



Machine Learning with AWS

Christian Petters, Solutions Architect, Amazon Web Services

27.09.2018

At Amazon,
we've been making
investments in ML for
the last 20+ years...



Fulfillment & logistics



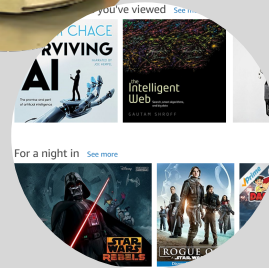
amazon
fulfillment

The image shows the interior of an Amazon fulfillment center. In the foreground, a white pillar with a red top section features the Amazon logo and the word "fulfillment". To the right, a tall metal shelving unit is filled with stacks of brown cardboard boxes. In the aisle between the shelving units, several orange and black Kiva mobile robots are visible, some moving along the floor. The background shows more shelving units and the structural steel beams of the warehouse ceiling.

At Amazon,
we've been making
investments in ML for
the last 20 years...



Fulfillment & logistics



Search & discovery



Welcome to Amazon.com Books!

*One million titles,
consistently low prices.*

(If you explore just one thing, make it our personal notification service. We think it's very cool!)

SPOTLIGHT! -- AUGUST 16TH

These are the books we love, offered at Amazon.com low prices. The spotlight moves EVERY day so please come often.

ONE MILLION TITLES

Search Amazon.com's [million title catalog](#) by author, subject, title, keyword, and more... Or take a look at the [books we recommend](#) in over 20 categories... Check out our [customer reviews](#) and the [award winners](#) from the Hugo and Nebula to the Pulitzer and Nobel... and [bestsellers](#) are 30% off the publishers list...

EYES & EDITORS, A PERSONAL NOTIFICATION SERVICE

Like to know when that book you want comes out in paperback or when your favorite author releases a new title? Eyes, our tireless, automated search agent, will send you mail. Meanwhile, our human editors are busy previewing galleys and reading advance reviews. They can let you know when especially wonderful works are published in particular genres or subject areas. Come in, [meet Eyes](#), and have it all explained.

YOUR ACCOUNT

Check the status of your orders or change the email address and password you have on file with us. Please note that you **do not** need an account to use the store. The first time you place an order, you will be given the opportunity to create an account.

Artificial Intelligence At Amazon (1995)



The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses Hardcover – September 13, 2011

by Eric Ries (Author)
★★★★☆ 1,043 customer reviews

See all 4 formats and editions

Hardcover \$16.49	Paperback from \$13.23	Audible \$0.00
------------------------------------	---------------------------	-------------------

63 Used from \$9.92	15 Used from \$13.43	Free with your Audible trial
92 New from \$10.00	4 New from \$13.23	

Most startups fail. But many of those failures are preventable. The Lean Startup is a new approach being adopted across the globe, changing the way companies are built and new products are launched.

Eric Ries defines a startup as an organization dedicated to creating something new under conditions of extreme uncertainty. This is just as true for one person in a garage or a group of seasoned professionals in a Fortune 500 boardroom. What they have in common is a mission to penetrate that fog of uncertainty to discover a successful path to a sustainable business.

The Lean Startup approach fosters companies that are both more capital efficient and that leverage human creativity more effectively. Inspired by lessons from lean manufacturing, it relies on "valid" > Read more

Best of the Month
See the Best Books of the Month
Want to know our Editors' picks for the best books of the month? Browse Best Books of the Month, featuring our favorite new books in more than a dozen categories.

Flip back

Listen



See all 4 images

Frequently Bought Together



Total price: **\$32.82**
[Add both to Cart](#)
[Add both to List](#)

- This item: The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful... by Eric Ries Hardcover \$16.49
- Zero to One: Notes on Startups, or How to Build the Future by Peter Thiel Hardcover \$16.33

Customers Who Bought This Item Also Bought

Page 1 of 10



Zero to One: Notes on Startups, or How to Build the Future by Peter Thiel



Business Model Generation: A Handbook of Visionary Thinking to Reinvent the Way You Business by Alexander Osterwalder



The Startup Owner's Manual: The Step-By-Step Guide to Building a Successful Business by Steve Blank



How to Win Friends & Influence People by Dale Carnegie



The \$100 Startup: Reinvent the Way You Make Money, Live Better, and Work Smarter by Chris Pinney



The Innovator's Dilemma: The Revolutionary Book that Will Destroy the Big Companies by Clayton Christensen



Value Proposition Design: How to Create Products Customers Want by Alexander Osterwalder



The Hard Thing About Hard Things: Building a Business that Lasts by Ben Horowitz



Traction: How Any Startup Can Achieve Explosive Customer Growth by Gino Wickham



The 4-Hour Workweek: Escape 9-5, Live Anywhere, and Join the New Rich by Timothy Ferriss

Share icons

Buy New \$16.49

Qty: 1 List Price: \$28.00 Save: \$11.51 (41%)

FREE Shipping on orders over \$35.

In Stock. Ships from and sold by Amazon.com. Gift-wrap available.

Yes, I want FREE Two-Day Shipping with Amazon Prime

Add to Cart

Turn on 1-Click ordering for this browser

Want it tomorrow, Feb. 16? Order within 12 hrs 7 mins and choose One-Day Shipping at checkout. Details

Ship to: Select a shipping address: -

Buy Used \$9.92

Add to List

Sell yours for a Gift Card We'll buy it for \$4.42 Learn More

Trade in now



Have one to sell? Sell on Amazon

...the most significant story of our time.

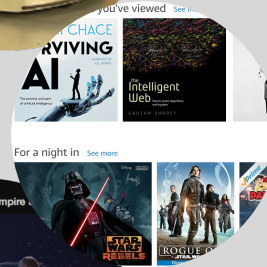
The Four Anthony Goulet ★★★★★ 13 \$4.99

Ad feedback

At Amazon,
we've been making
investments in ML for
the last 20 years...



Fulfillment & logistics




Search & discovery



Existing products

Page	Chapter	Book	Done
A Dance with Dragons: A Song of Ice... All (461) People (440) Terms (21)			
Daena			
Daenerys Targaryen			
"Queen Daenerys" - "Dany" - "Daenerys Stormborn" - "The Unburnt" - "Mother of Dragons" - "Queen of Meereen" - "Queen of the Andals and the Rhoynar and the First Men" - "Khaleesi of Great Grass Sea" - "Breaker of Shackles" - 15 years old; Sole heir to the Targaryen dynasty who ruled Westeros for 300 years u...			
Dagmer Cleftjaw			
Dagon Codd			
Dagon Greyjoy			
Dalla			
Wildling Queen; Mance Rayder's wife; Perished giving birth to Mance's son			
Damon			
One of the Bastard's Boys; Is fair haired and boyish			
Dannel			
Danny Flint			
Dareon			
Sworn Brother of the Night's Watch; A handsome young singer with fair hair and hazel eyes; Was sent to the Wall accused of rape but he has always claimed he is innocent; Was sent to Eastwatch after taking his vows; Jon & Sam's friend; Jon named him a recruiter to replace Yoren who vanished on his way to the wall. He will...			
Daryn Hornwood			
Davos Seaworth			

X-Ray	Khal Drogo	Done
Dany's "sun-and-stars" - Daenerys' deceased husband; A great khal of the Dothraki; Was never defeated in battle; Died of a wound that festered and was help along by Mirri Maz Duur		
From Shelfari.com		
		
Chapter 2		
They had been given to her when she wed Khal Drogo . It was Drogo who had given her the pelt she wore, the head and hide of a hrakkar, the white lion of the Dothraki sea.		
Page 31 • Loc 648		
Chapter 5		
"That did not stop you selling her to Khal Drogo ... "		
"Dothraki neither buy nor sell. Say rather that her brother Viserys gave her to Drogo to win the khal's friendship.		
Page 73 • Loc 1470		
Chapter 23		
"Bring the grey linen gown with the pearls on the		


Star Wars: Episode V - The Empire Strikes Back



2:01:29/2:04:19 **HD**

[View All X-Ray](#)

[Options](#)




amazon echo


At Amazon,
we've been making
investments in ML for
the last 20 years...




Fulfillment & logistics



Search & discovery



Existing products



New products



fire TV



amazon

amazon

NO LONGER
NECESSARY
**JUST
WALK
OUT**
SHOPPING

amazon



ML @ AWS: OUR MISSION

**Put machine learning in
the hands of every
developer and data
scientist**

The Amazon Machine Learning stack

Application Services

Platform Services

Frameworks & Infrastructure

Bottom layer: frameworks & infrastructure



NVIDIA
Tesla V100 GPUs

5,120 tensor cores

128 GB of memory

1 petaflop of compute

NVLink 2.0


~14X faster than P2



AWS Deep Learning AMI




Meeeeelions of uncorrelated workloads



Collective action

When everyone comes together in the cloud to share the resource, and **only pays for what they use**, the efficiency is huge.

Spot Market



Spot Market

Our ultimate space filler.

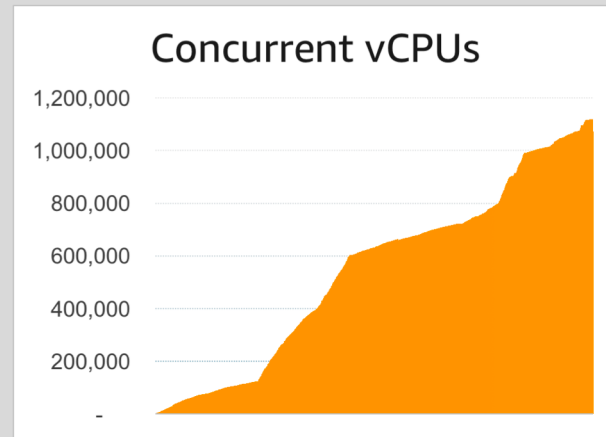
Great for workloads that aren't time sensitive, and especially popular in research (hint: **it's really cheap**).

Spot Instances: Spare compute capacity at *scale*

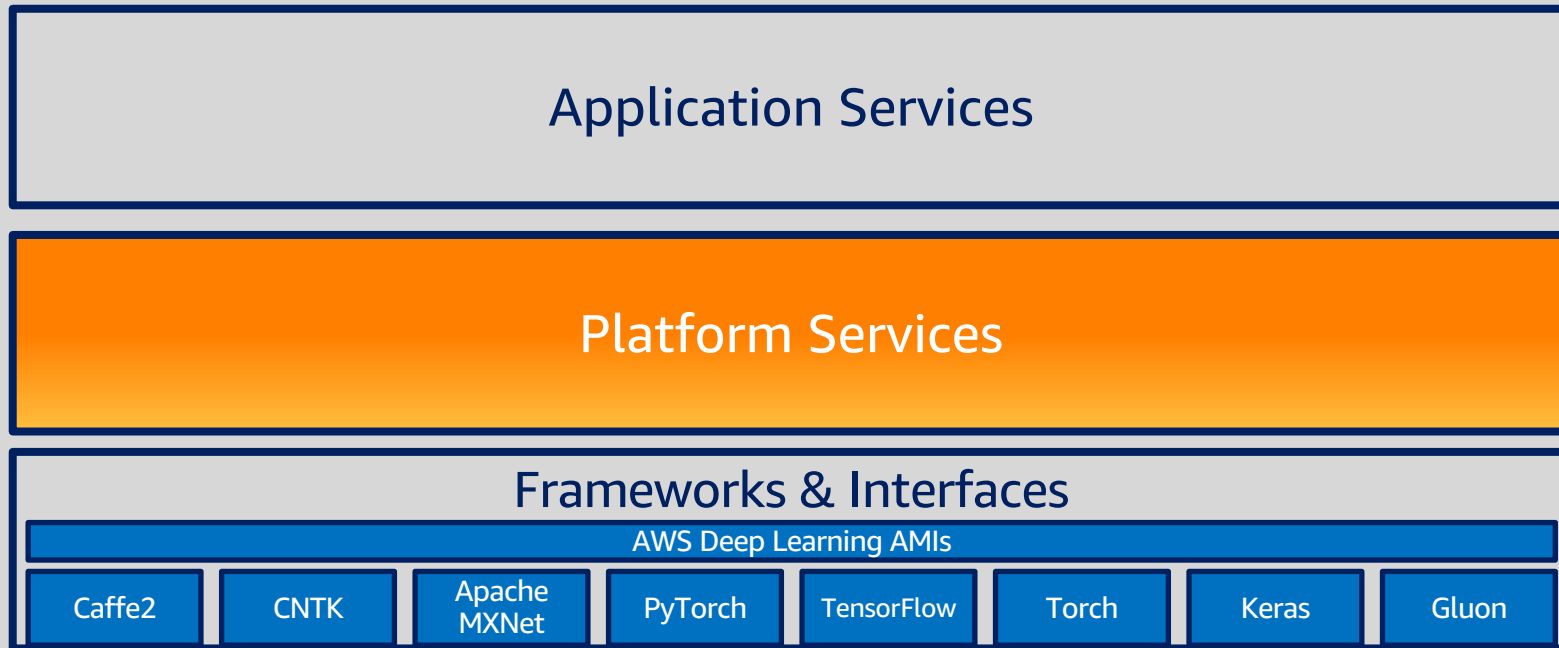


ON AVERAGE, **EVERY WEEK**, AWS CUSTOMERS ARE USING MORE COMPUTE CAPACITY ON AMAZON EC2 SPOT INSTANCES THAN CUSTOMERS IN 2012 WERE RUNNING ACROSS **ALL OF AMAZON EC2**.

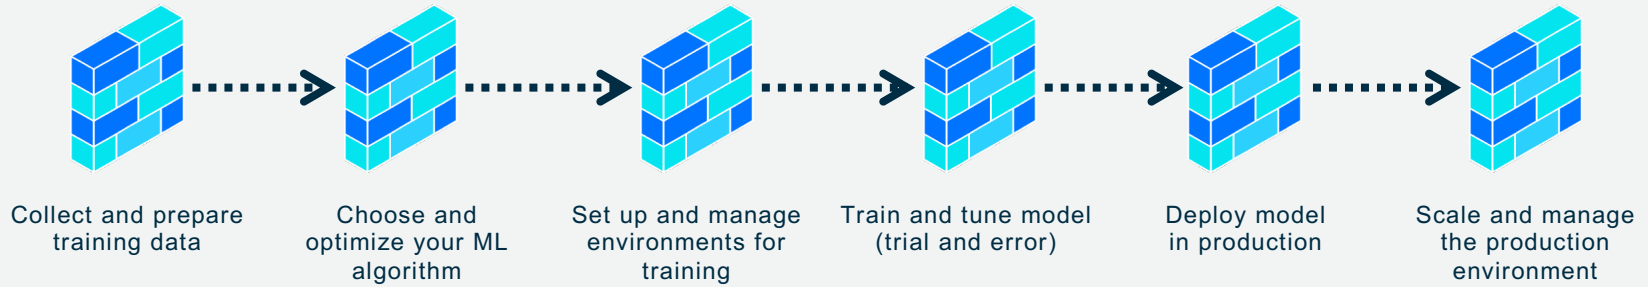
CLEMSON UNIVERSITY RAN THE LARGEST HIGH PERFORMANCE CLUSTER ON SPOT WITH **1,100,000 CONCURRENT vCPUs**



The Amazon Machine Learning stack



Amazon SageMaker



Amazon SageMaker



Pre-built notebooks for common problems



Built-in, high performance algorithms

ALGORITHMS

K-Means Clustering	XGBoost
Principal Component Analysis	Latent Dirichlet Allocation
Neural Topic Modelling	Image Classification
Factorization Machines	Seq2Seq
Linear Learner - Regression	Linear Learner - Classification

FRAMEWORKS

Apache MXNet	Caffe2, CNTK,
TensorFlow	PyTorch, Torch



Set up and manage environments for training



Train and tune model (trial and error)




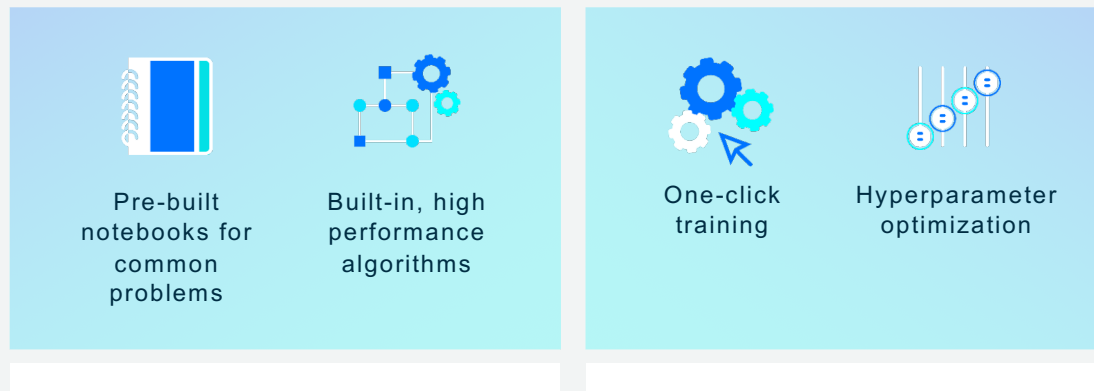
Deploy model in production



Scale and manage the production environment

BUILD

Amazon SageMaker



Amazon SageMaker



Pre-built notebooks for common problems




Built-in, high performance algorithms

BUILD



One-click training



Hyperparameter optimization

TRAIN



One-click deployment




Fully managed hosting with auto-scaling

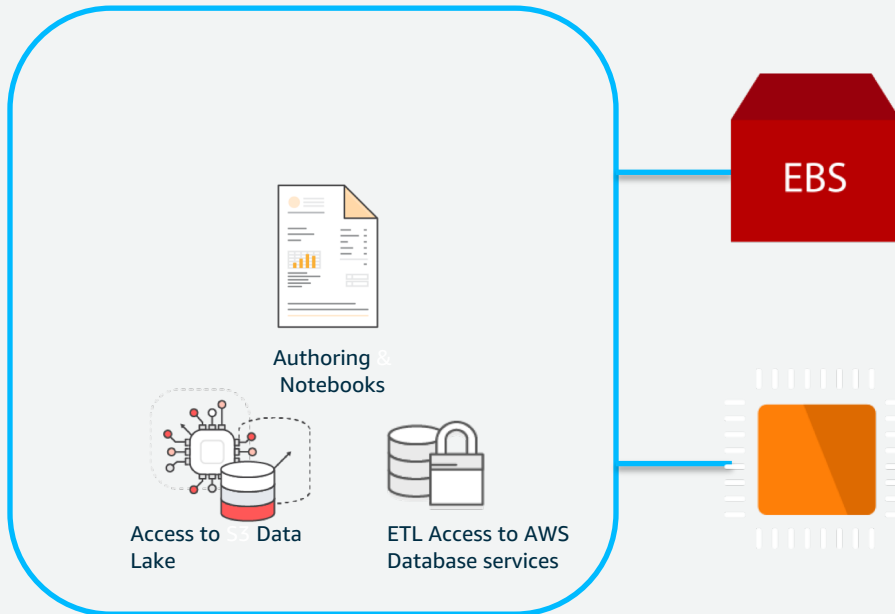
DEPLOY

1

Zero Setup For Exploratory Data Analysis



Notebook Instances



"Just add data"

- Recommendations/Personalization
- Fraud Detection
- Forecasting
- Image Classification
- Churn Prediction
- Marketing Email/Campaign Targeting
- Log processing and anomaly detection
- Speech to Text
- More...

2

Amazon SageMaker: 10x better algorithms




Algorithms



Streaming datasets,
for cheaper training



Train faster, in a
single pass



Greater reliability on
extremely large
datasets



Choice of several ML
algorithms



Our Customers use ML at a massive scale!

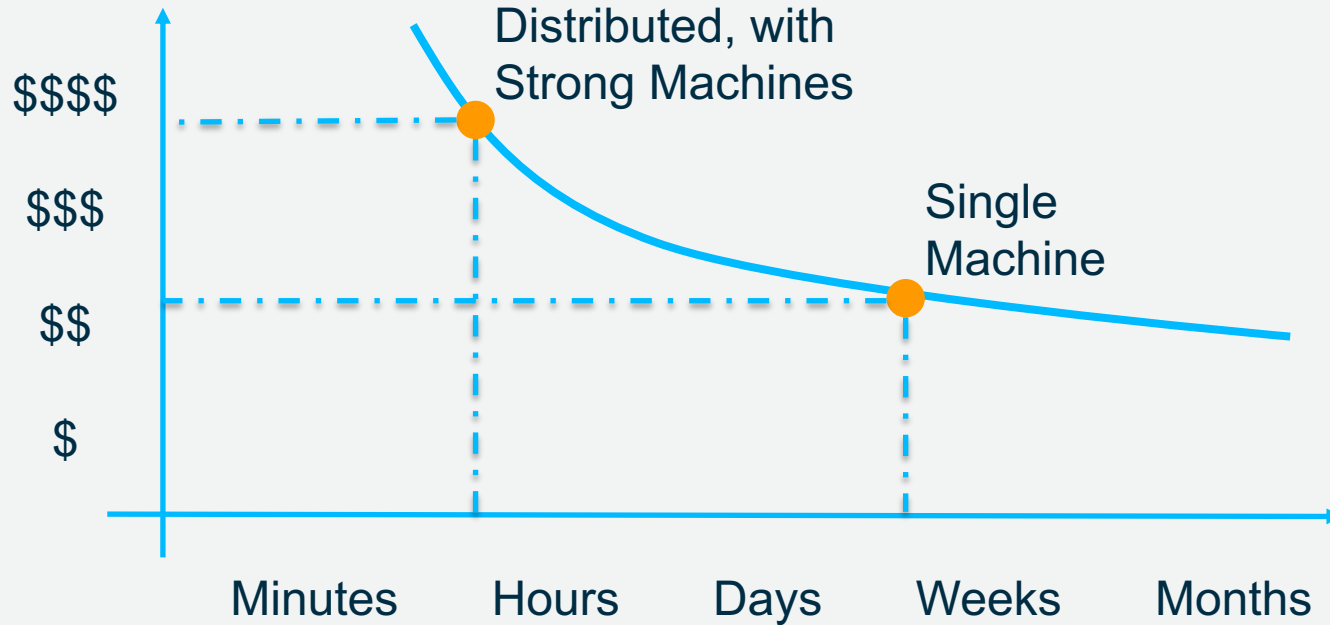


"We process 3 million ad requests a second, 100,000 features per request. That's 250 trillion per day. Not your run of the mill Data science problem!"
Bill Simmons, CTO




"We collect 160M events daily in the ML pipeline and run training over the last 15 days and need it to complete in one hour. Effectively there's 100M features in the model"
Valentino Volonghi, CTO


Cost vs. Time



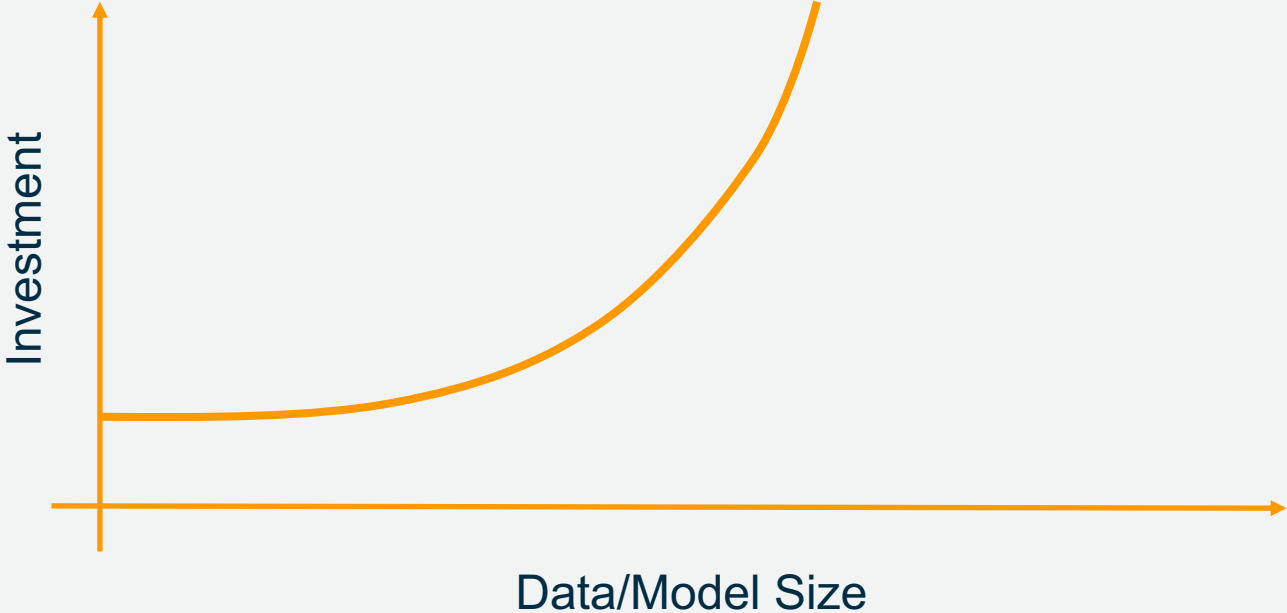
Model Selection



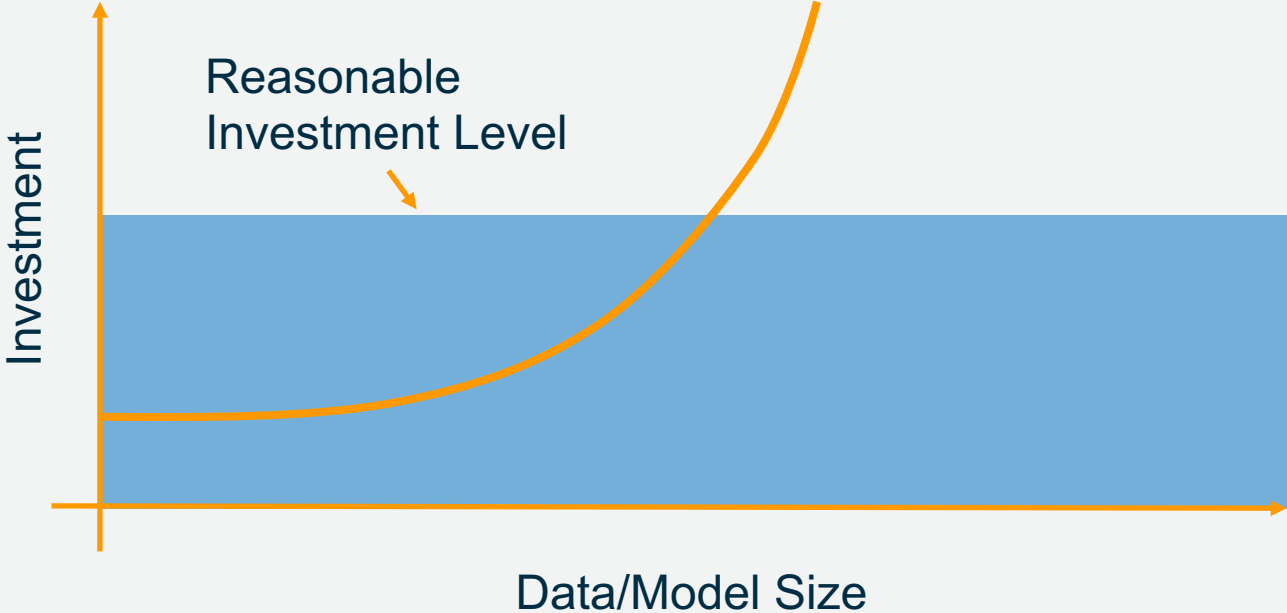
Incremental Training




Production Readiness



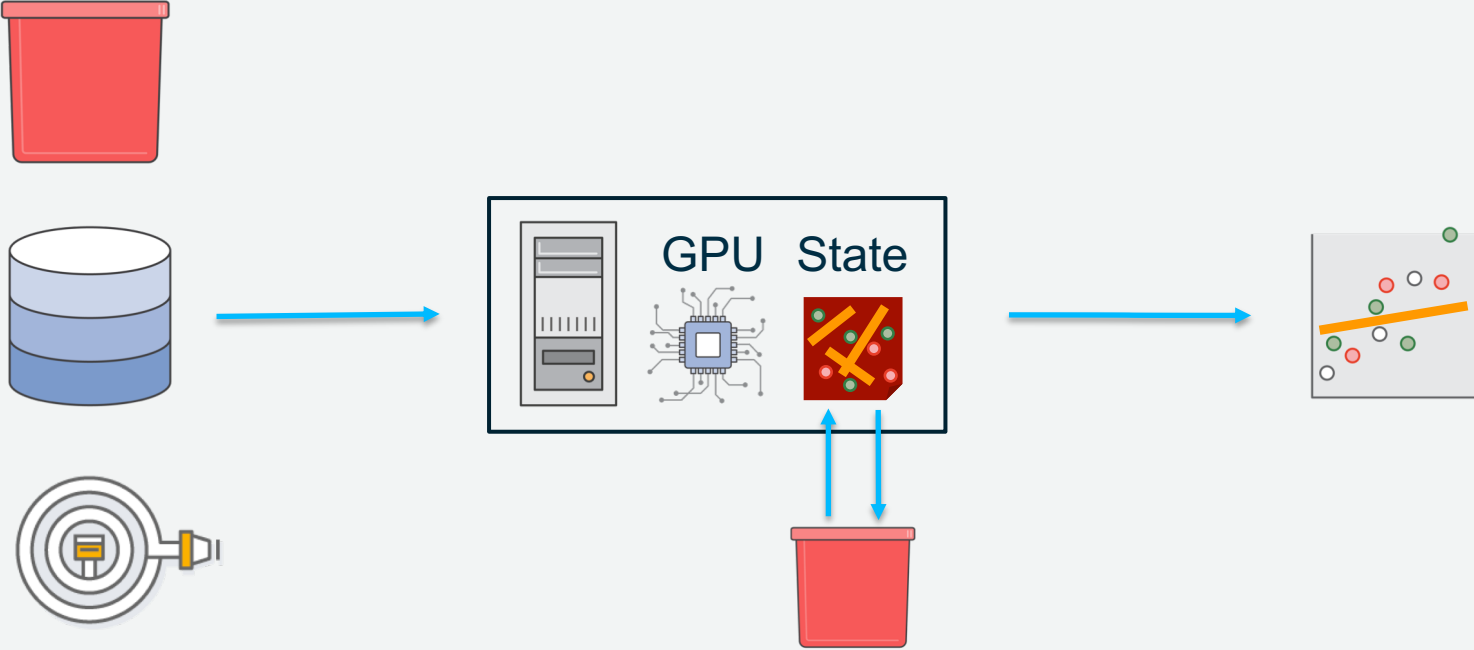
Production Readiness




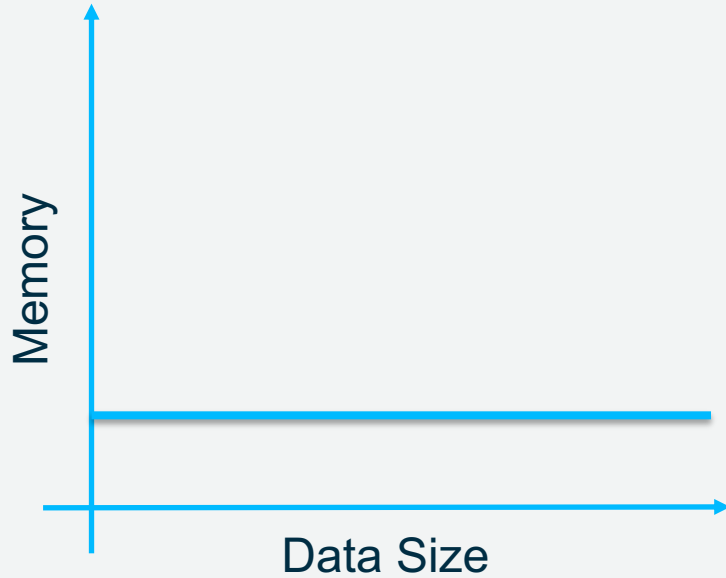
Production Readiness




Streaming




Streaming




Incremental Training



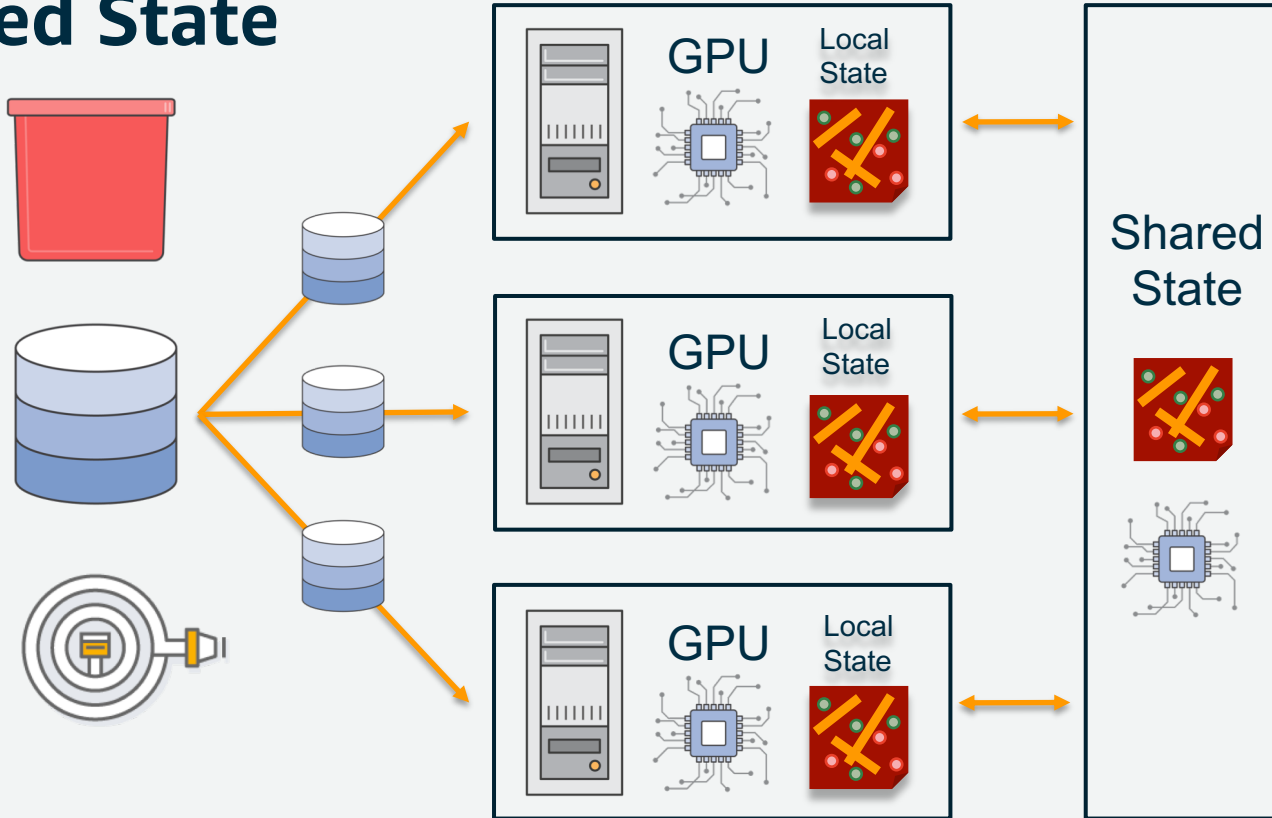
Incremental Training



Distributed



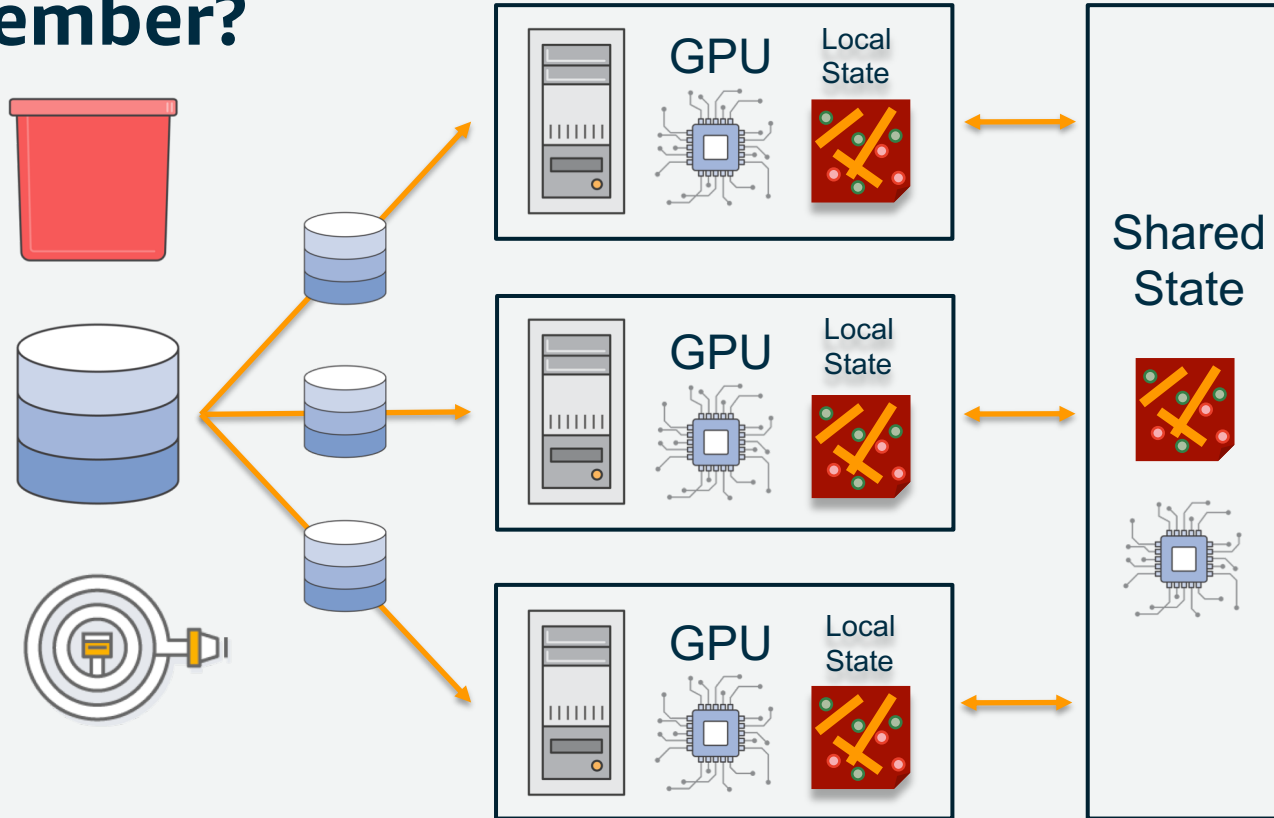
Shared State




Cost vs. Time



Remember?




Pipe Mode



Pipe Mode

- Data flows on-the-fly during training
- Shorter startup times
- Higher I/O throughput with optimized agent
- No dataset size limits
- Requires protobuf recordIO format
- Available for most first party algorithms

PCA + K-Means on NY Taxi Dataset

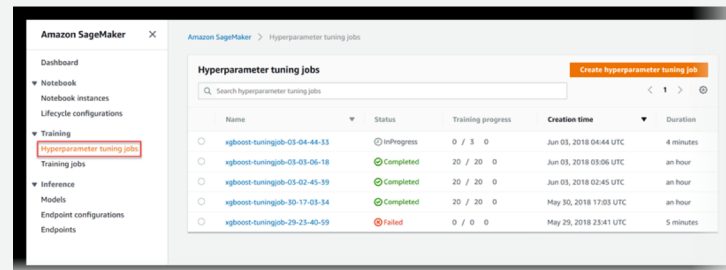


<https://aws.amazon.com/blogs/machine-learning/using-pipe-input-mode-for-amazon-sagemaker-algorithms/>

Automatic Model Tuning

1. Define objective metrics
2. Define hyperparameter ranges
3. Let automatic tuning do the job!

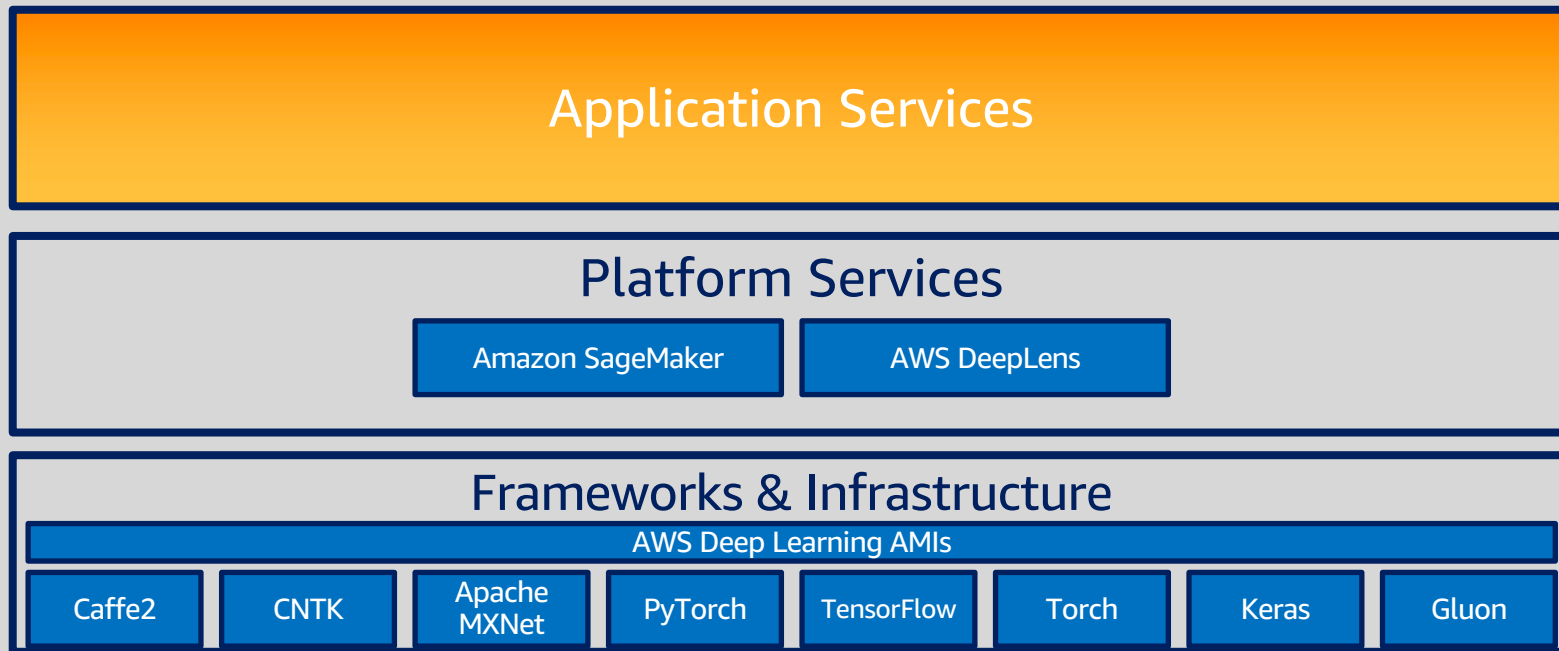
```
"ParameterRanges": {  
  "CategoricalParameterRanges": [  
    {  
      "Name": "tree_method",  
      "Values": ["auto", "exact", "approx", "hist"]  
    },  
    {  
      "Name": "eta",  
      "MaxValue": "0.5",  
      "MinValue": "0"  
    }  
  ],  
  "IntegerParameterRanges": [  
    {  
      "Name": "max_depth",  
      "MaxValue": "10",  
      "MinValue": "1",  
    }  
  ]  
}
```



The screenshot shows the Amazon SageMaker console interface for Hyperparameter tuning jobs. The left sidebar contains navigation options: Dashboard, Notebook, Notebook instances, Lifecycle configurations, Training (with 'Hyperparameter tuning jobs' selected), Inference, Models, Endpoint configurations, and Endpoints. The main content area displays a table of hyperparameter tuning jobs with columns for Name, Status, Training progress, Creation time, and Duration. A 'Create hyperparameter tuning job' button is visible in the top right.

Name	Status	Training progress	Creation time	Duration
xgboost-tuningjob-03-04-44-33	InProgress	0 / 3 0	Jun 03, 2018 04:44 UTC	4 minutes
xgboost-tuningjob-03-06-18	Completed	20 / 20 0	Jun 03, 2018 05:06 UTC	an hour
xgboost-tuningjob-03-02-45-39	Completed	20 / 20 0	Jun 03, 2018 02:45 UTC	an hour
xgboost-tuningjob-30-17-03-34	Completed	20 / 20 0	May 30, 2018 17:03 UTC	an hour
xgboost-tuningjob-29-23-40-59	Failed	0 / 0 0	May 29, 2018 23:41 UTC	5 minutes

The Amazon Machine Learning stack



Amazon ML application services

Vision



Amazon Rekognition
Amazon Rekognition
Video

Speech



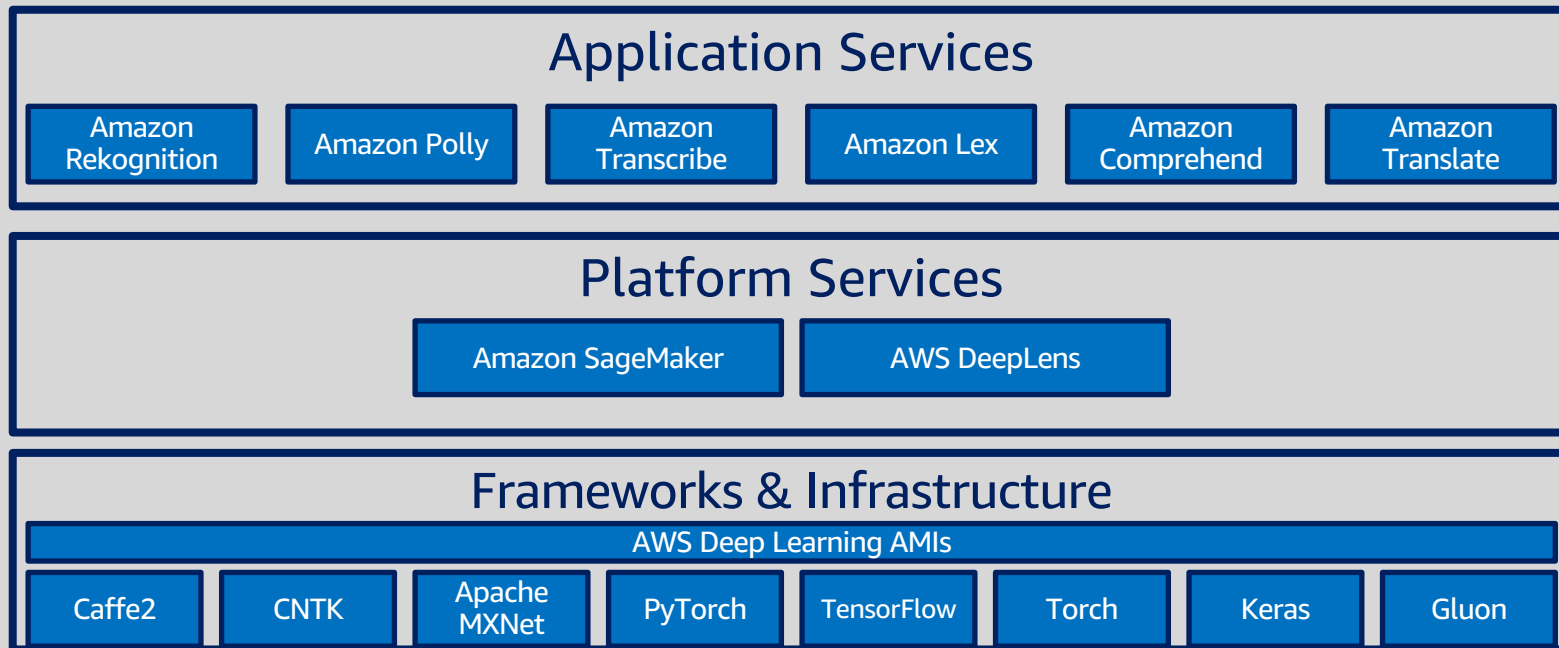
Amazon Polly
Amazon Transcribe

Language



Amazon Lex
Amazon Translate
Amazon Comprehend

The Amazon Machine Learning stack





Teach Tomorrow's Cloud Workforce Today

With the increasing demand for cloud employees, AWS Educate provides an academic gateway for the next generation of IT and cloud professionals. AWS Educate is Amazon's global initiative to provide students and educators with the resources needed to accelerate cloud-related learning.

Join AWS Educate

<https://aws.amazon.com/education/awseducate/>

AWS Academy

AWS Cloud computing curricula specifically designed for higher education institutions.
AWS Academy connects the jobs of the future to the next generation of IT professionals.

Apply to become an AWS Academy Member

Institution application

<https://aws.amazon.com/training/awsacademy>

AWS Cloud Credits for Research

The AWS Cloud Credits for Research Program (formerly AWS Research Grants) supports researchers who seek to:

1. Build cloud-hosted publicly available science-as-a-service applications, software, or tools to facilitate their future research and the research of their community.
2. Perform proof of concept or benchmark tests evaluating the efficacy of moving research workloads or open data sets to the cloud.
3. Train a broader community on the usage of cloud for research workloads via workshops or tutorials.



<https://aws.amazon.com/research-credits>

Call to action!

cpetters@amazon.de

