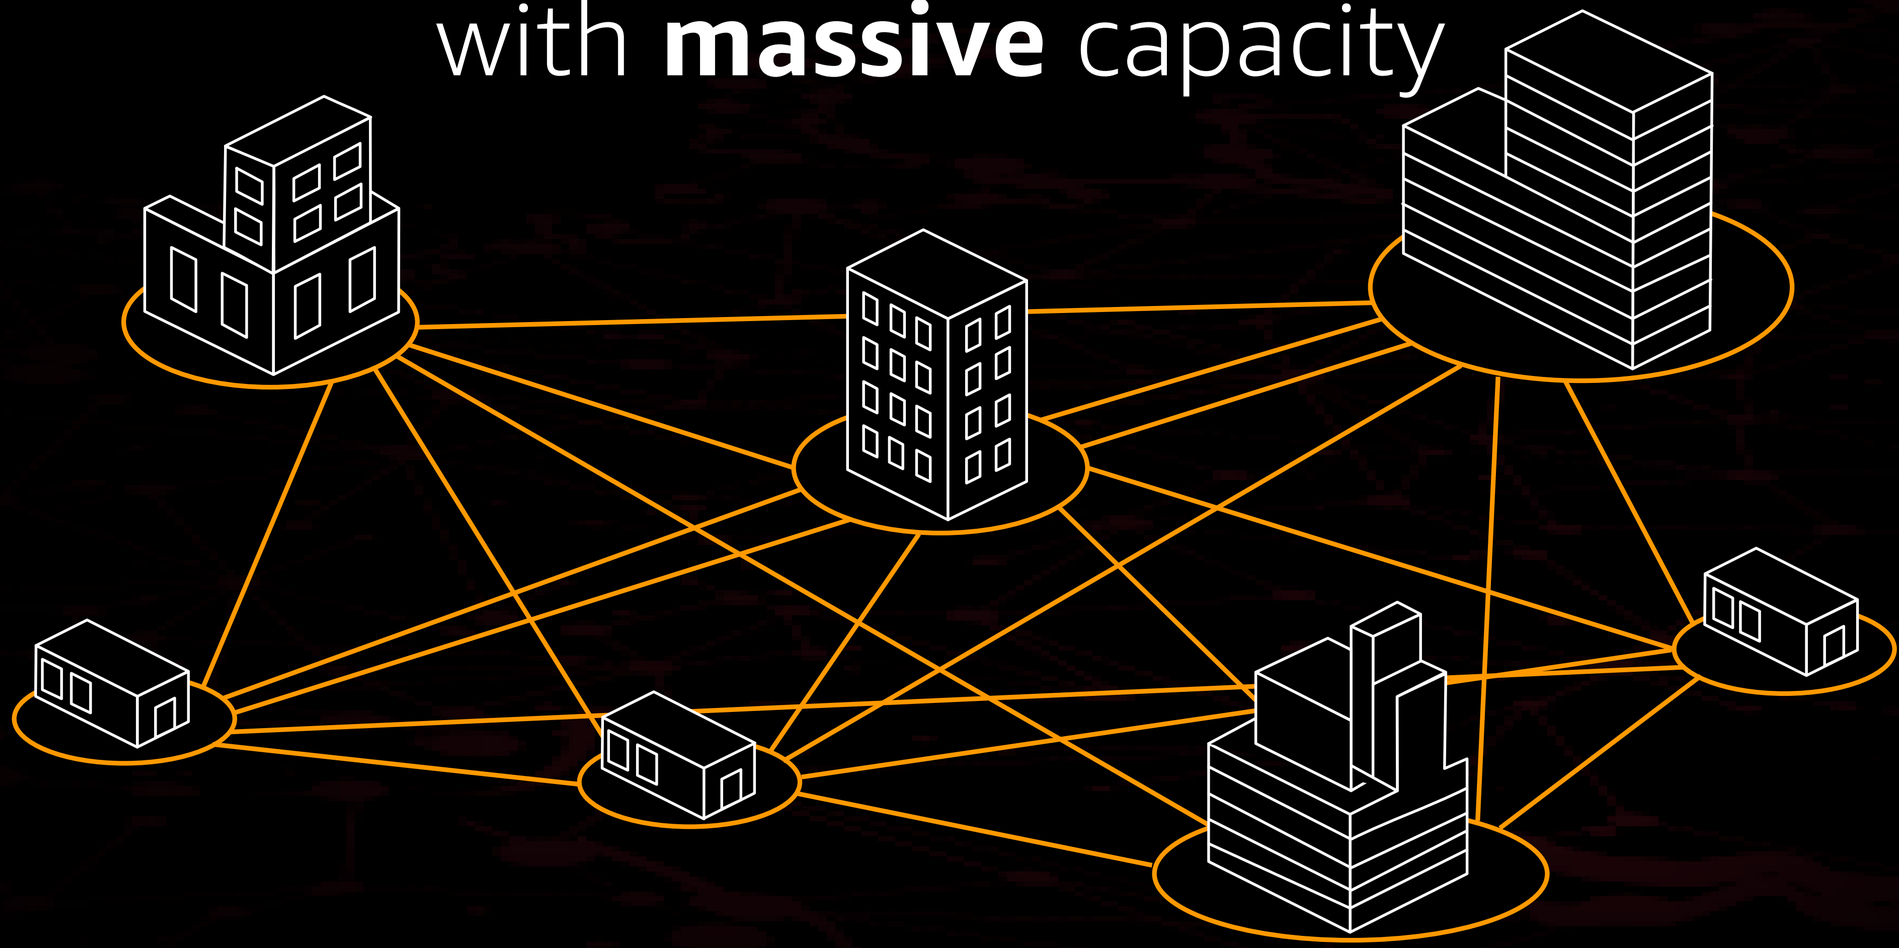

亞馬遜雲端運算服務 (AWS)

AWS Positioned as a Leader in the IDC MarketScape: Worldwide Infrastructure as a Service 2017 Vendor Assessment*

"AWS has effectively defined and led the core offering portfolio in the public cloud IaaS market."

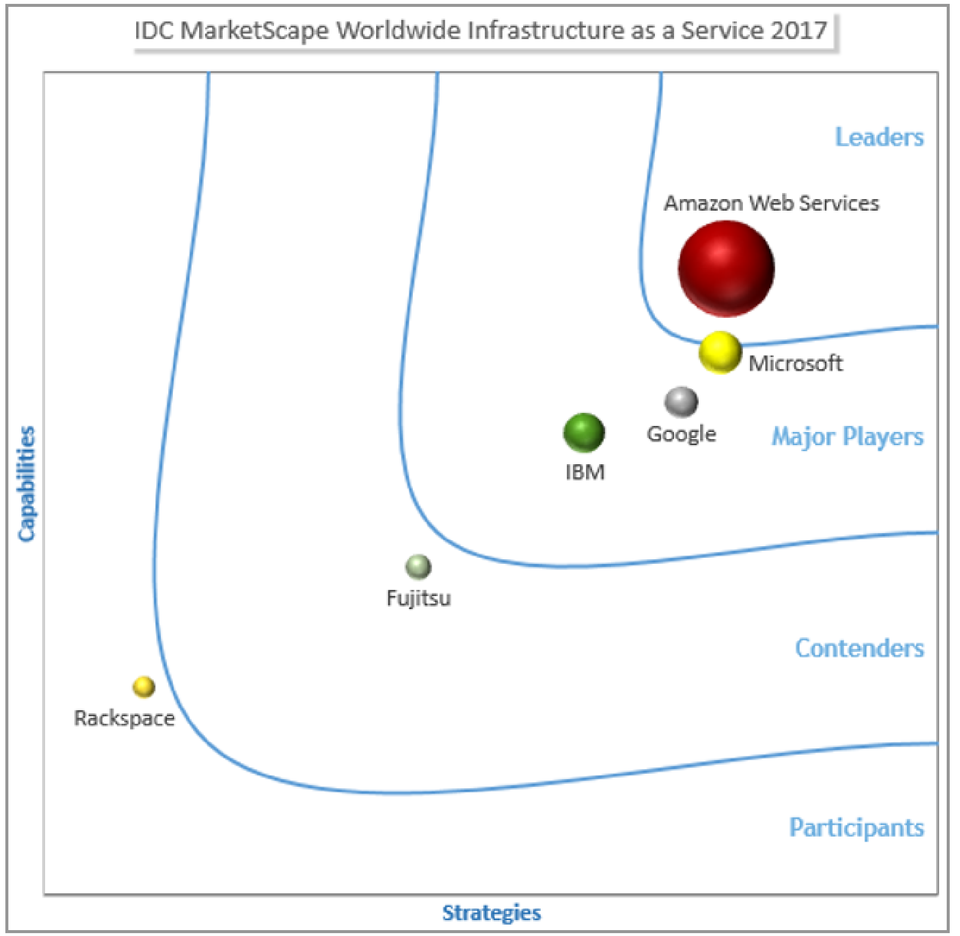

AWS (Amazon Web Services) 亞馬遜網路服務

- 2006年開始提供服務
- 全球雲端運算服務的領導者
- 年營業收入超過267 億美金 (Q3, 2018, 46.7% YoY)
- 服務于遍佈190個國家數百萬活躍使用者
- 全球基礎設施

aws 44.1%

Market Share: Public Cloud Services, Worldwide, 2016, 19 September 2017, Analyst(s): Sid Nag, Chad Eschinger, Fred Ng, David Edward Ackerman, Fabrizio Biscotti, Ed Anderson, Mike Dorosh, Colleen Graham

Building A Day 1 Business

It’s always Day 1

Amazon.com passed many milestones in 1997…

But this is Day 1 for the Internet and, if we execute well, for Amazon.com. Today, online commerce saves customers money and precious time. Tomorrow, through personalization, online commerce will accelerate the very process of discovery.

Amazon.com uses the Internet to create real value for its customers and, by doing so, hopes to create an enduring franchise, even in established and large markets.

貝佐斯在1997年寫下了「第1天」的意思：
…如果我們做得還算不錯的話，這算得上是造就了Amazon.com與互聯網的「第一天」。在今天，線上商務為客戶節省了大量的金錢和寶貴的時間。 透過個人化，明天的線上商務將進一步加速這樣的過程。 Amazon.com使用網際網路為我們的客戶創造了實際的價值，藉此並希望創建一個即便面對更成熟與大型的市場，仍然可長可久的經銷模式。

It’s always Day 1

True Customer Obsession

Resisting Proxies

Embracing External Tailwinds

High Velocity Decision Making

amazon marketplace

It's always Day 1

- True Customer Obsession
- Resisting Proxies
- Embracing External Tailwinds
- High Velocity Decision Making


It’s always Day 1
True Customer Obsession
Resisting Proxies
Embracing External Tailwinds
High Velocity Decision Making
INTRODUCING
amazon go

Day 2

“Day 2 is decline, followed by death. That is why it is always Day 1.”

Jeff Bezos, Amazon

對於貝佐斯來說，他絕不想讓Amazon進入到「Day 2」，因為：

“Day 2是停滯不前的，其次是無關緊要的。隨之而來的是令人難以忍受的痛苦與衰退，隨後死亡。 這就是為什麼我們總是要活在「Day 1」。“

Stay as the Day1Company with AWS

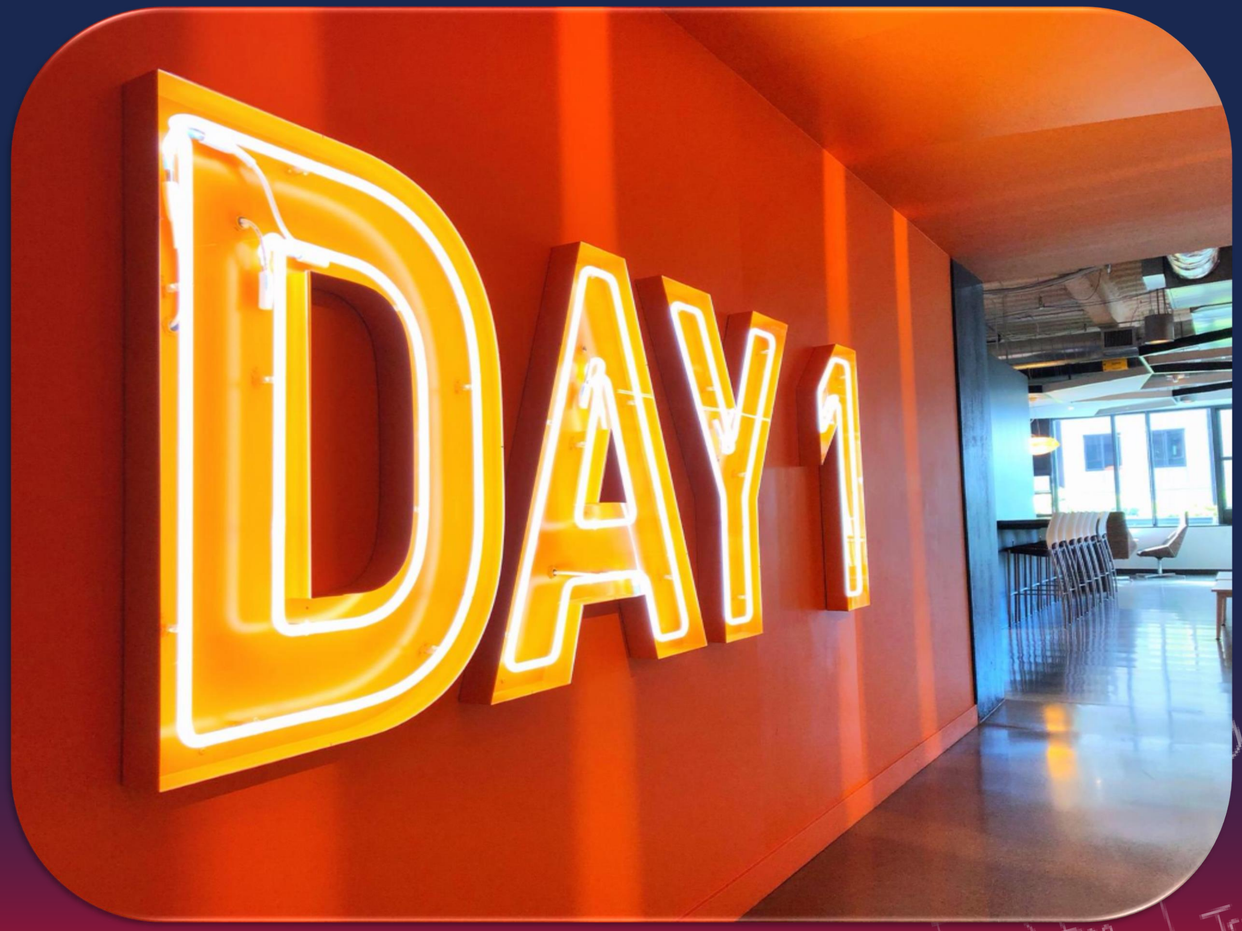

Digital Transformation and Cloud Journey

Amazon Web Services since 2006

- Any decision can be undone
- Disagree and commit with yourself
- Act with 70% information

Typical weekly traffic to Amazon.com

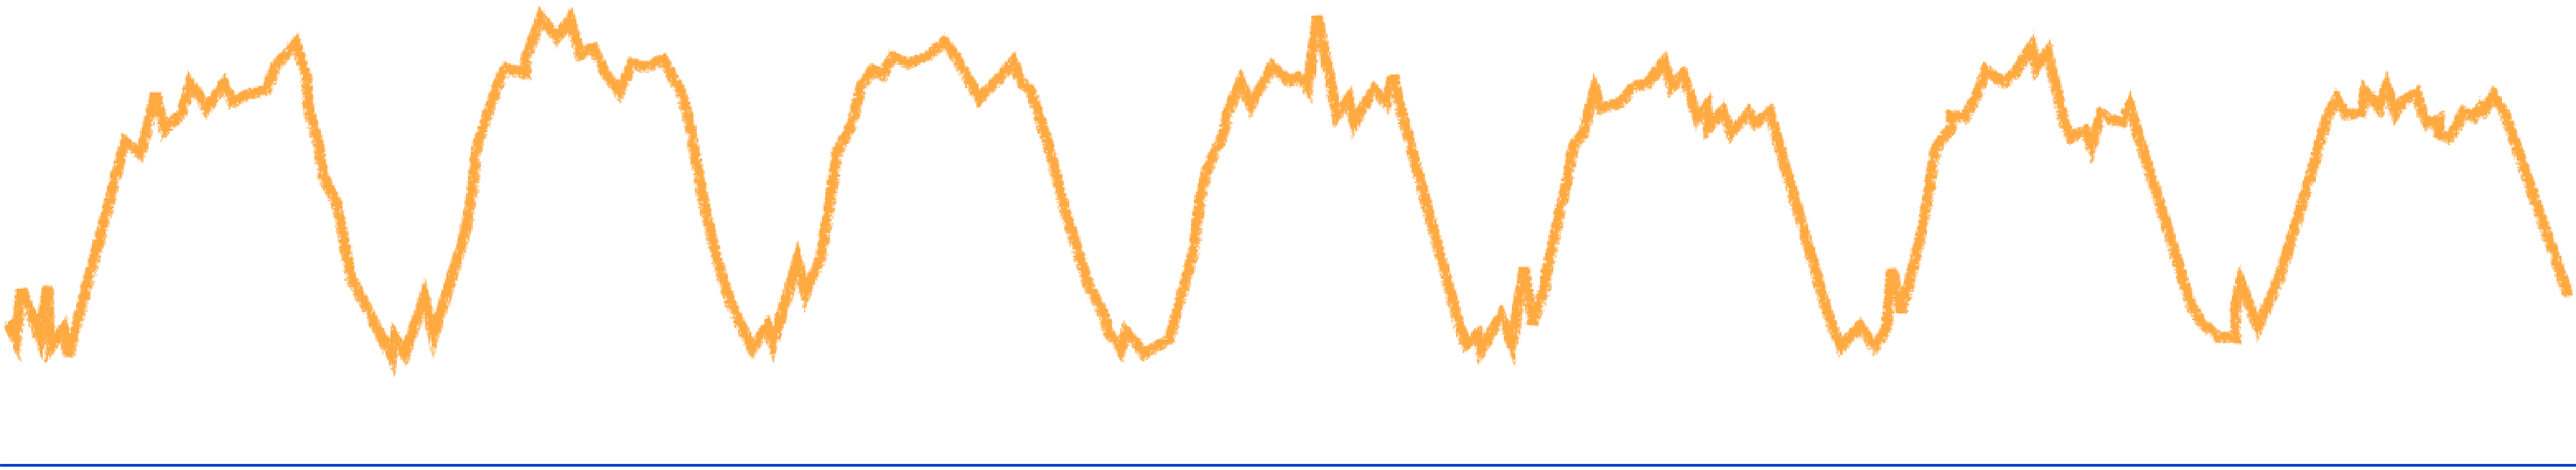

Sunday Monday Tuesday Wednesday Thursday Friday Saturday

aws

November traffic to Amazon.com

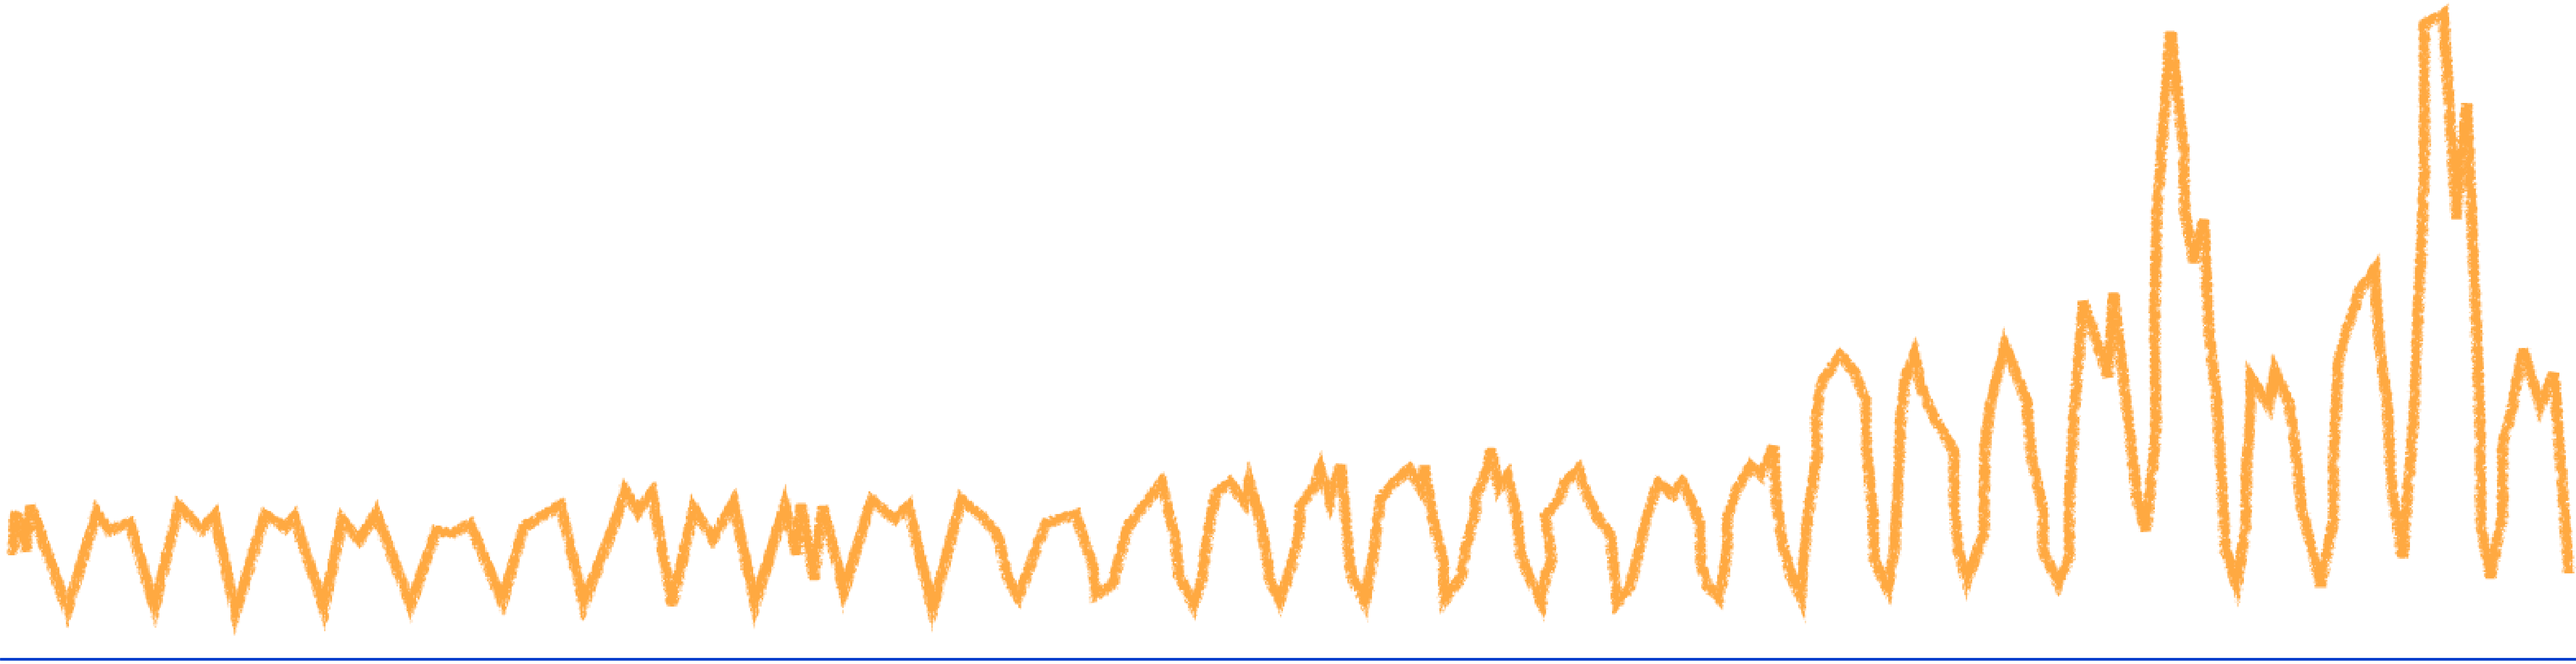

November

November traffic to Amazon.com

Provisioned capacity

November

November traffic to Amazon.com

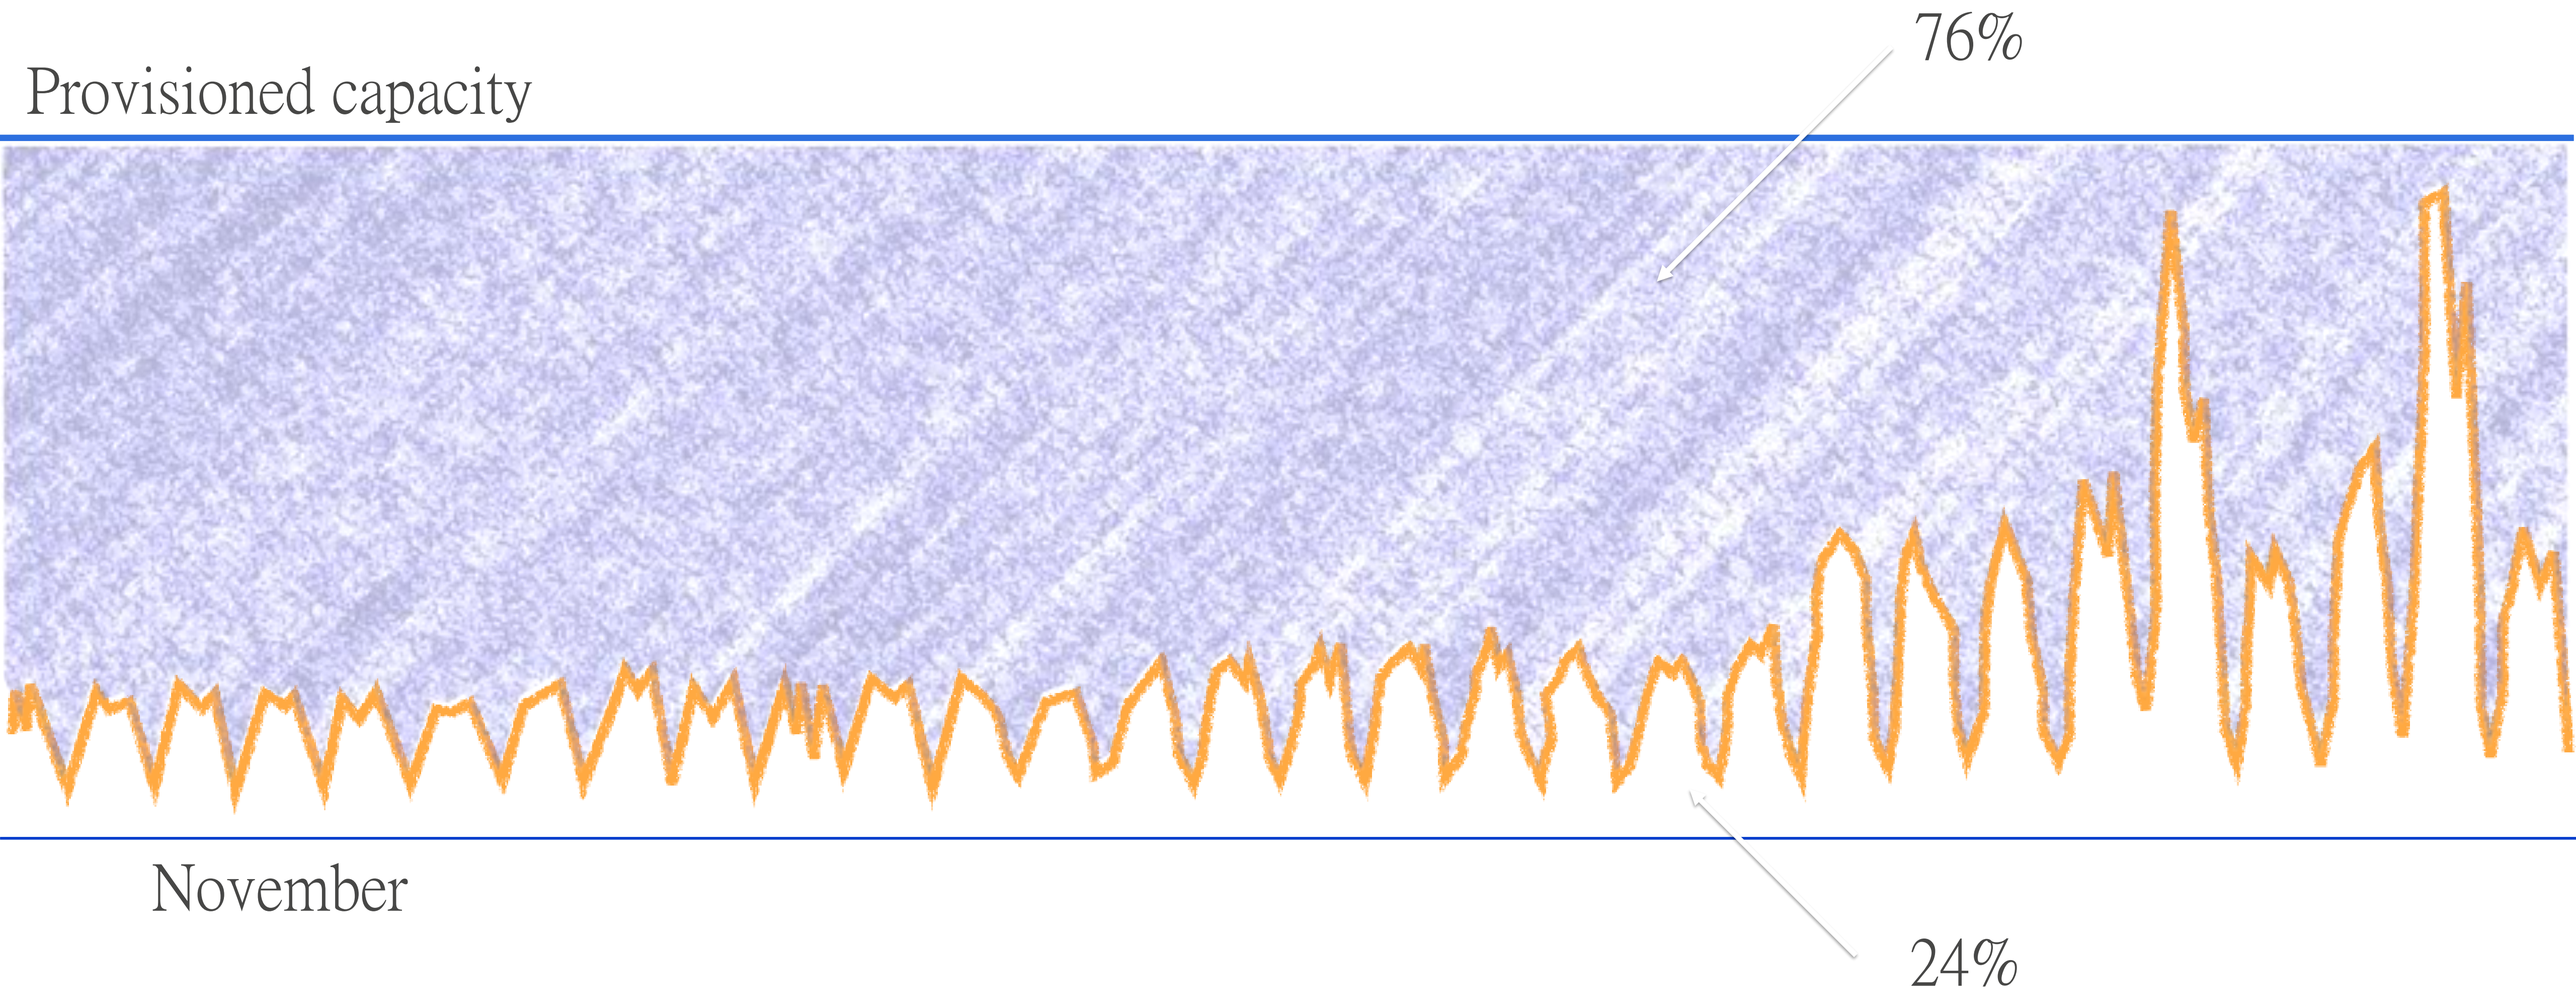

Amazon is innovating across many domains

Drone Development

Video Streaming

Kindle

In-house Entertainment

Grocery Delivery

Advanced Shopping

Cloud Computing

Home Automation

“Why Amazon Is The World’s Most Innovative Company Of 2017”
fastcompany.com

(Culture * Mechanisms)

F(I)=Organization * Architecture

Culture

Amazon Leadership Principles

Customer Obsession
Ownership
Invent and Simplify
Are Right, A Lot
Hire and Develop the Best
Insist on the Highest Standards
Think Big
Bias for Action
Frugality
Learn and Be Curious
Earn Trust
Dive Deep
Have Backbone; Disagree and Commit
Deliver Results

Amazon Leadership Principles

Customer Obsession
Ownership
Invent and Simplify
Are Right, A Lot
Hire and Develop the Best
Insist on the Highest Standards
Think Big
Bias for Action
Frugality
Learn and Be Curious
Earn Trust
Dive Deep
Have Backbone; Disagree and Commit
Deliver Results

SPEED VS QUALITY OF DECISION

“…most decisions should probably be made with somewhere around 70% of the informatio you wish you had.
If you wait for 90%, in most cases, you’re probably being slow”

“If you’re good at course correcting, being wrong may be less costly than you think, whereas being slow is going to be expensive for sure.”

Letter to Shareholders - April 12, 2017

Mechanisms


"WE'RE NOT COMPETITOR OBSESSED, **WE'RE CUSTOMER OBSESSED**. WE START WITH WHAT THE CUSTOMER NEEDS AND WE WORK BACKWARDS."

Press Release

When you write your press release, imagine that your customer is going to read it.
It’s a one-page narrative explaining your vision using customer-centric language

- No Marketing Buzz Words (e.g. simple, easy, fast). The reader decides if it's “easy” to use.
- The most important info first. No one reads past the first paragraph.
- Only metrics and data that matters to the customer (e.g. their time, their money).

Innovation and the Working Backwards Process

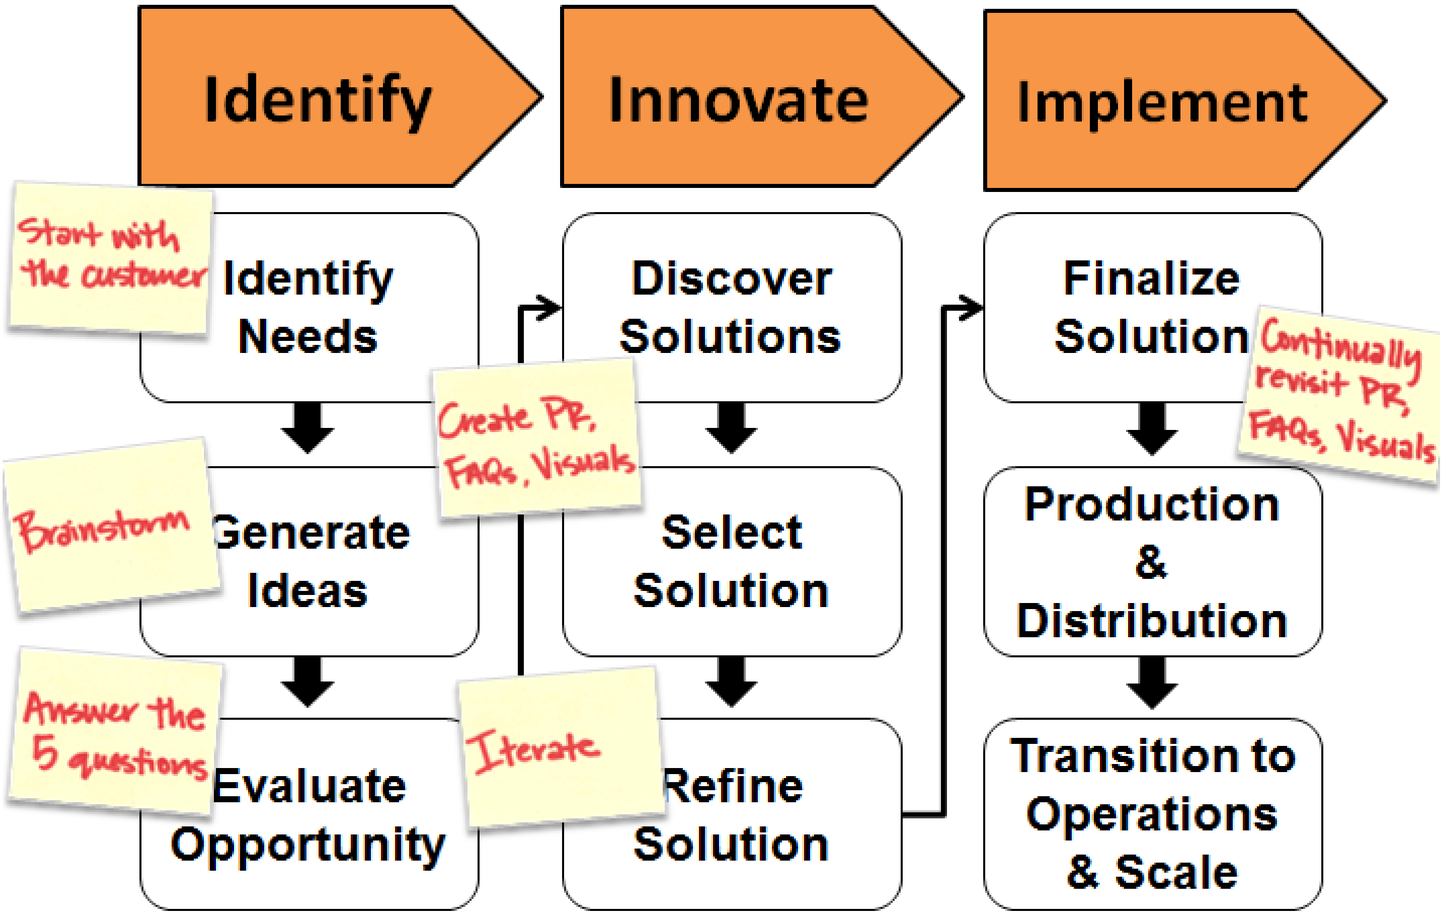

Organisation

Pushing Decisions Down


2 PIZZA TEAMS

Characteristics

Outcome-based

Iterative

Optimize for Responsiveness

Empowerment & Autonomy

Feedback & Learning

Safety & Trust

amazon web services

適用於教育及學術研究機構之AWS計畫

適用於教育及學術研究機構之AWS計畫

教學及證照

Educate: 旨在為學生和教育人員提供所需的資源，以大幅加速與雲端相關的學習，並協助提升未來企業家、人力和研究人員的能力。

Academy: 為大專院校提供 AWS 課程，以協助學生所需的雲端運算技能，並做好取得業界認可 AWS 證照的準備。

資料及研究

Public Data Sets: 託管可開放大眾免費使用的資料集。在AWS平台上分享數據，可讓世界各地的使用者據此進行研究及建置服務，而無需購置伺服器或儲存容量。

Credits for Research: 提供AWS雲端資源給研究暨學術單位從事研究計畫暨研究開發與應用。

新創及商轉

Activate: 為新創公司提供快速開始使用 AWS 所需的資源 – 包括雲端服務、培訓和支援。

Market Place: 透過AWS 的全球佈建，可以迅速將你的服務或解決方案擴展到全球。

Solution Architect 解決方案架構師
協助客戶在學習、研究、新創及佈署上給予架構上的建議及優化
以達到雲端資源的最佳運用

AWS Educate program benefits

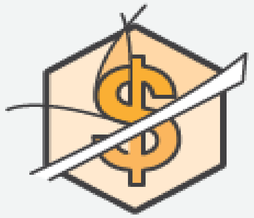

Grants for free usage of AWS services

Labs and training on cloud topics and AWS products

Open course content by leading professors and AWS

AWS Educate Cloud Career Pathways tied to industry job roles

- ✓ Holistic program to positively impact students around the world
- ✓ Curriculum change
- ✓ Student app development and entrepreneurship
- ✓ Accelerate career opportunities and corporate hiring pipeline

+JOBS*

*AWS Educate Job Board is for 18 years old and above

Cloud Skills Readiness with AWS Educate

aws educate

- **For students**
  - Starter Account (without credit card), Standard Account (with AWS Account) to receive **30-100 USD credits** for AWS service renew annually
  - **Free** access to **AWS Technical Essential Training Learning Course** (valued at $600USD)
  - NEW CAREER PATHWAY - **1 Cloud Intro Class and 11 Highly Demanded Job Roles Training** – Machine Learning Scientist, Big Data Scientist, Cloud Support Engineer, Web Developer, DevOps Engineer… etc.
  - **5 Badges** - gaming, startup, alexa ….
  - Job Board for **internship and full-time opportunities** – including Hong Kong jobs
- **For educators**
  - NEW Starter Account (without credit card), Standard Account (with AWS Account) to **50-200 USD credits** on AWS renew annually
  - **Free** access to **AWS Technical Essential Training Learning Course** (valued at $600USD)
  - **1500+ Institutes** Teaching Material Sharing
  - NEW FEATURE – **AWS Educate classroom** to offer hands-on lab in class with additional credits
  - **50% off AWS certifications exams**

New Educator's Portal

Search by keywords or specialty

Request for credits

Online training, explore discounted certifications, view students' career pathways

Feature box that can be highlights latest news in the Educate community to the educators

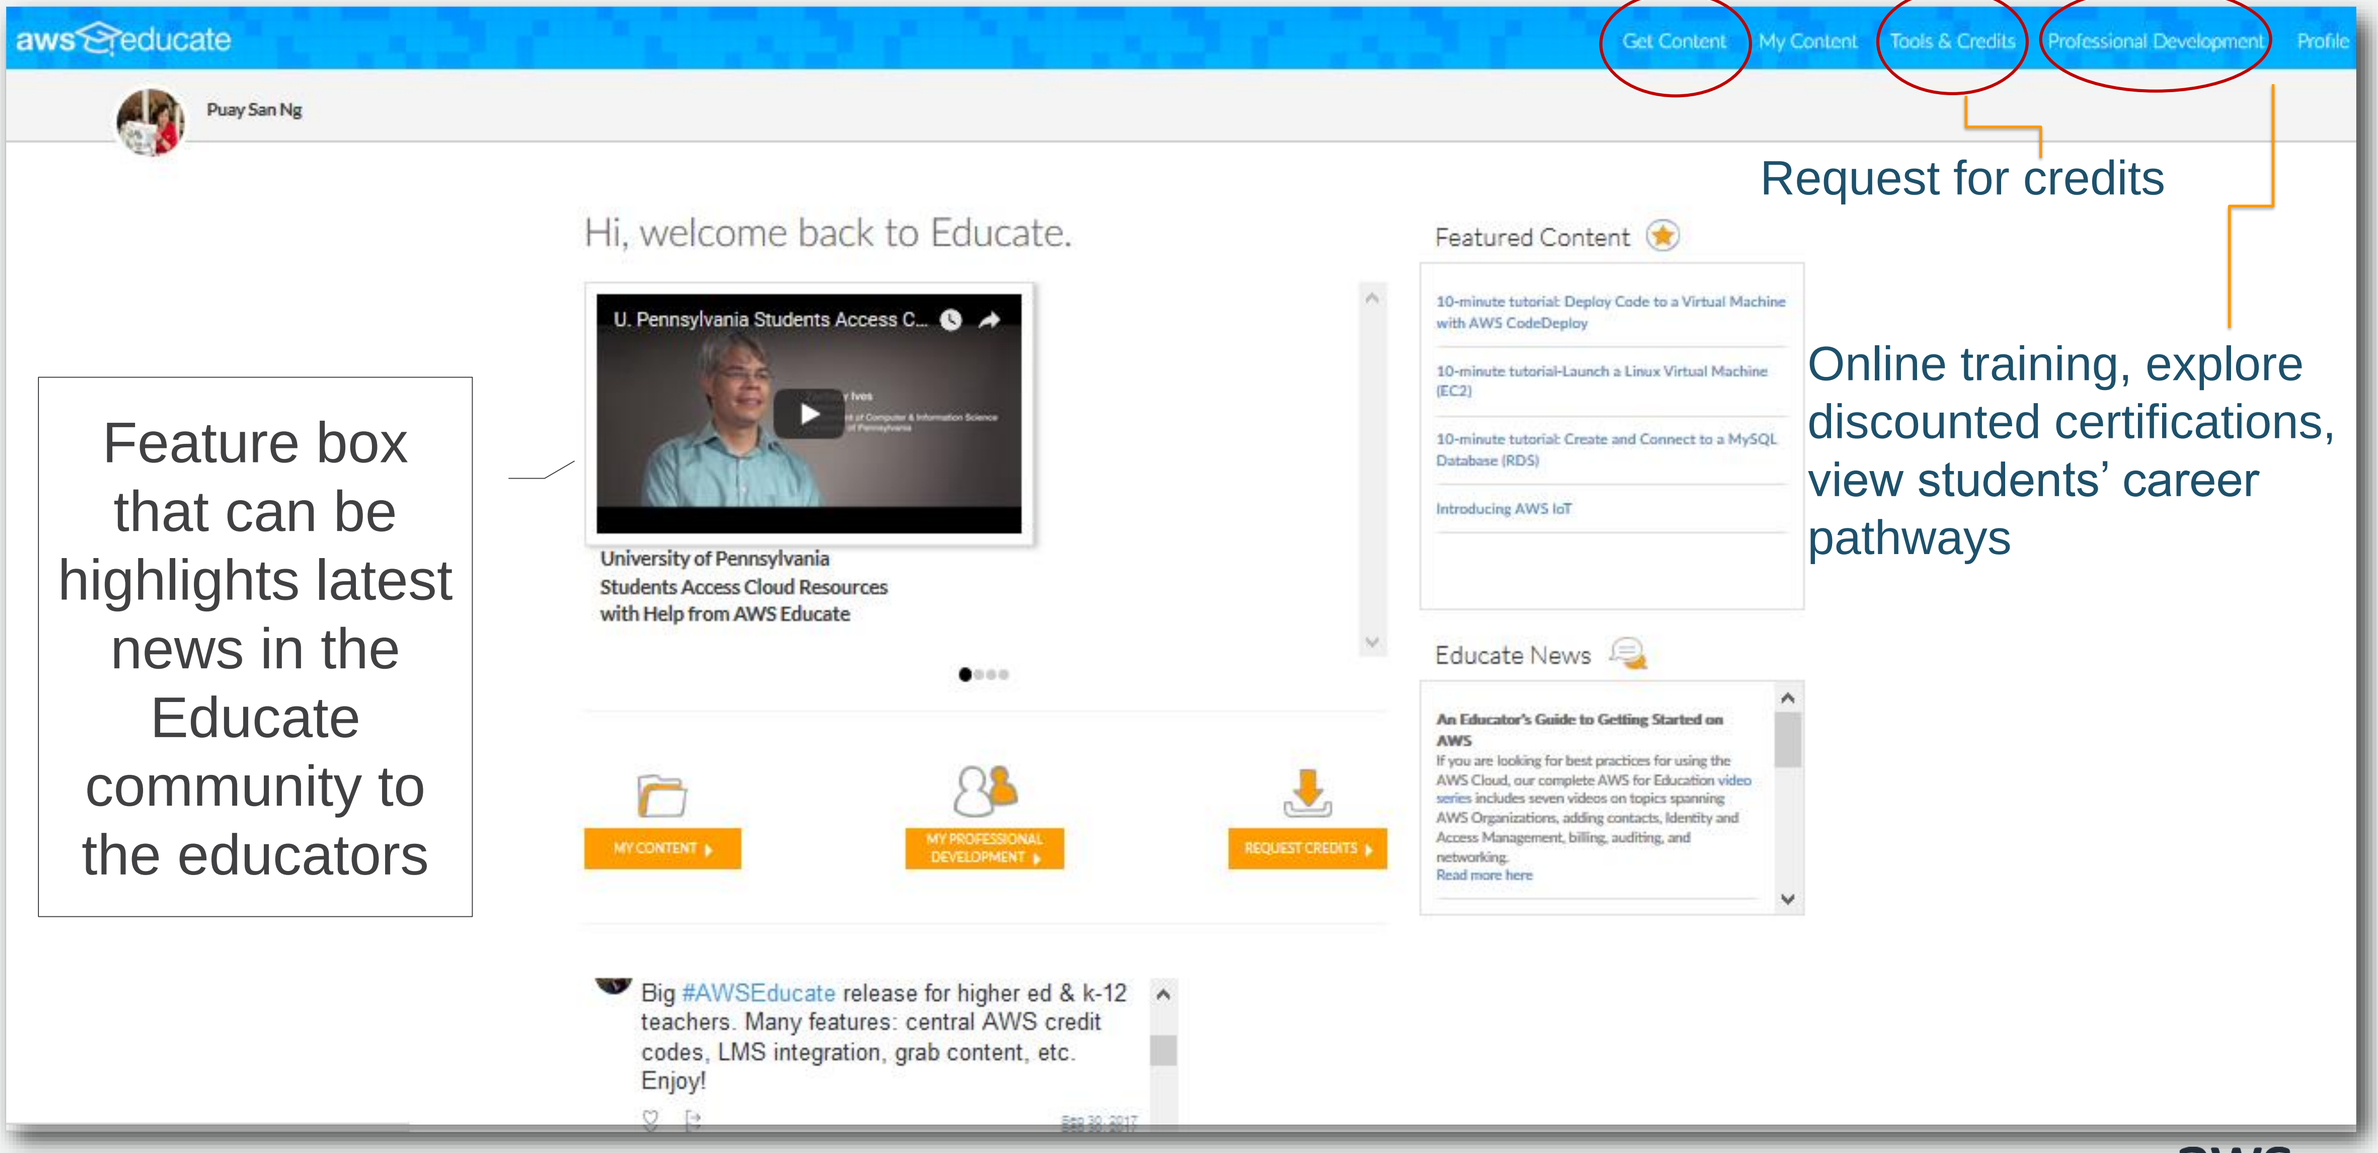

Get Content

Educator Portal allows educators to access teaching content and material from other professors from more than **+1500 institutions and AWS itself**

- Tutorials From AWS
- Slideshare From Other Professors
- Homework/ Lab References
- Content from Udacity and Udemy
- Request Credits for Your Class

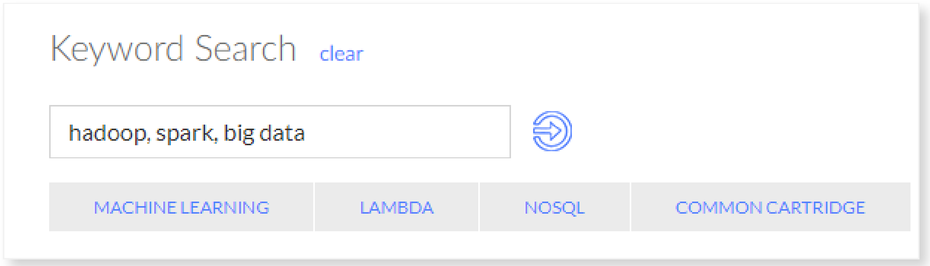

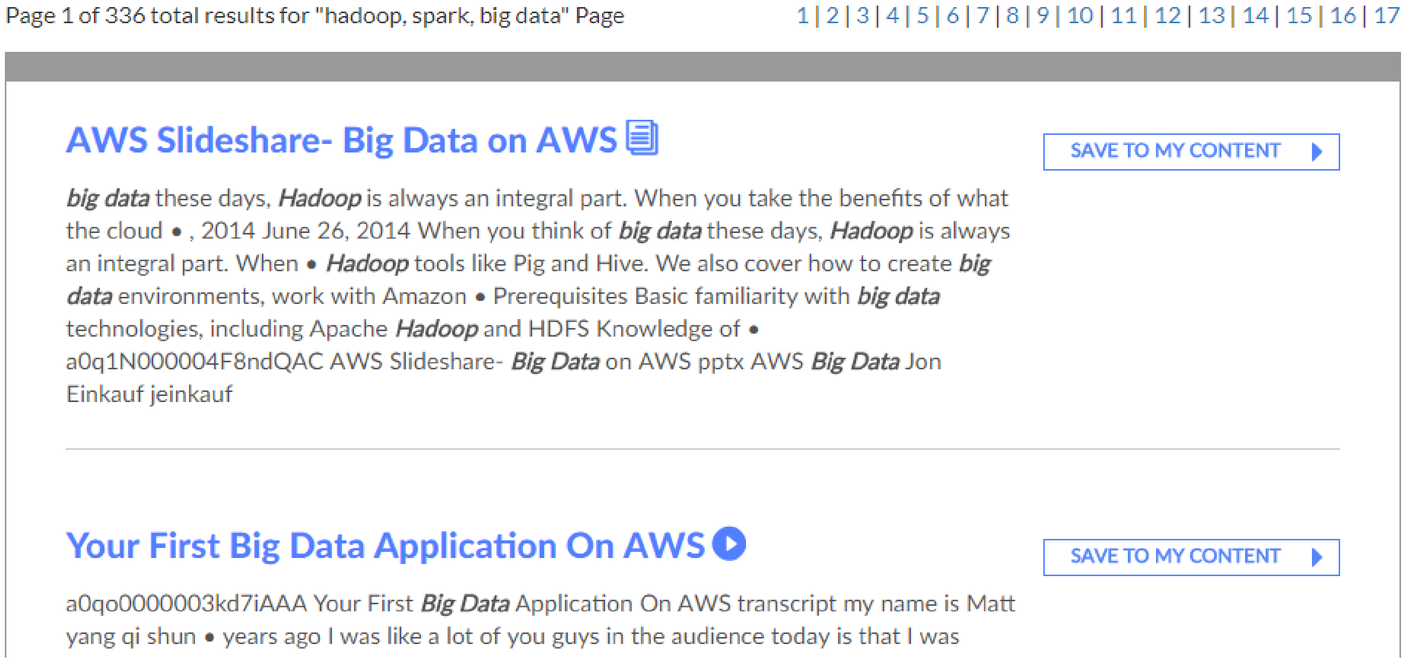

CSE344 Introduction to Data Management: Homework -- Hadoop and Pig

To view and download the file associated with this content, click on the content title below.

CSE344 Introduction to Data Management: Homework -- Hadoop and Pig

Magdalena Balazinska, University of Washington

This 6-page homework assignment, which is the core Amazon Web Services (AWS) assignment for this class, requires students to perform some analysis over a graph drawn from **billion triple dataset**. This is an RDF dataset that contains a billion (add or take a few) triples from the Semantic Web. Some Web pages on the Web have a machine-readable description of their semantics stored as RDF triples: our dataset was obtained by a crawler that extracted all RDF triples from the Web.

The first problem in the assignment gets students started with Pig, by running on an AWS cluster with 10 nodes. The second problem involves the writing of a Pig script that groups tuples by the subject column and creates/stores histogram data showing the distribution of counts per subject, then generating a scatter-plot of this histogram. The third problem involves computing a join on a subgraph. And the fourth problem, for extra credit, involves computing a histogram on the entire dataset.

Associated files:

- **Setup AWS for the Homework**
- Click **Project code archive** to download

AWS Educate Example

- EC2 Set Up Instructions with AWS Educate From Stanford

AWS Tutorial

CS224D Spring 2016

April 17, 2016

1 Introduction

This tutorial explains how to set up your EC2 instance using our provided AMI which has TensorFlow installed. Our AMI is cs224d_tensorflow (ami-d8433cb8). We've installed on it:

- CUDA 7.0
- cuDNN 4.0
- TensorFlow 0.7

2 Create an AWS account and apply for AWS Educate Program

2.1 Create an AWS account

Go to AWS homepage `http://www.aws.amazon.com`. Click the *Sign In to Console* or *Create an AWS account* button on the top right corner. This will bring you to the sign in/sign up page. Create your account there with your email and password.

2.2 Apply for AWS Educate Program

AWS Educate Program gives $35 AWS credits per student. You can apply for it here: `https://aws.amazon.com/education/awseducate/`.

- EC2 Sample HW From Virginia Tech

Virginia Tech.
Computer Science

CS 5614 – (Big) Data Management Systems
Spring 2017, Prakash

Amazon Web Services Setup Guidelines for Homework 3

Goals:

1. Create an AWS instance (to get access to EC2, Elastic MapReduce and S3 storage).
2. Get free credit (and monitor how much you have left)
3. Create storage buckets on S3 (to save outputs and logs of MR jobs)
4. Create a key pair (required for running MR jobs on EC2)
5. Get Access keys (also required for running MR jobs on EC2)
6. Familiarize yourself with S3, EC2 and EMR (by doing a sample MR run)

1. Create an AWS and AWS Educate account

a. Go to https://aws.amazon.com/account/. Click "AWS Free Tier" to sign up for an account, if you don't have one already.
b. Follow all the required steps and enter all the details. Follow the instructions to create your account using the "AWS Free Tier". You'll have to enter payment details (you'll need **a valid credit/debit card**). You'll also have to validate using your phone.
c. Once everything has been verified and created, you should have access to the AWS management Console.
d. After you are signed up, from the drop down next to your name, select My Account. In the Account Settings section, you will notice the Account Id. This is the AWS Account ID to be used next when signing up for AWS Educate.
e. To sign up for AWS Educate and get a $40 credit go to:
http://aws.amazon.com/education/awseducate
Click on the Apply Now button. Click on the Apply for AWS Educate for Students button. On Step 1, select I am a Student. On Step 2, fill out the form appropriately. In the AWS Account ID field, enter the ID from your initial AWS signup. Click Next and finish the sign up process.

After your application is reviewed and approved, you will receive a welcome e-mail from AWS Educate Support, which includes details about the $40 promotional credit, as shown below in the example (so **do this setup EARLY!**):

AWS Promotional Credit
It's our pleasure to issue you an Amazon Web Services (AWS) promotional credit code in the amount listed below.

AWS Educator Guide

- DEEPLENDS PROJECT

aws educate

Course Project

EDUCATOR GUIDE

Title: Animal Recognition (object detection) with DeepLens

There is no shortage of challenges and issues that can arise on the job. Having the necessary knowledge and skills to identify solutions to complex problems is one of the top skills employers look for in employees especially in cloud roles. This course project simulates real-world tasks encountered by cloud professionals and challenges students to apply their AWS knowledge and technical skills to evaluate and recommend an efficient, secure and scalable solution.

Purpose:

To simulate real-world practice by providing a complex, authentic task designed to challenge students to apply their AWS and technical knowledge, critical thinking and problem solving skills in a real-world context.

Objectives:

The objective(s) of the course project are:

- Identify the most efficient solution to address the challenge;
- Compare and contrast alternative AWS solutions and services;
- Determine and apply the appropriate AWS architecture to create a stable, fault tolerant environment; and
- Summarize in writing the strengths and challenges of proposed solution(s).

- Associate Cloud Archetiect

aws educate

Course Project

EDUCATOR GUIDE

Title: Regions and Availability Zones

There is no shortage of challenges and issues that can arise on the job. Having the necessary knowledge and skills to identify solutions to complex problems is one of the top skills employers look for in employees especially in cloud roles. This course project simulates real-world tasks encountered by cloud professionals and challenges students to apply their AWS knowledge and technical skills to design an efficient, secure and scalable solution.

Purpose:

To simulate real-world practice by providing a complex, authentic task designed to challenge students to apply their AWS and technical knowledge, critical thinking and problem solving skills in a real-world context.

Objectives:

The objective(s) of the course project are:

- Identify the most efficient solution to address the computing challenge;
- Compare and contrast alternative AWS solutions;
- Determine and apply the appropriate AWS architecture to create a stable, fault tolerant environment; and
- Summarize in writing the strengths and challenges of proposed solution(s).

Classroom

AWS Educate Classroom is an **educational cloud lab environment** provided by Vocareum. Through AWS Educate classroom, Educators can experience secured lab environment with **50 more promotional credits** by default.

This classroom offers **teacher ability to get a full picture of every students credits usage and services used** (Activity or Resource Report) or **accessing individual students work area** (AWS Console)

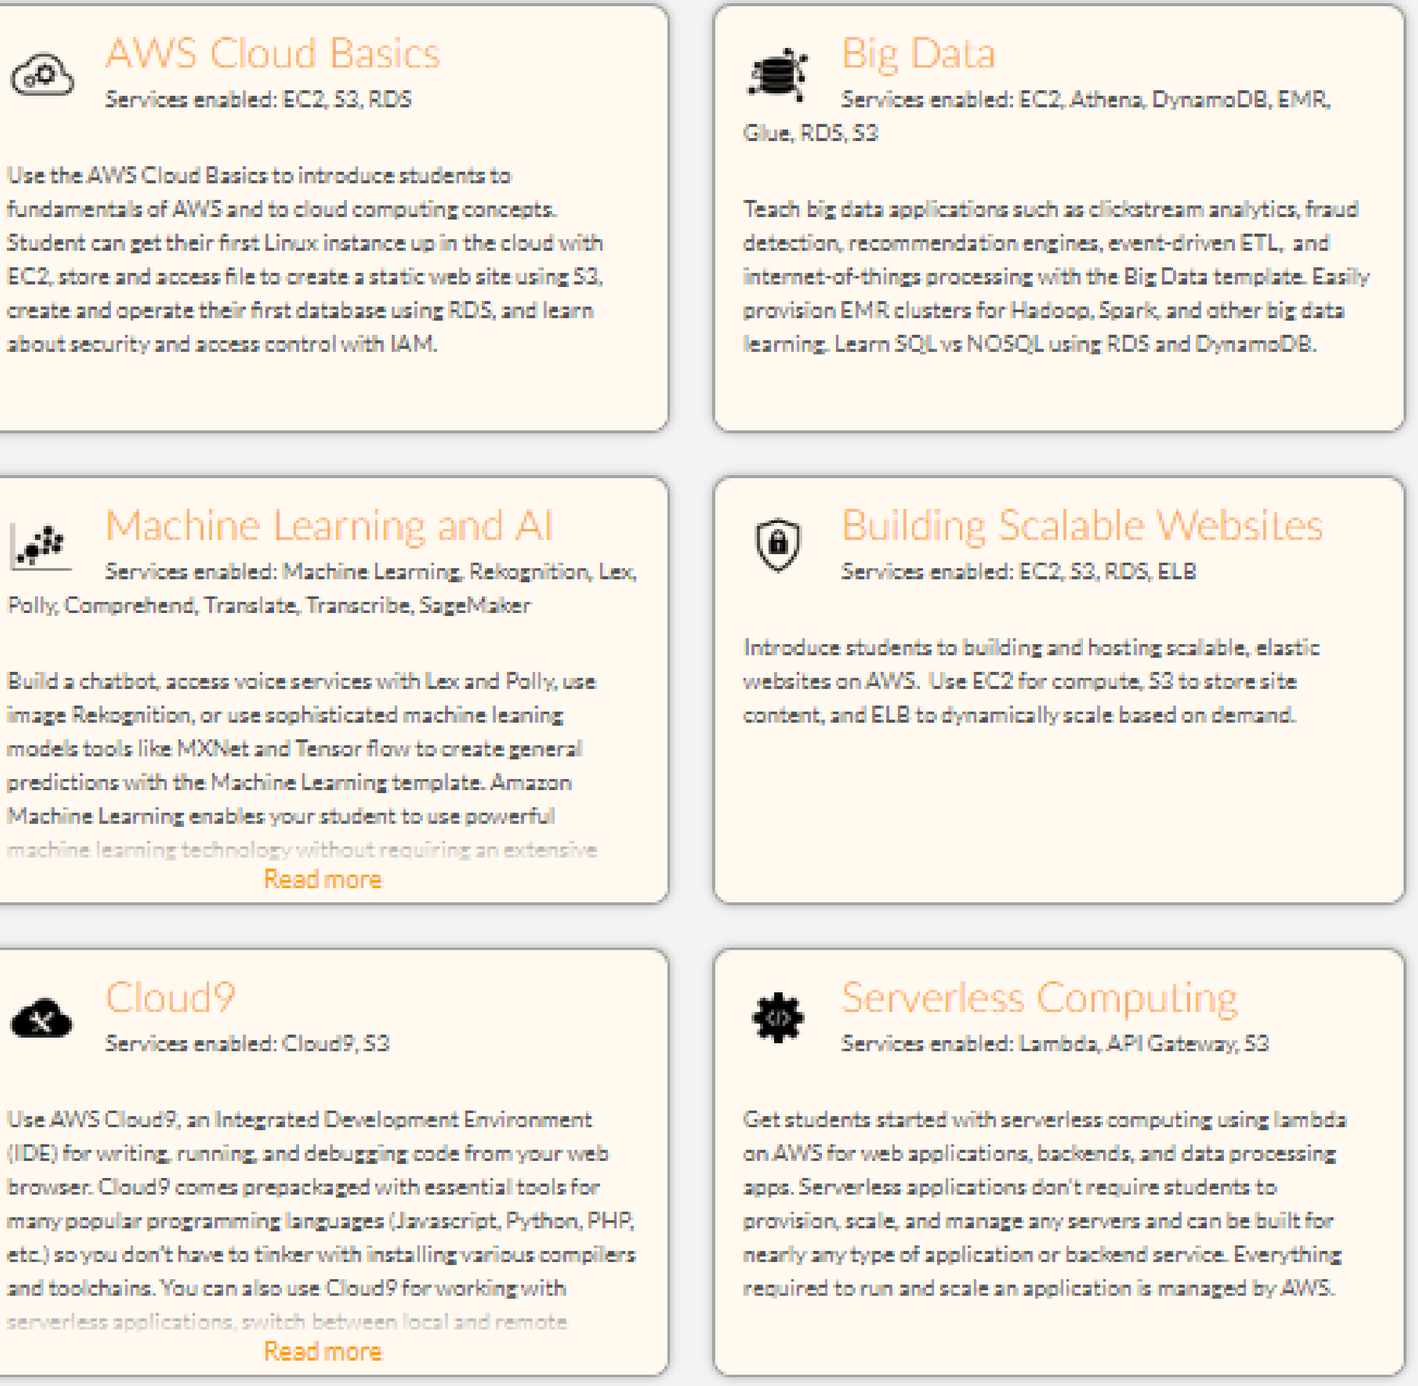

Used by educational institutions worldwide

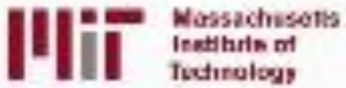

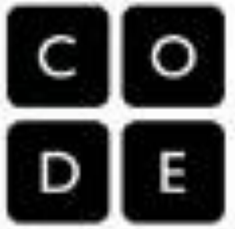

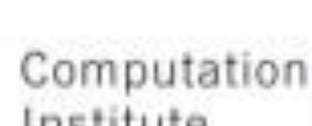

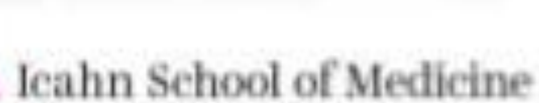

什麼是 AWS Academy 計劃

aws academy

為高等教育機構提供AWS課程，以協助學生開發所需的雲端運算技能，並做好取得業界認可AWS認證的準備

教育機構和教育工作者可獲得完整的課程，該課程的設計可輕鬆整合到現有的教育計劃中

AWS 負責建立和更新課程內容，讓教育工作者有更多時間專注在學生上

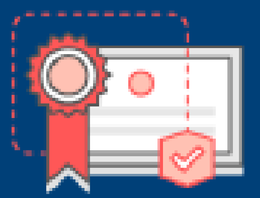

教育工作者成員可透過講師認證獲得免費 AWS 培訓及 AWS 認證的折扣

aws academy

AWS Public Data Sets

提供兩年資料儲存及資料下載的費用

We work with data providers who seek to:

- Democratize access to data by making it available for analysis on AWS.
- Develop new cloud-native techniques, formats, and tools that lower the cost of working with data.
- Encourage the development of communities that benefit from access to shared datasets.

Putting Data to AWS

Examples of how data shared on AWS is accelerating research and creation of new applications.

Visit **Earth on AWS** to learn about building planetary-scale applications in the cloud with open geospatial data.

All public data from the **Hubble Space Telescope**'s active instruments are available for large-scale analysis on Amazon S3.

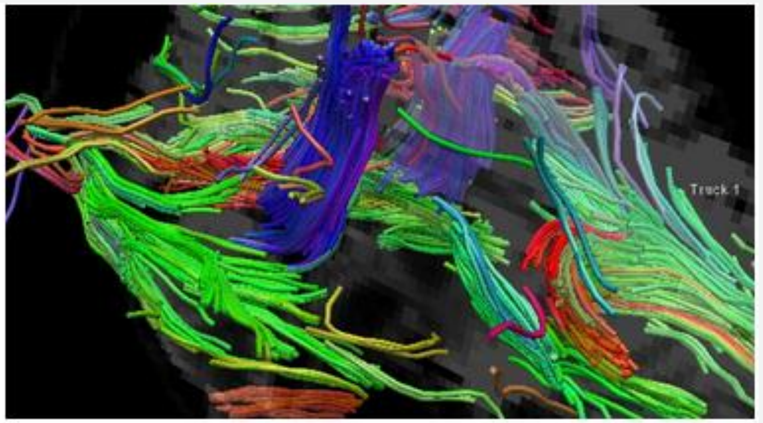

The **Allen Institute for Brain Science** and the **University of Washington** provided students with 35TB of data with Amazon S3.

You can query billions of **OpenStreetMap** features with Amazon Athena without needing to download data or set up a server.

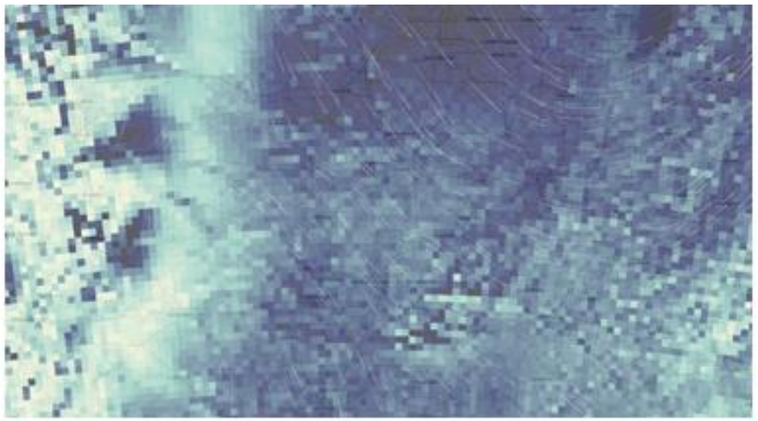

The **National Renewable Energy Laboratory (NREL)** makes a 500 TB open weather model dataset available to the world on Amazon S3.

```
      sv.referenceallele ,
      sv.alternateallele ,
      sv.genotype0 ,
      sv.genotype1
FROM demo.samplevariants sv
CROSS JOIN
  (SELECT count(1) AS numsamples
   FROM
     (SELECT DISTINCT sampleid
      FROM demo.samplevariants
      WHERE sampleid LIKE 'NA12%'))
JOIN demo.clinvar cv ON sv.chromosome = cv.chromosome
AND sv.startposition = cv.startposition - 1
AND sv.endposition = cv.endposition
AND sv.referenceallele = cv.referenceallele
AND sv.alternateallele = cv.alternateallele
WHERE assembly='GRCh37'
  AND cv.clinicalsignificance LIKE '%response%'
```

Learn how to prepare **1000 Genomes** data for fast interactive analysis using Amazon Athena.

AWS Credits for Research

IT'S ABOUT SCIENCE, NOT SERVERS.

We recognize that whilst research is often a compute-intensive activity, most researchers are not IT experts.

We want to simplify research in the cloud with easy-to-use tools for researchers and their students, and share the catalogue of "researcher-obsessed" products and services created by many of our partners.

aws.amazon.com/rcp

The Cancer Genome Atlas (TCGA)

NIH National Institutes of Health

- Biospecimen Repositories
- Genomics Data Analysis Centers
- Analysis Working Groups
- Proteome Characterization Centers
- Genome Characterization Centers
- Data Coordinating Center

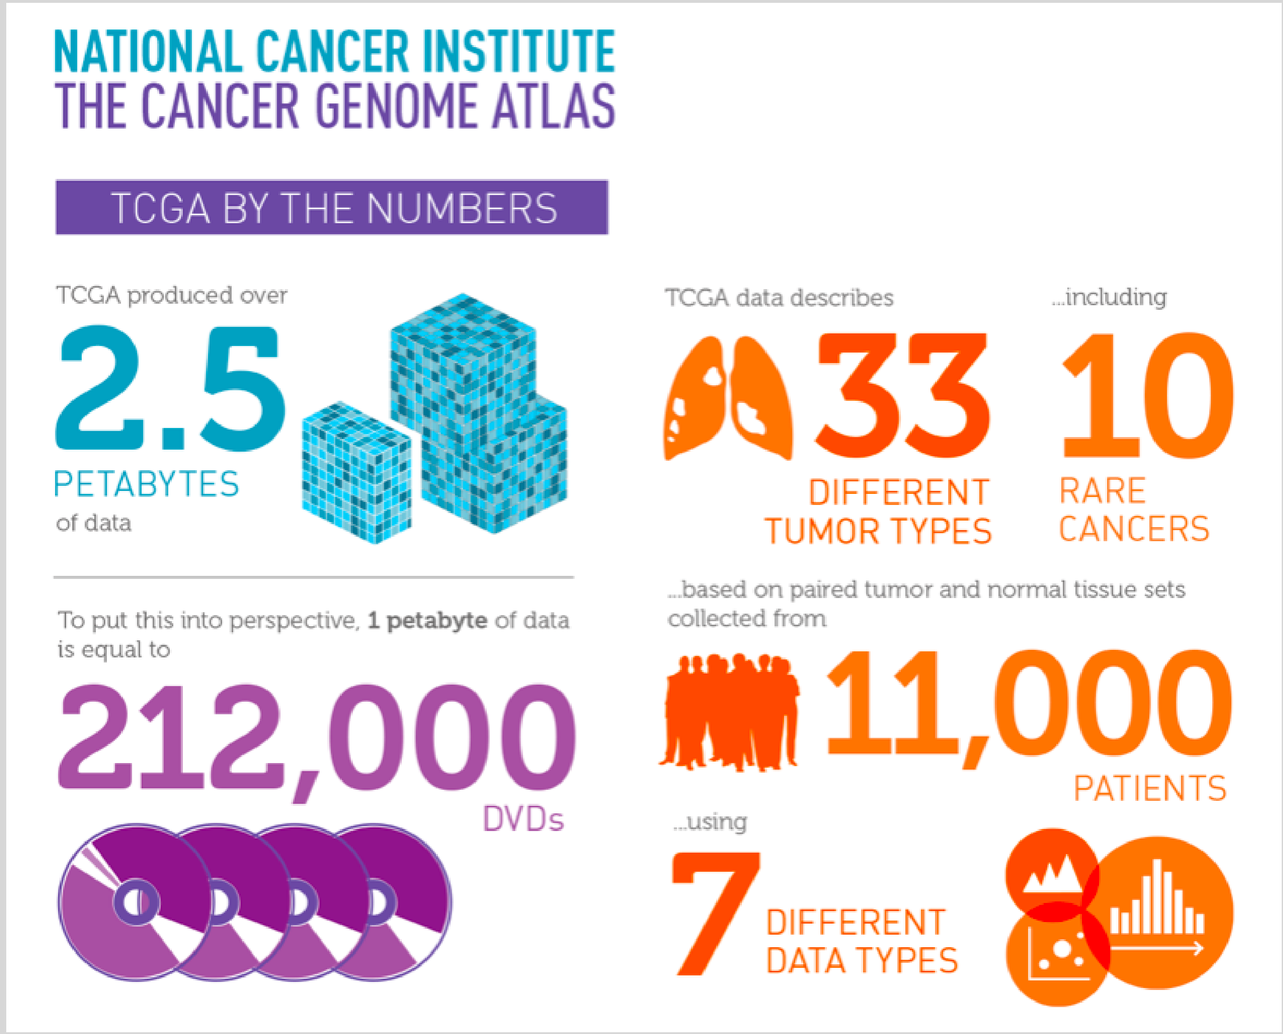

National Science Foundation
WHERE DISCOVERIES BEGIN

Forbes

FEB 18, 2018 @ 02:42 PM 1,933

The Little Black Book of Billionaire Secrets

NSF's New Initiative To Bring The Cloud Era To Academic Big Data Research

Amazon Web Services
@awscloud
Follow

Researchers awarded to push boundaries of modern science using AWS via @NSF BIGDATA program. amzn.to/2nKsbpw

aws Research Initiative

NSF Partners With 3 Cloud Providers for Data Science Research Support Program

The National Science Foundation has partnered with Amazon Web Services, Microsoft and Google to support research projects in the data science and engineering field. NSF said Wednesday it will obligate nearly $30 million to the Critical Techniques, Technologies and Methodologies for Advancing Foundations and Applications of Big Data Sciences and Engineering program...

Read More

Collaborative program with the National Science Foundation (NSF)

- The big data program supported by multiple directorates at NSF, provides funds up to $26.5 million in addition to Cloud Credits to perform **cutting edge big data research on cloud for a period of 3-4 years** (up to 2021)
- Big Data Award (2017): Out of 8 awards – 5 were awarded to researchers using AWS for Research

aws Research Initiatives - https://amzn.to/2GVxx9a

Natural Language Processing at Clemson University – 1.1 Million vCPUs & EC2 Spot Instances

- 1.1M vCPUs at peak execution in a single AWS Region
- Total of 1,832,923 vCPU hours
- 2 hours of run time
- $32,423 total execution cost
- **Average of 1.72¢ per core hour**

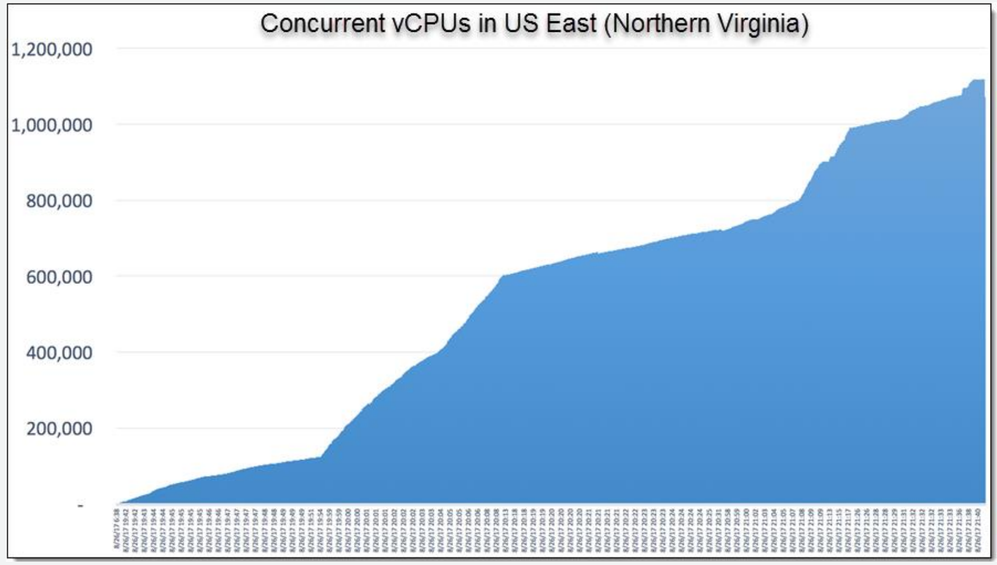

I am absolutely thrilled with the outcome of this experiment. The graduate students on the project are amazing. They used resources from AWS and Omnibond and developed a new software infrastructure to perform research at a scale and time-to-completion not possible with only campus resources. Per-second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

GPU (nVidia DGX-1 (v100) vs AWS P3 (v100)

nVidia DGX-1 (w/ 8 v100) 建議售價USD149,000

- 每天24小時不眠不休執行, 可以執行p3.2xlarge 371天 (RI)
- 如果每天執行8小時, 可以執行p3.2xlarge 1,383天 (Spot Block)
- 如果每天執行3小時, 可以執行p3.2xlarge 4,479天 (Spot)

Instance Size	GPUs - Tesla V100	GPU Peer to Peer	GPU Memory (GB)	vCPUs	Memory (GB)	Network Bandwidth	EBS Bandwidth
p3.2xlarge	1	N/A	16	8	61	Up to 10 Gbps	1.5 Gbps
p3.8xlarge	4	NVLink	64	32	244	10 Gbps	7 Gbps
p3.16xlarge	8	NVLink	128	64	488	25 Gbps	14 Gbps

AWS Activate Program

Amazon Web Services 為新創公司提供擴展各種業務規模所需之低成本、易用型基礎設施。全球最熱門的新創公司，包括 Airbnb、Slack 和 Robinhood，都利用 AWS 的強大功能來迅速擴展事業版圖。AWS Activate 是為您的新創公司提供所需資源的一個計劃，讓您開始利用 AWS，躋身全球發展最快速的新創公司之列。

Portfolio / Portfolio Plus 套裝服務

AWS Activate 為精選 Accelerator、種子/創投基金等組織中之符合資格的新創公司提供客製化的 Portfolio 和 Portfolio Plus 套裝服務：

最高 **15,000 USD** 的 AWS 促銷積分 (抵扣卷)，可供使用 2 年 (金額依組織而有不同) 或 100,000 USD 可供使用 1 年 (此選擇僅在 Portfolio Plus 套裝)

最高 10,000 USD 的 AWS Support 積分，可用於 AWS Support (商業計劃) 2 年
AWS Business Essentials 線上培訓或面授培訓 (價值 600 USD)
AWS Technical Essentials 線上培訓或面授培訓 (價值 600 USD)
80 點的自主進度實驗室積分 (價值 80 USD)

AWS Market Place

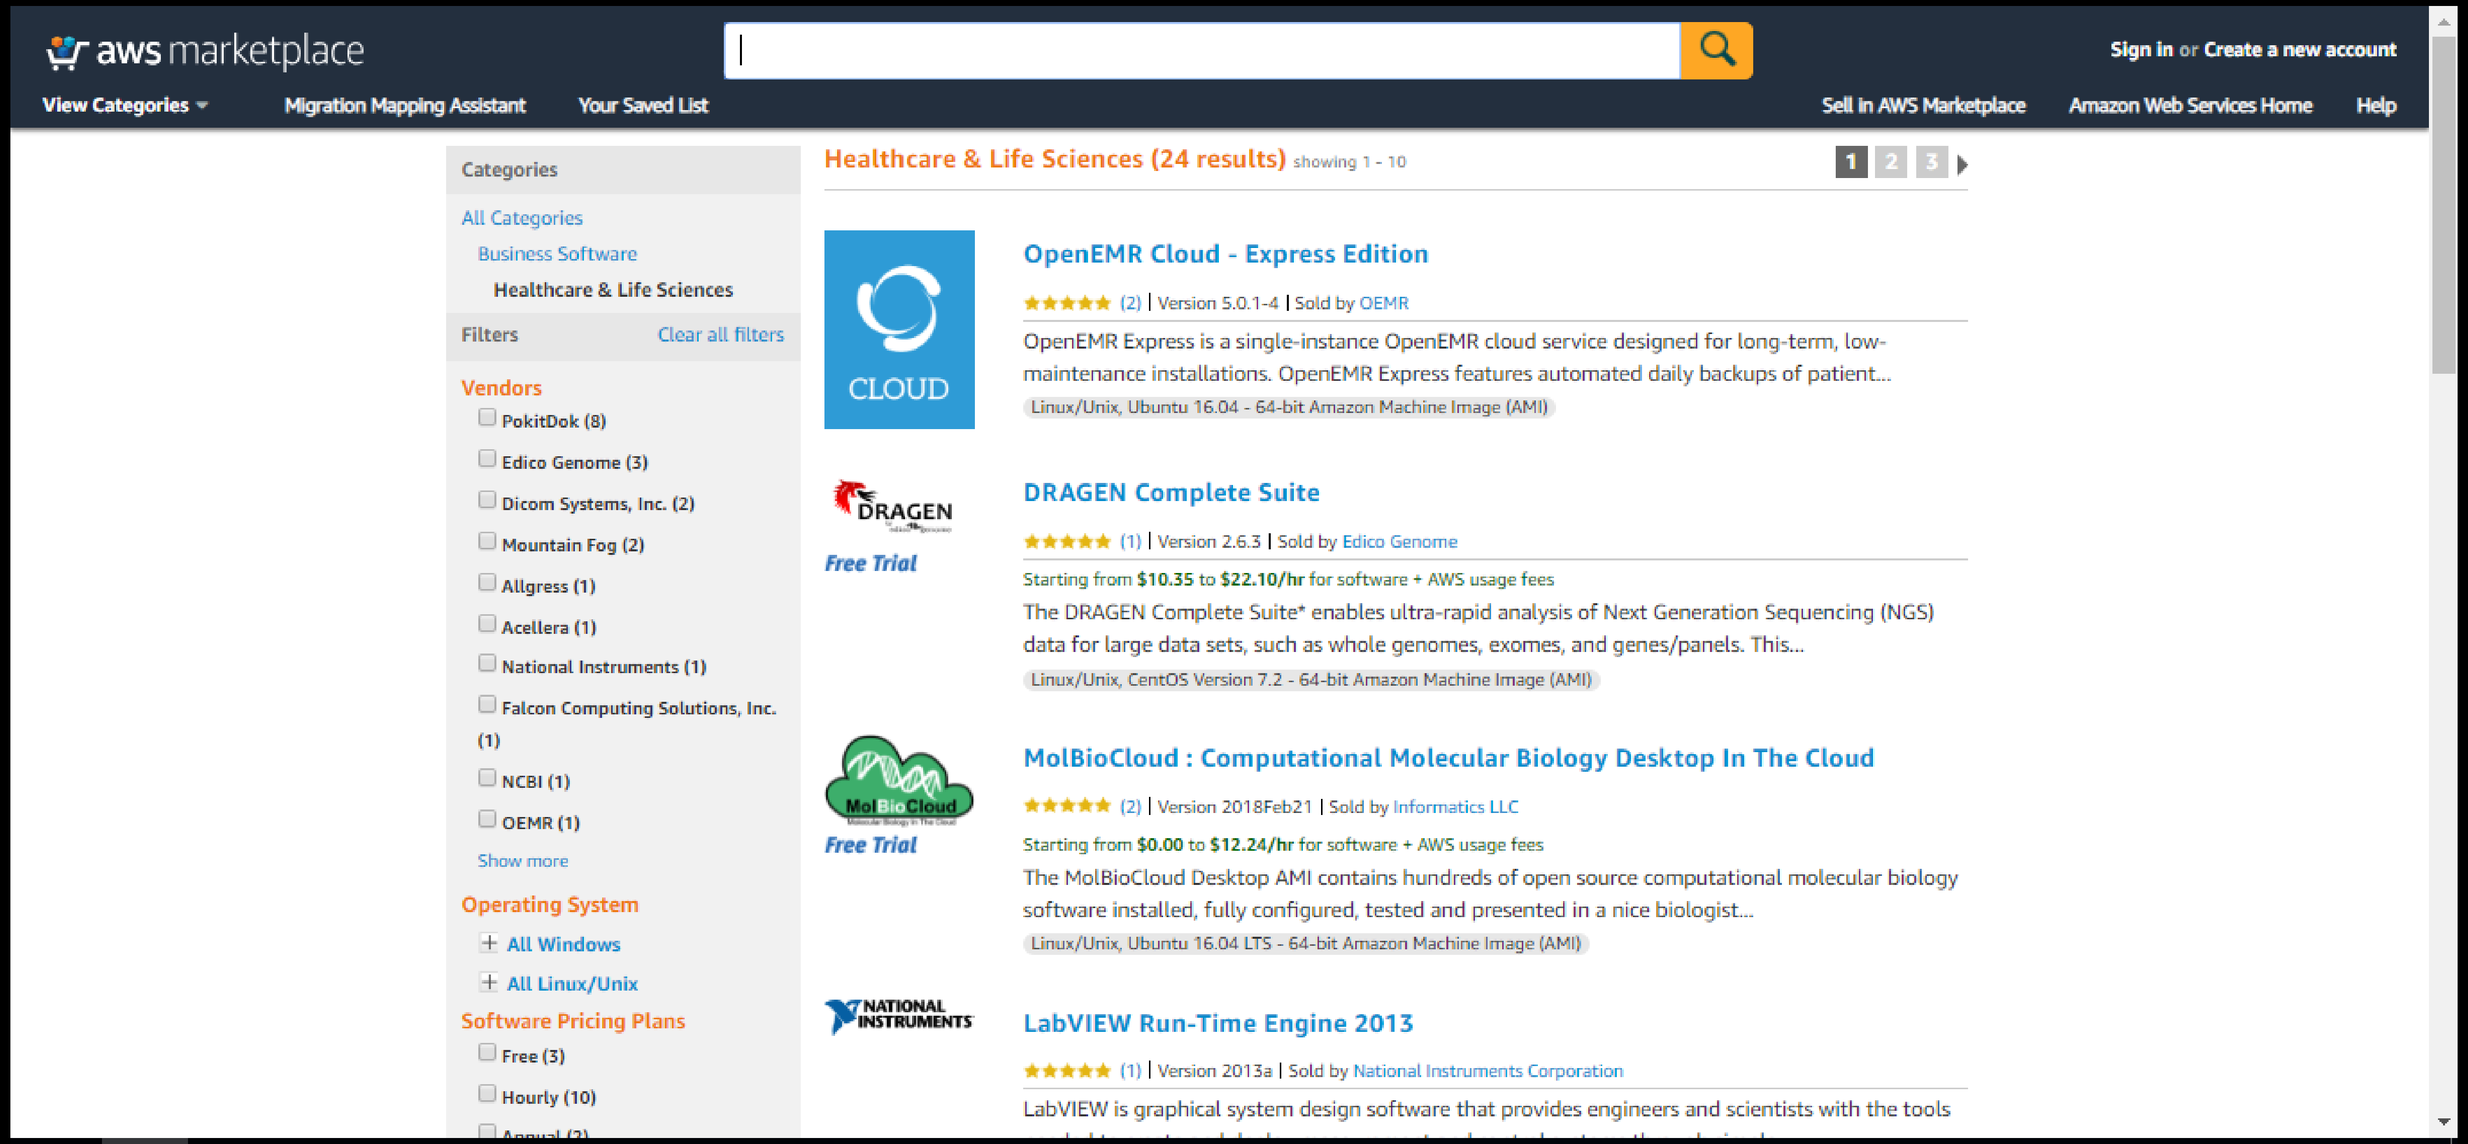

Partner

- 伊雲谷 eCloudValley https://www.ecloudvalley.com/zh-hant/
  - 陳心妍 Mandy Chen

    mandy.chen@ecloudvalley.com

    手機: 0910236963

- 博弘 Nextlink https://www.nextlink.com.tw/
  - 劉取哲 Fred Liu

    Fred.liu@nextlink.com.tw

    手機: 0910676501

Thank You

Appendix

全球雲端人才需求若渴

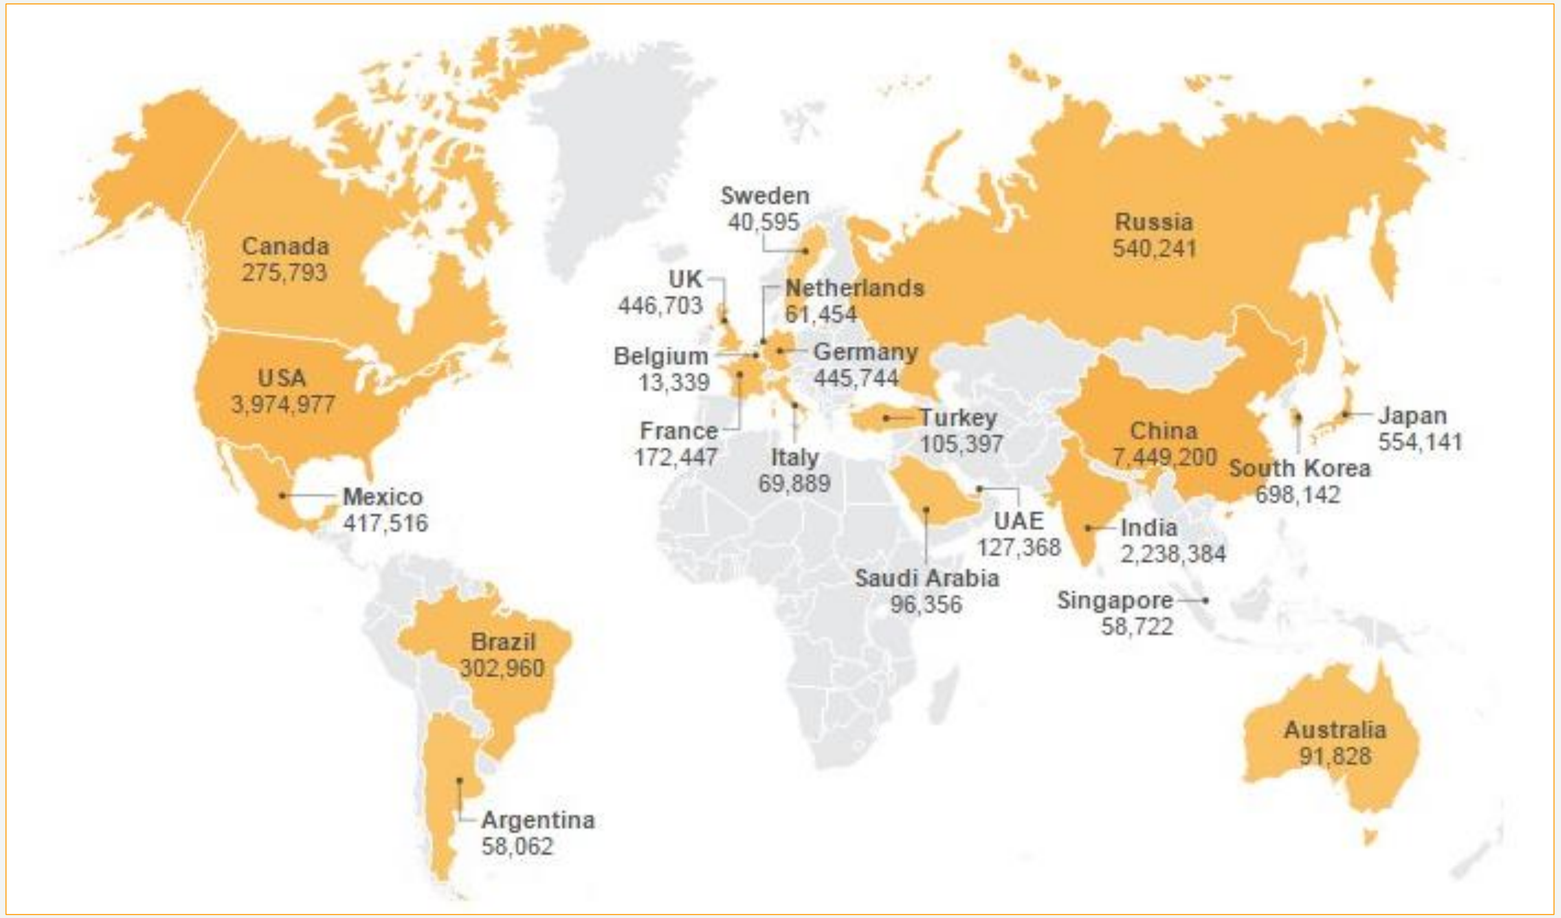

Source: WANTED Analytics

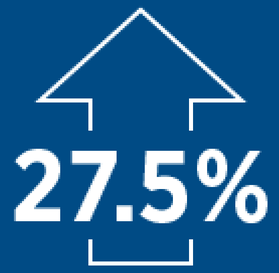

The average salary of AWS-certified IT professionals is 27.5% higher than the norm.

AWS-certified IT decision-makers and professionals make an average of $113,932.

AWS Certified Solutions Architect (Associate) is the third highest paying certification in the U.S. and Canada.

Certification Roadmap

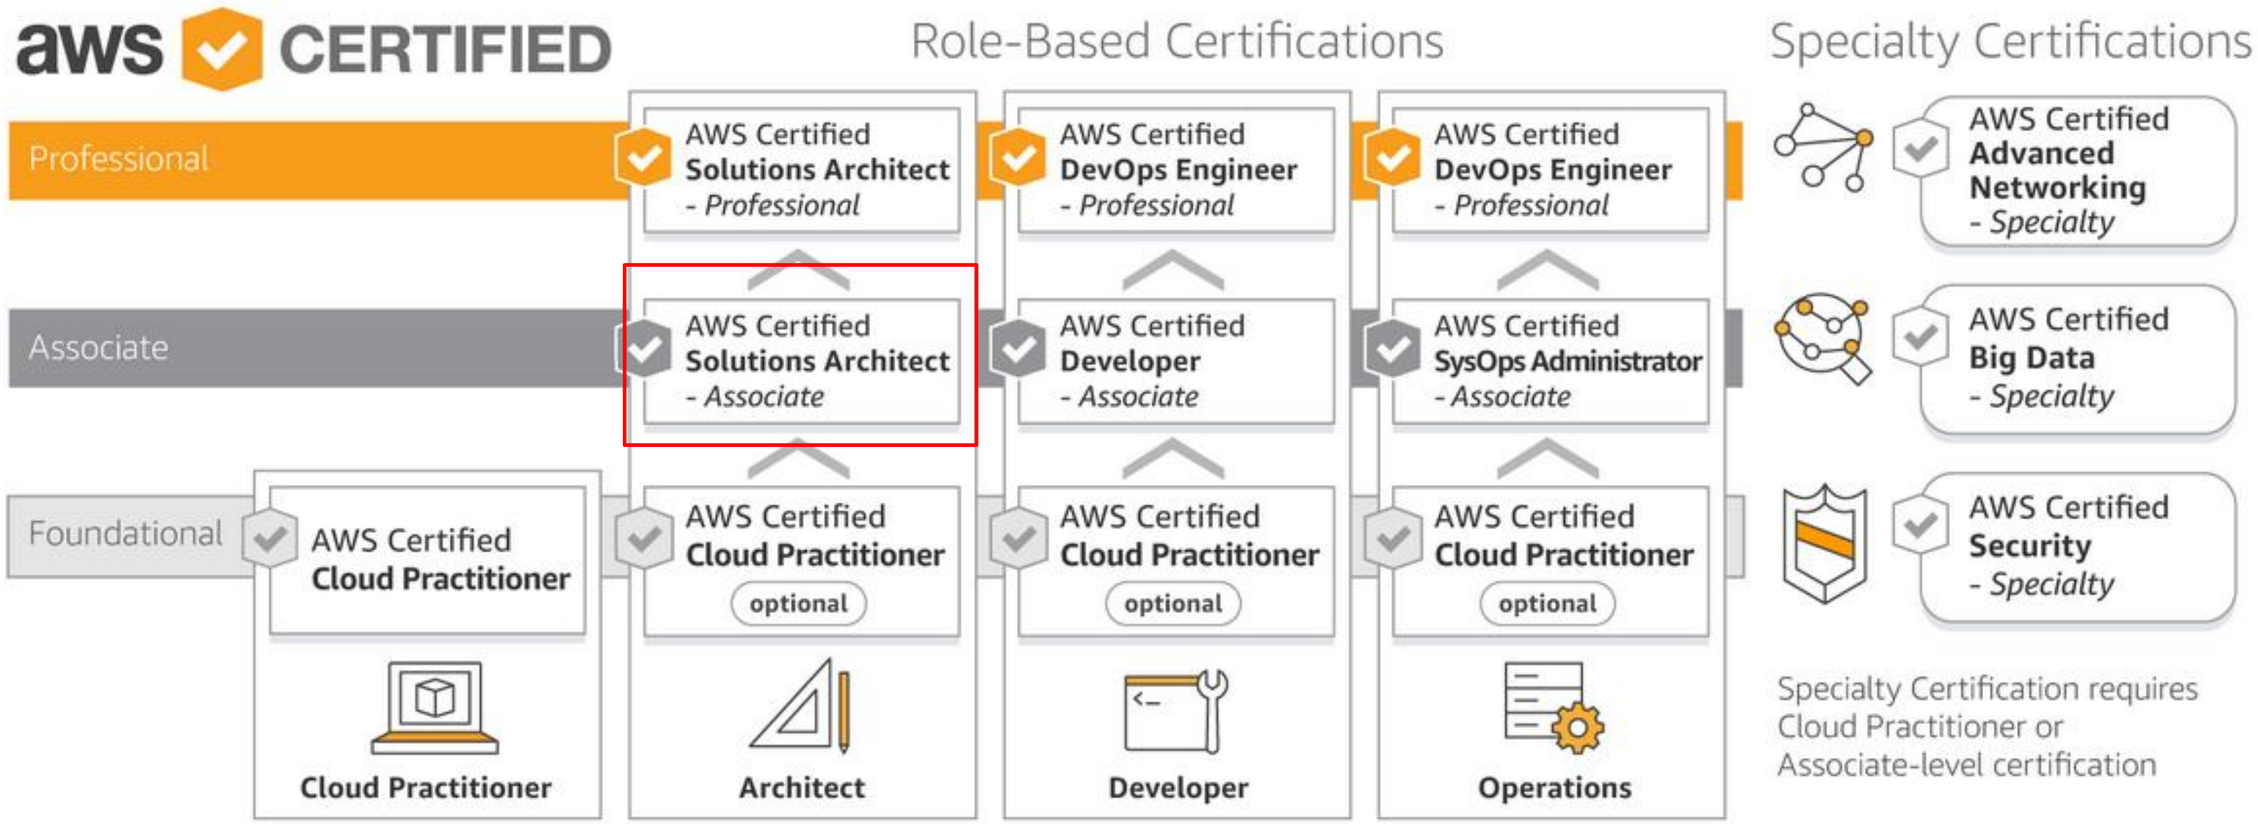

AWS Research Cloud Program

Science first, not servers. Researchers are not professional IT people (nor do they wish to be).

Simple and easily explained procedures to get set up with cloud access.

Budget management tools to ensure that over-spends do not happen.

Best practices to ensure both data and research budgets are safe and privacy is protected.

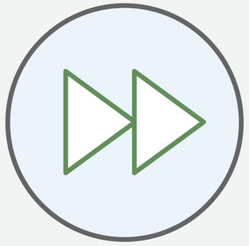

Fast track to invoice-backed billing & Egress Waiver.

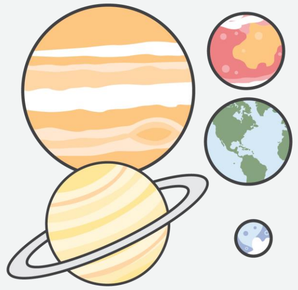

Large catalog of scientific Solutions from partners, including instant clusters from AWS Marketplace.

Defining HPC – example use cases

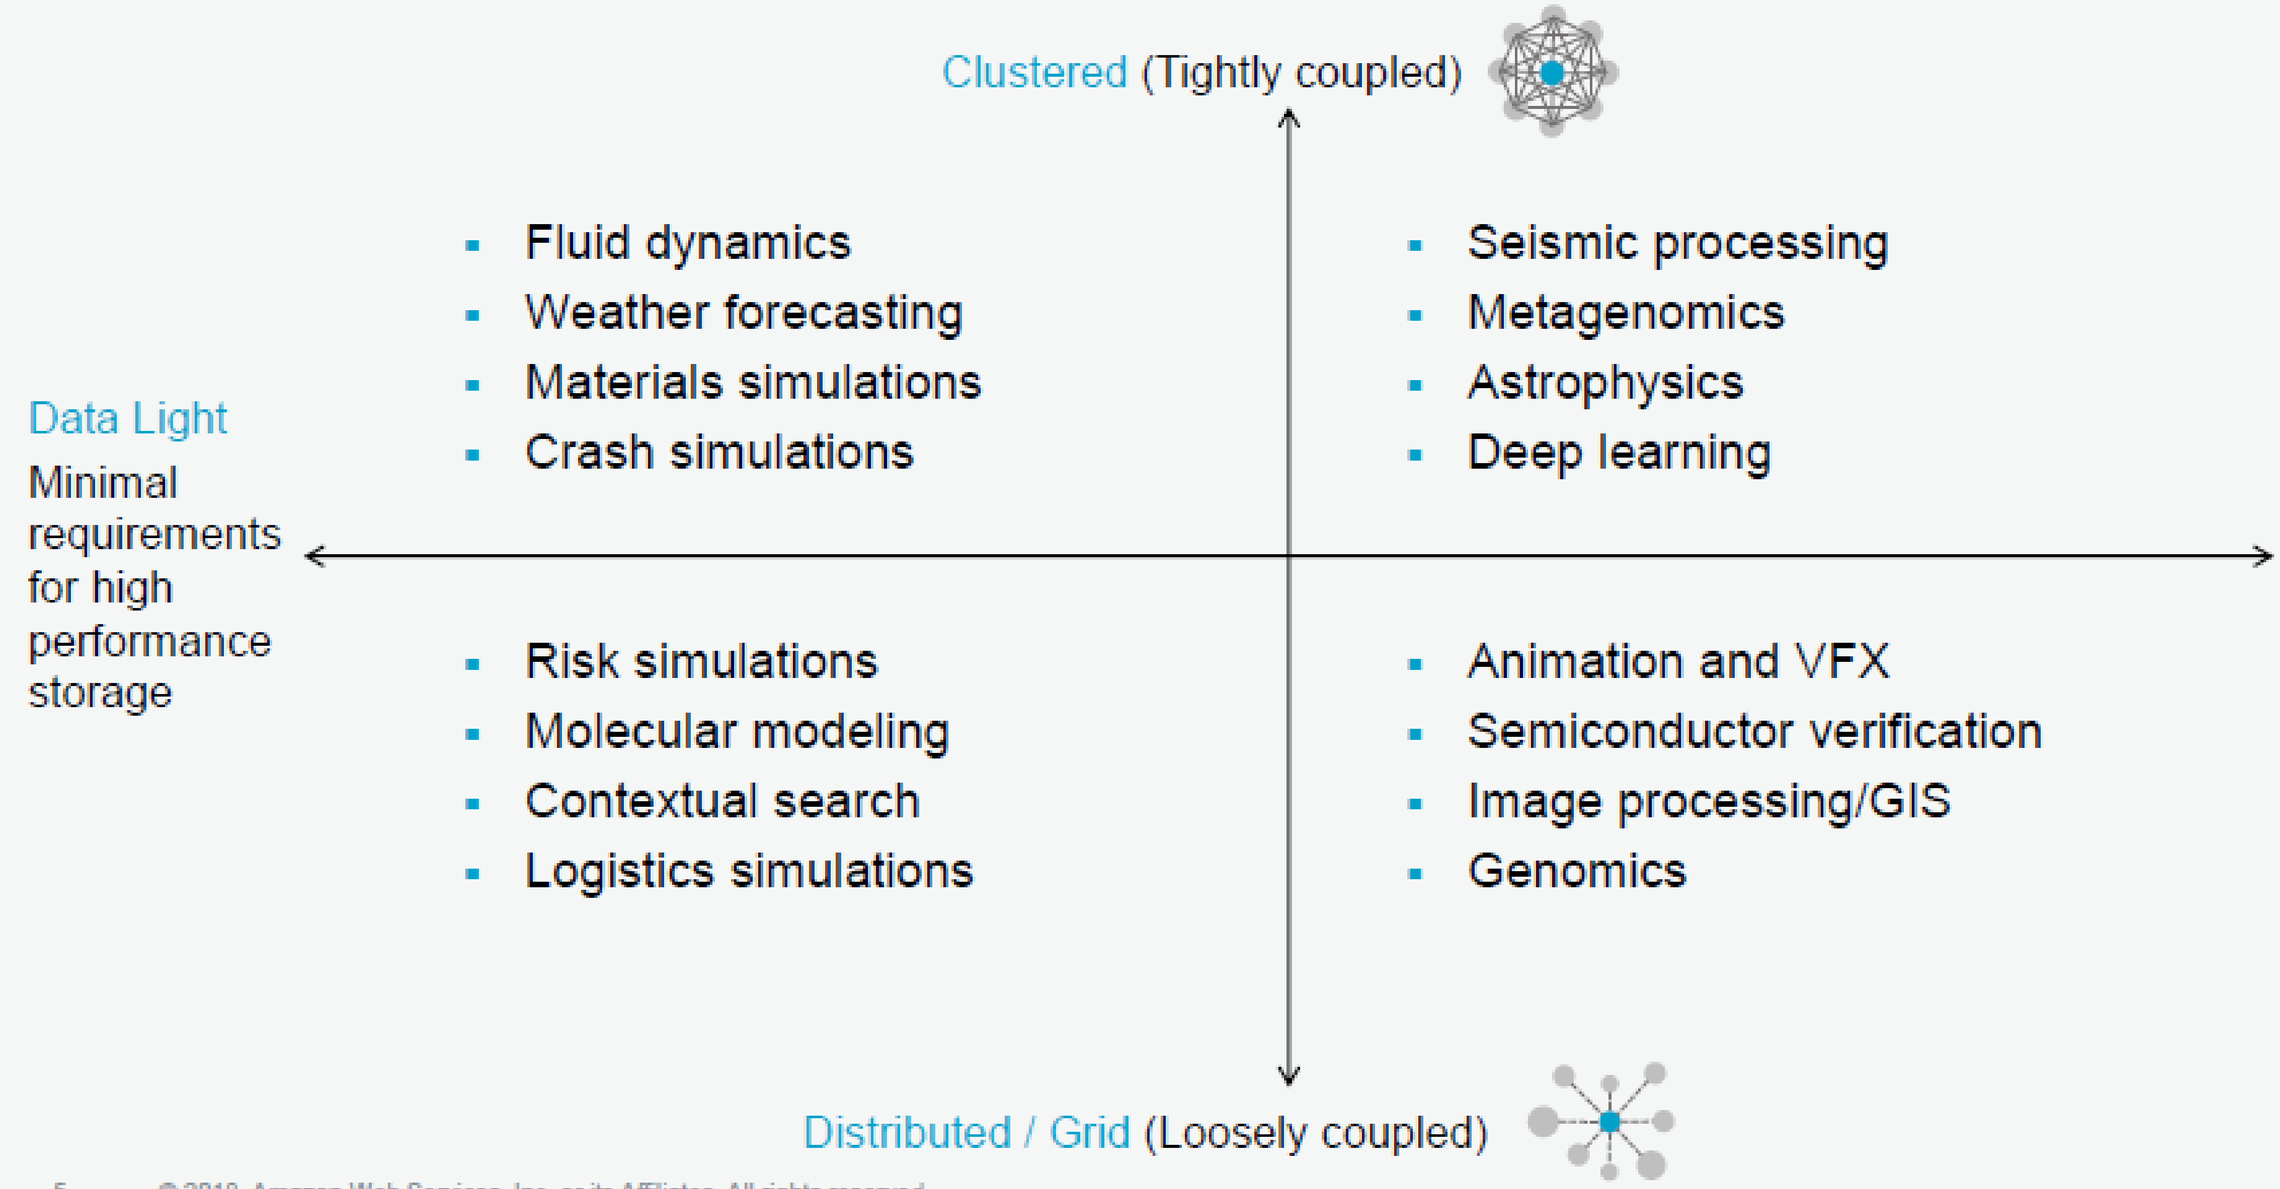

Important factors for application performance

- Compute performance – CPUs, GPUs, FPGAs
- Memory performance – high RAM requirements in many applications
- Network performance – throughput, latency, and consistency
- Storage performance – including shared filesystems
- Automation and cluster/job management
- Remote graphics for interactive applications
- ISV support – including license management

…and SCALE

Parallel processing: GPU

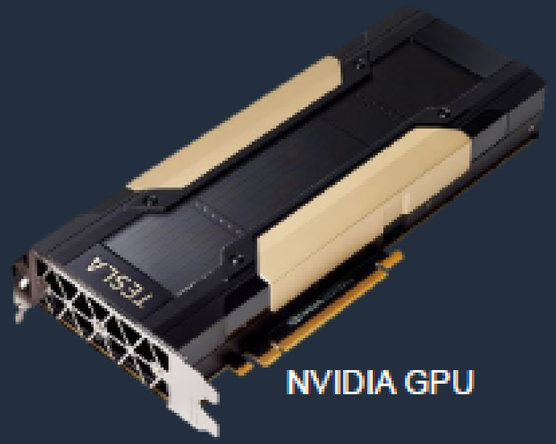

NVIDIA GPU

A GPU is effective at processing the same instruction in parallel, for example calculating pixel values in parallel for graphics shading, or running many parallel financial computations.

A GPU has a well-defined instruction-set, and fixed word sizes.

FPGA

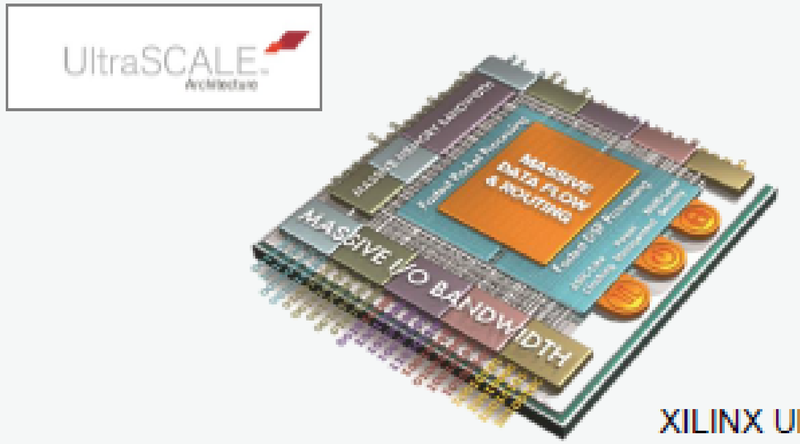

XILINX ULTRASCALE+ FPGA

An FPGA is effective at processing the same or different instructions in parallel, for example creating a complex pipeline of parallel, multistage operations on a video stream, or performing a sequence of dependent calculations and data manipulations for genomics processing.

An FPGA does not have a predefined instruction-set, or a fixed data width.

Questions

?

Thank you

aws | Professional Services