

Preparation to exam AI 100

« Designing and Implementing an Azure AI Solution »

Module 1 - Analyze Solution Requirement (25-30%)

Mohammed Arif Senior Technical Consultant



What is the Microsoft AI 100 exam?

 AI-100 Certification is for engineers who have intermediate to an expert level understanding of Machine Learning and are responsible for deploying end-to-end AI solutions on Azure using various AI & ML services. Part 1

Recommend Cognitive Services APIs to meet business requirements



Microsoft Cognitive Services

- Microsoft Cognitive Services is a set of machine learning algorithms available as a service for anyone to use
- Instead of building this intelligence for our apps from scratch, we can use these services for vision, speech, language, knowledge, and search
 - · You can try each service for free
 - When you decide to integrate a service into your app, you sign up for a paid subscription

Cognitive Services

Microsoft



Cognitive Services Save for later



Microsoft

Create

Cognitive Services is a product bundle that enables customers to access multiple services with a single API key.

Product features:

Access to Vision, Language, Search, and Speech services using a single API

Quickly connect services together to achieve more insights into your content

Easily integrate with other services like Azure Search

Dashboard > New > Cognitive Services > Create
Create All Cognitive Services
* Name
PrepaAl100-CognitiveServices
* Subscription
Conso interne de la plateforme Windows Azure 2
* Location
(Europe) West Europe
Location specifies the region only for included regional services. This does not specify a region for included non-regional services. See service status page for more details.
Eocation specifies the region only for included regional services. This does not specify a region for included non-regional services. See service status page for more details.
* Pricing tier (View full pricing details)
S0 ×
* Resource group
RG-Prepa-AI-100
Create new
✓ * I confirm I have read and understood the notice below.
Microsoft will use data you send to Bing Search Services to improve Microsoft products and services. Where you send personal data to this service, you are responsible for obtaining sufficient consent from the data subjects. The Data Protection Terms in the Online Services Terms do not apply to Bing Search Services.
Please refer to the Online Services Terms for details. Microsoft offers policy controls that may be used to disable new deployments.

Cognitive Services Area

- Vision: image classification and recognition
- · Speech: conversion between speech and text
- · Language: Spell check, Analytics, translation
- Knowledge : Q&A
- · Search: text, image, video

Vision API

- Vision services to resolve imaging or visual content :
 - Computer vision
 - Content moderator
 - Face API (attribute detection, recognition)
 - Video Indexer
 - Custom Vision Service
 - Emotion API

Develop solutions by using intelligent algorithms that identify items from images and videos

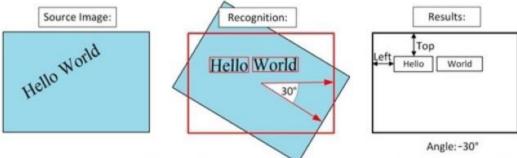
- Use the Computer Vision API to analyze images for insights, extract text from images, and generate high-quality thumbnails
 - 1. Create a Cognitive Services account.
 - 2. Get information about the visual content found in an image.
 - 3. Generate a thumbnail of an image.
 - 4. Detect and extract printed text from an image.
 - 5. Detect and extract handwritten text from an image.

You can use the Computer Vision API to

- Analyze images (JPEG, PNG, GIF, BMP) for insight
- Extract printed text from images using optical character

recognition (OCR)

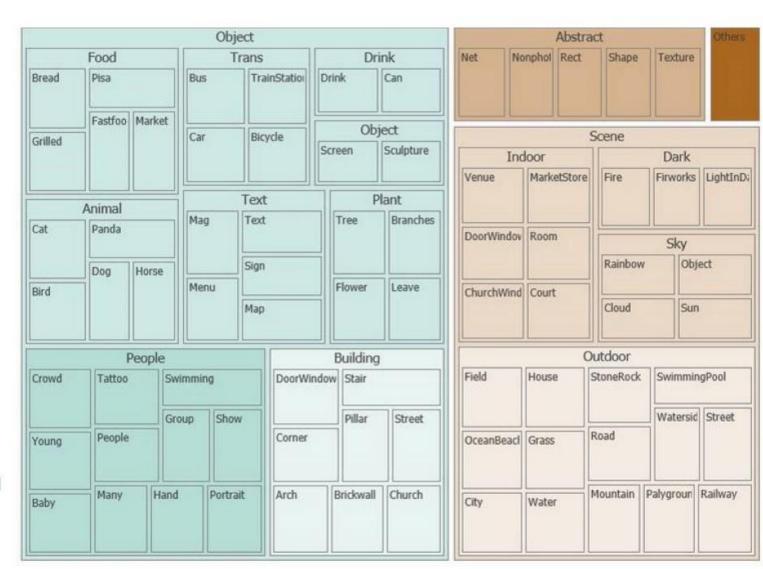
· Support 25 languages



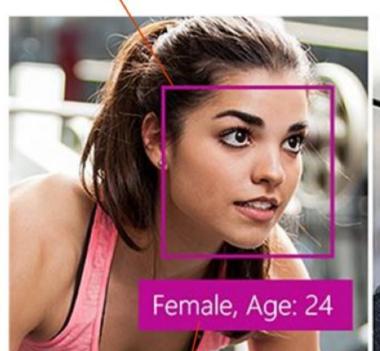
- · Recognize printed and handwritten text from images
- Recognize celebrities and landmarks
- Analyze video
- Generate a thumbnail of an image

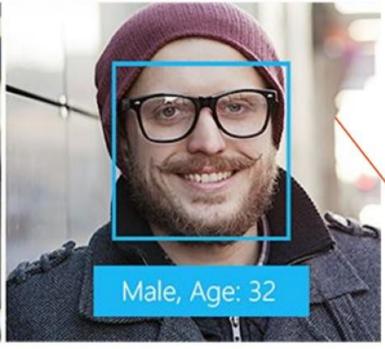
Tagging & Categorizing for Computer vision API

- Tagging images
 - 2000 recognizables objects, living beings, scenary and actions
- Categorizing images
 - · 86-category concept
- Domain-specific content
 - Celebrity or landmark recognition



Face location





Submitted faces

Metadata about face

Custom Vision Service

- Image classification service
- User provided, custom, domain specific model
 - · Increased specificity
 - Reduced distractors
- Help for classification challenge like
 - · Image deformation, background distractions, Partial obscuring of object...
- Not a « black Box model » to analyze an image because you provided your specific model that you have trained

Image Classification Pipeline

Input

 Pre-classified, training set of images



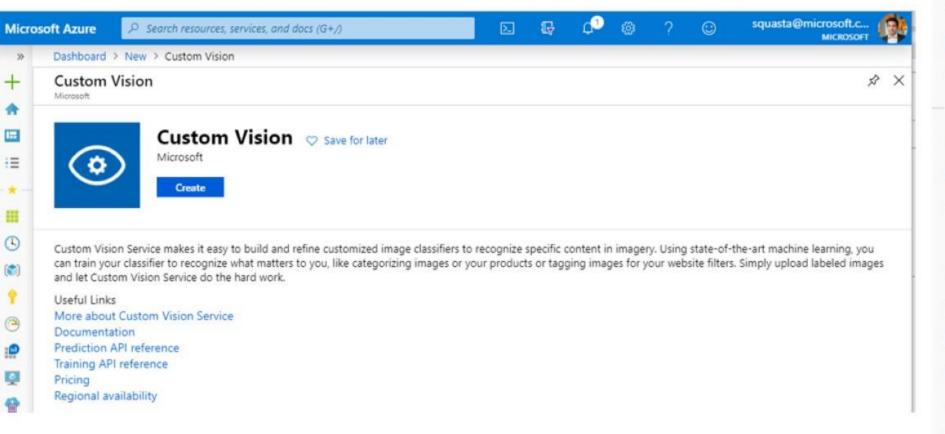
Learning

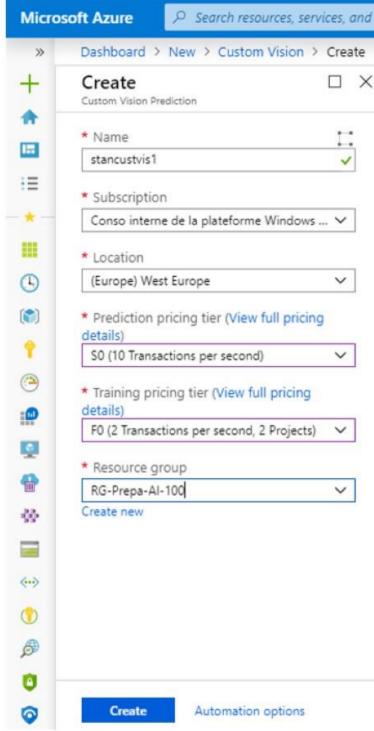
 Train the model to distinguish between each category

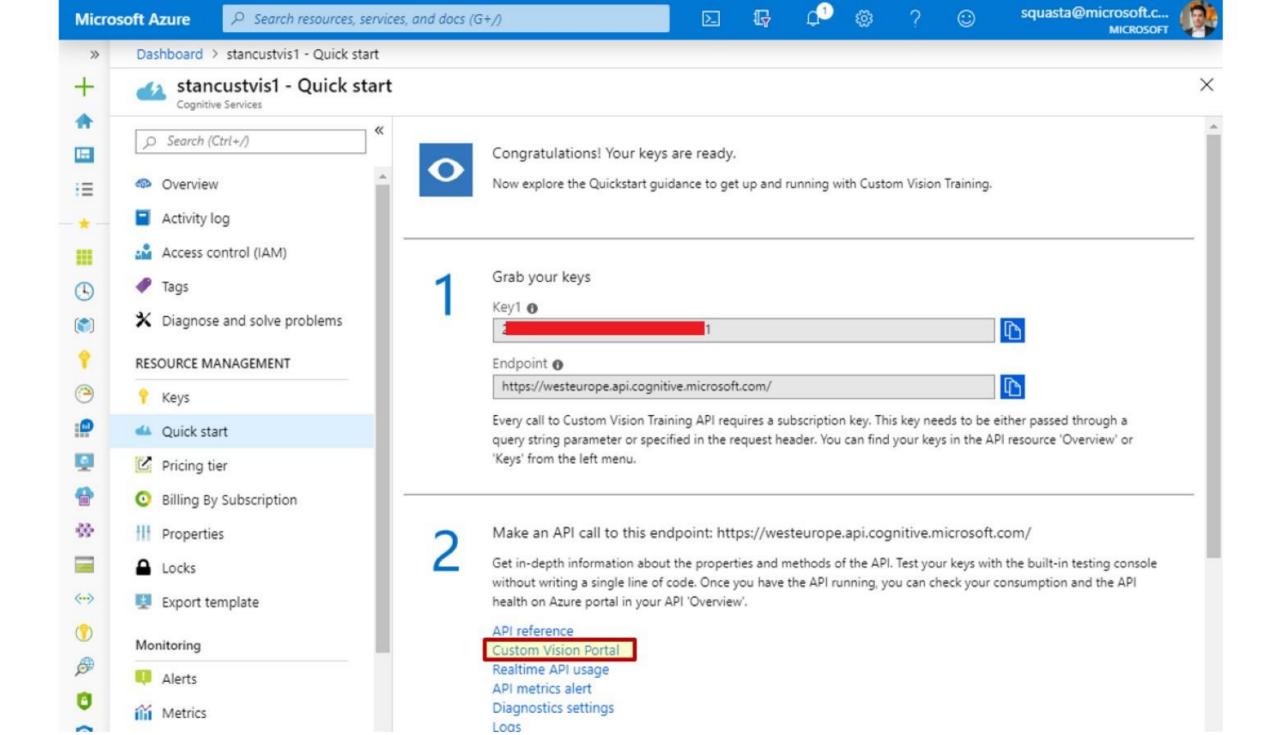


Evaluation

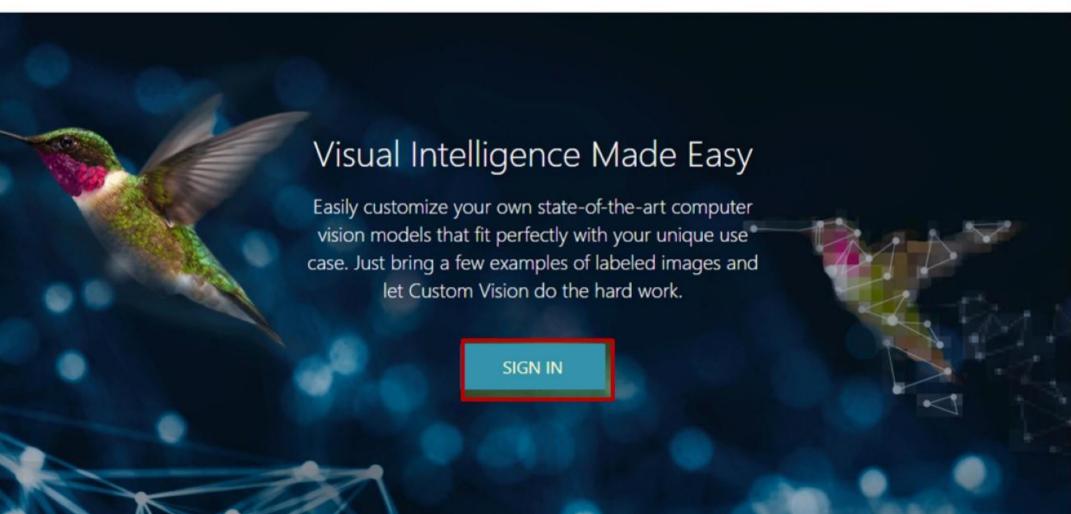
 Validate the model produces accurate predictions with new images

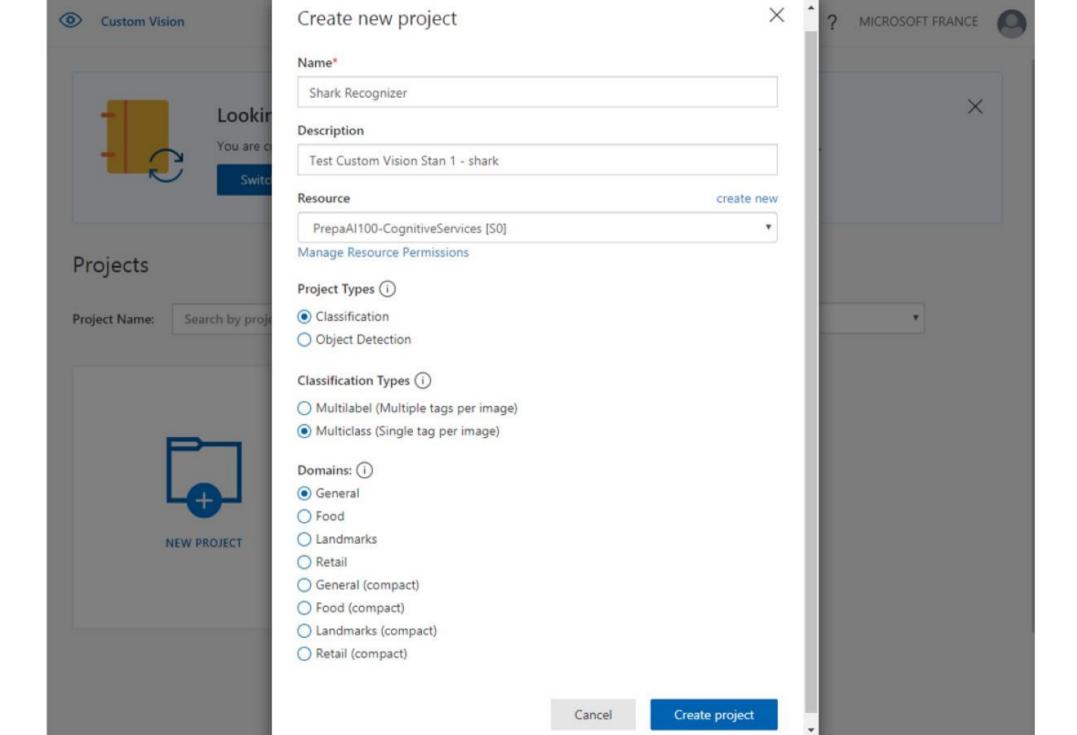


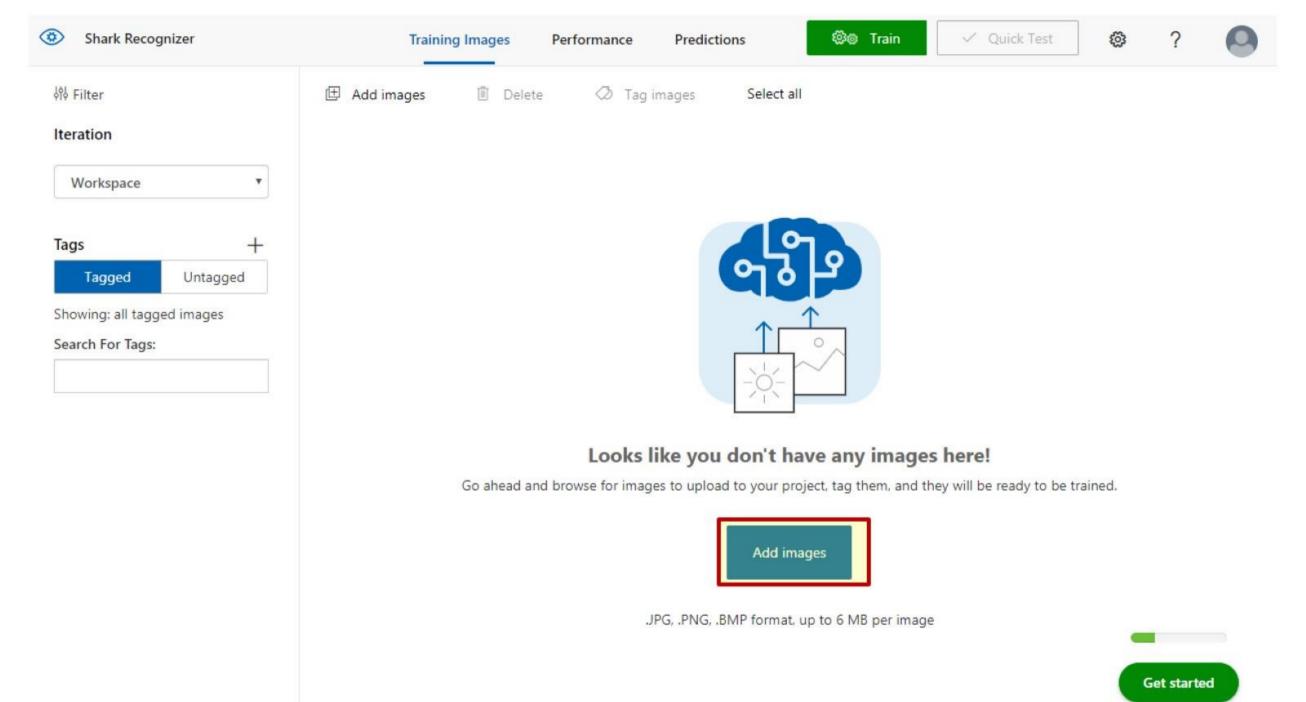












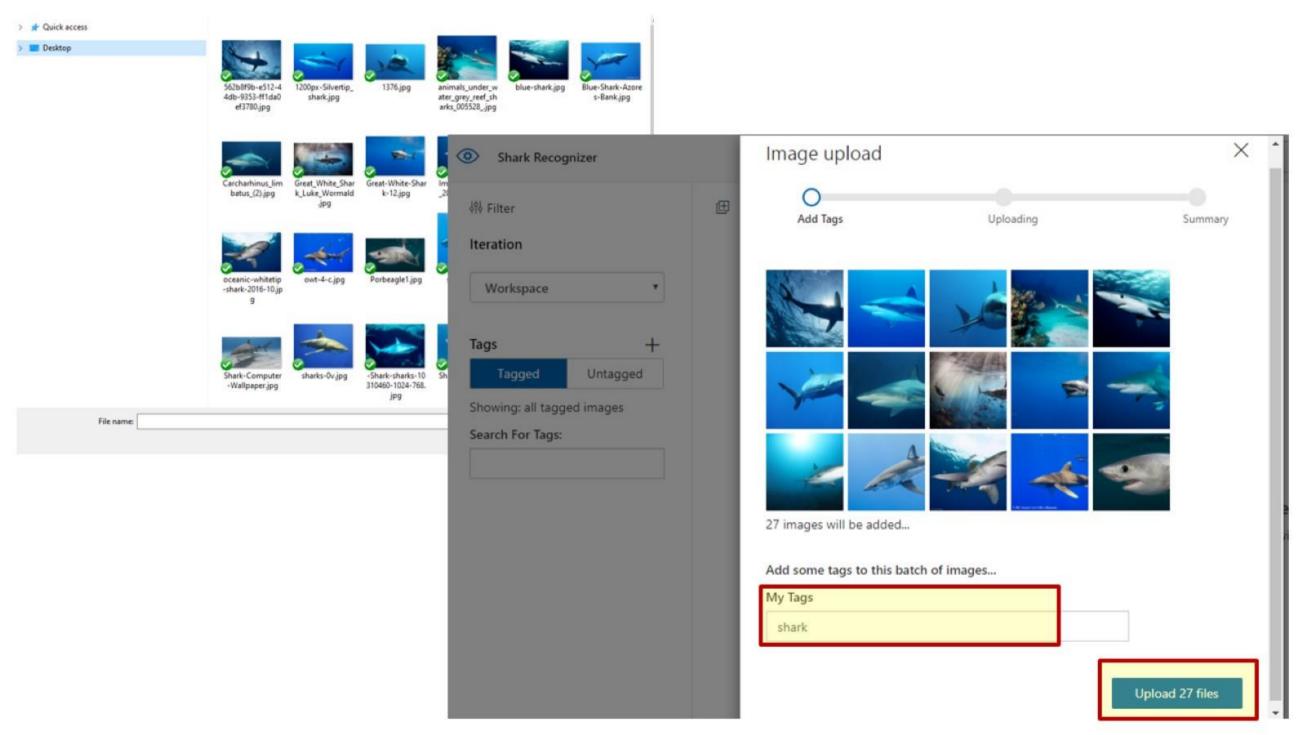


Image upload





Add Tags

Uploading

Summary

26 images uploaded successfully

1 file(s) ignored because they exceeded 6MB or are not images

Done





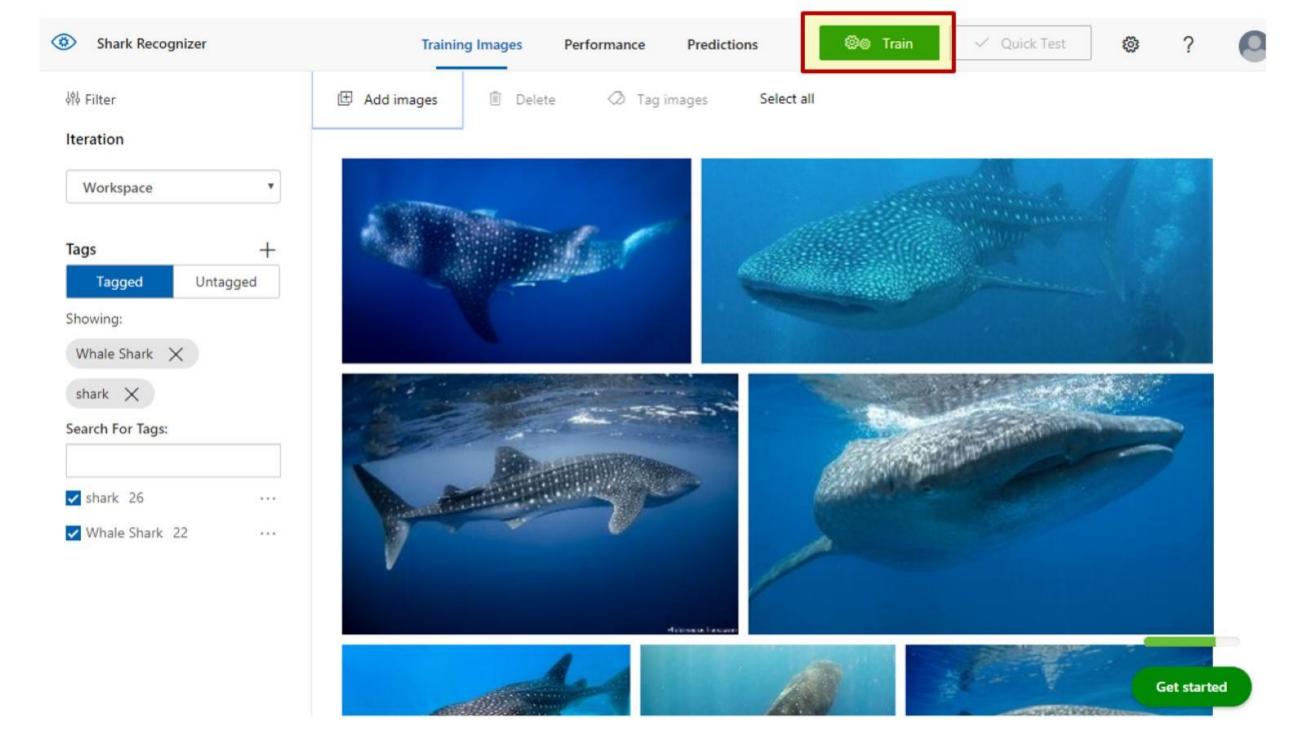
23 images will be added...

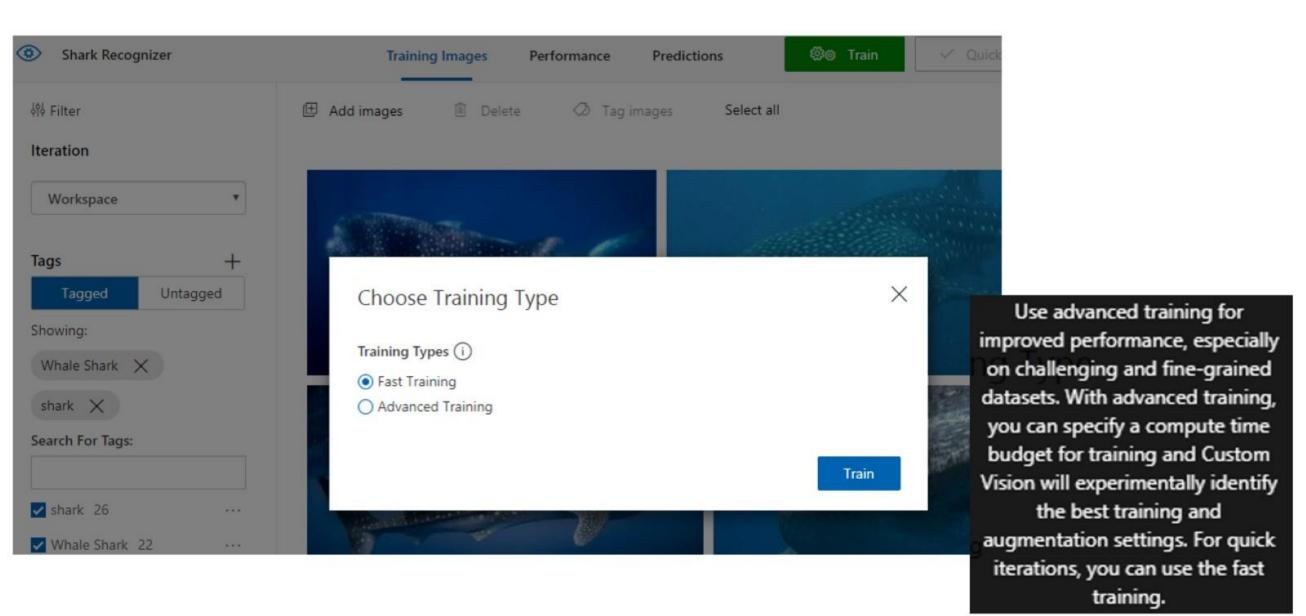
Add some tags to this batch of images...

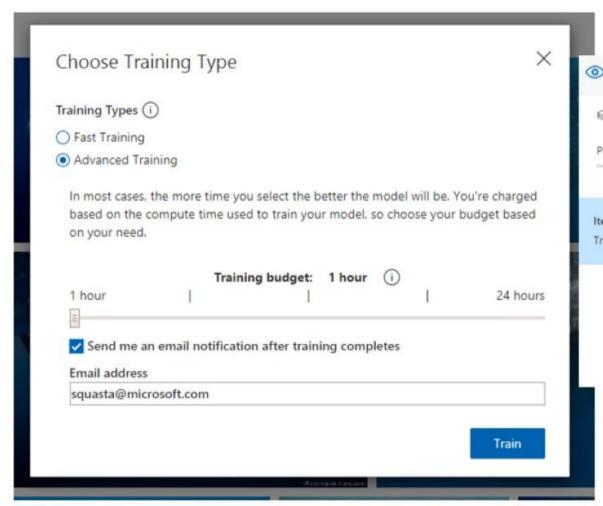
My Tags

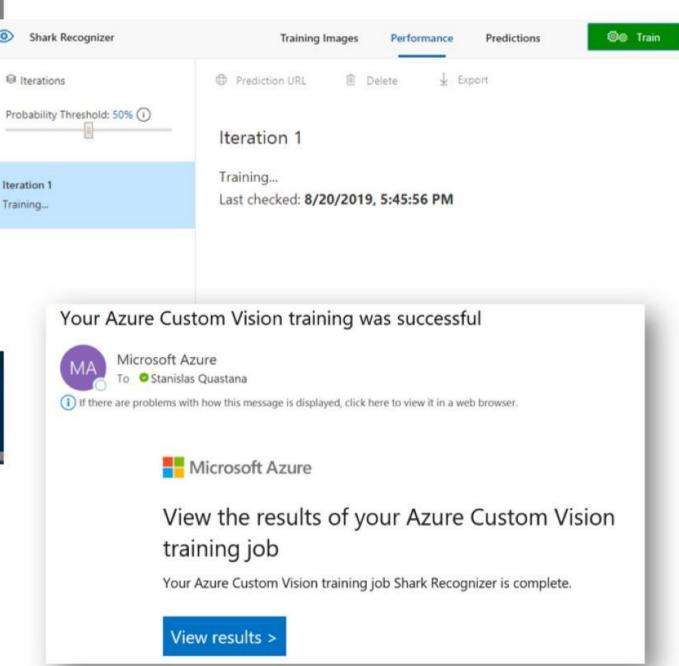
Whale Shark

Summary











Probability Threshold: 50% (i)

Iteration 1

Advanced Trained : 2 minutes ago with General domain, Training Budget: 1 hour





Performance Per Tag

Tag	Precision	^	Recall	A.P.	Image count
Whale Shark	100.0%		100.0%	100.0%	22
shark	100.0%		100.0%	100.0%	26

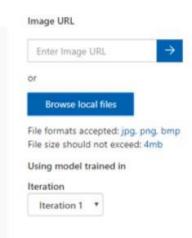


Precision vs Recall

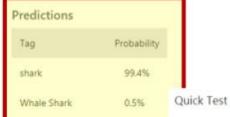
 Precision: represents that ability for a classifier to label all instances correctly

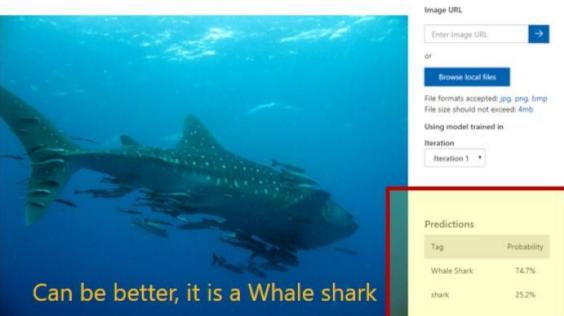
 Recall: represents the ability for a classifier to find all instances of a particular label Quick Test





Wrong, you can improve model by training with this picture





Improving Custom Vision model

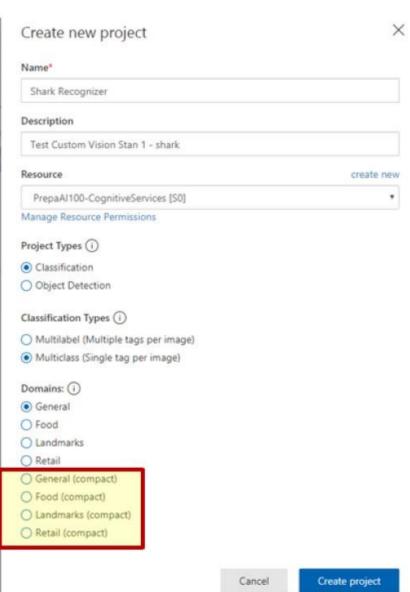
- Adding more pictures for training
 - Minimum 50 images per tag
 - · Use the result of positive or negative prediction
- Diverse type of images : angle, light, background...
- Balance number of images across tags
- Negative image handling

Working with Custom Vision Service Prediction API

- Use code instead of customai portal
- Use information from Azure Portal
 - API key
 - Project ID
 - Prediction URL

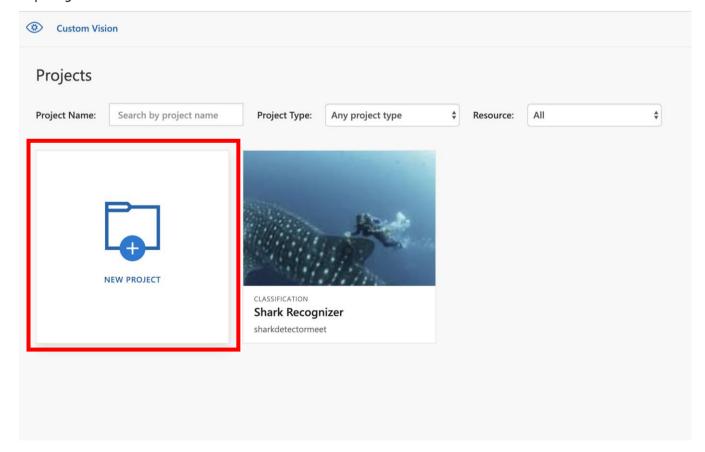
Using a Custom Vision Service Model domains

- Use standard domains
 - Higher accuracy
 - Support Online use only (HTTP calls to the API)
- Use compact domains
 - · Lower accuracy
 - Support Online and offline use

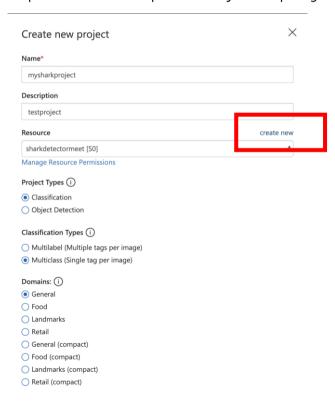


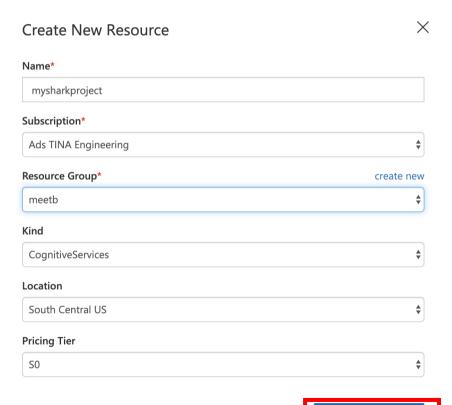
Step 1: Go to https://www.customvision.ai/

Step 2: Create new project!



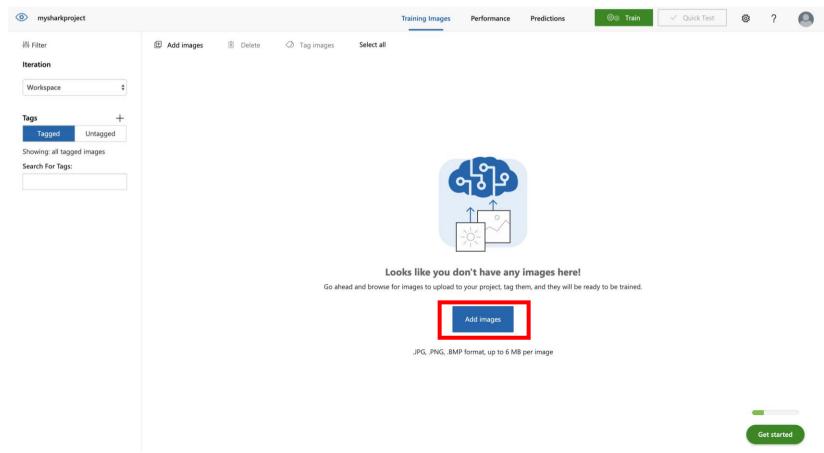
Step 3: Enter input for your project





Create resource

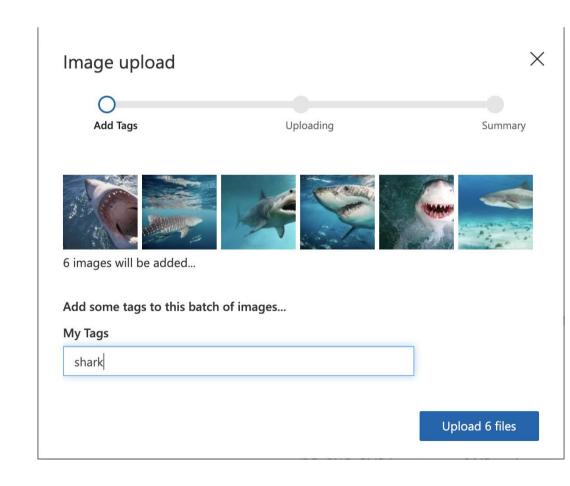
Step 4: Add images



Step 5: Download images and upload them.

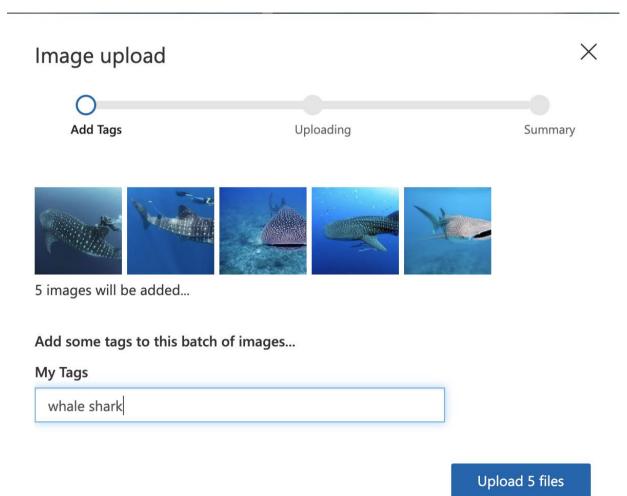
To download go to https://github.com/arifmarias/ai-100-prep/blob/master/whalesharks.zip

Once you unzip, only upload the images from the sharks.zip

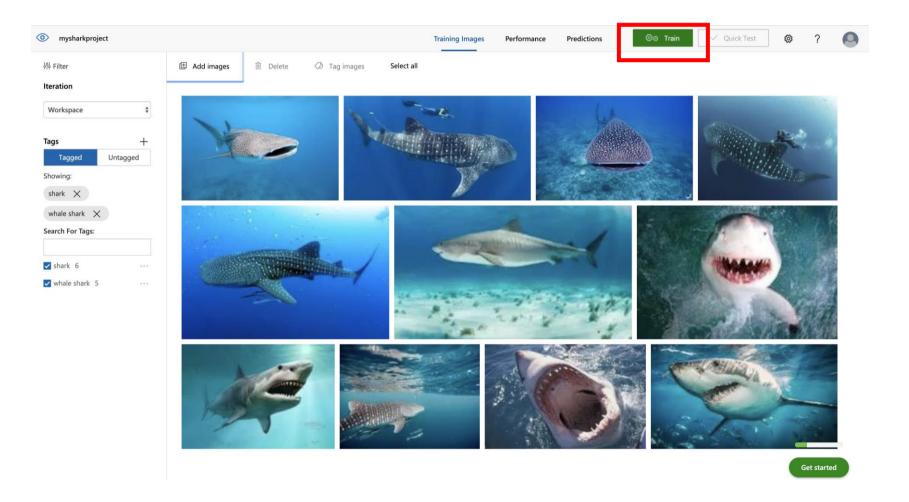


Step 6: Download images and upload them.

Now upload the images from the whale shark folder

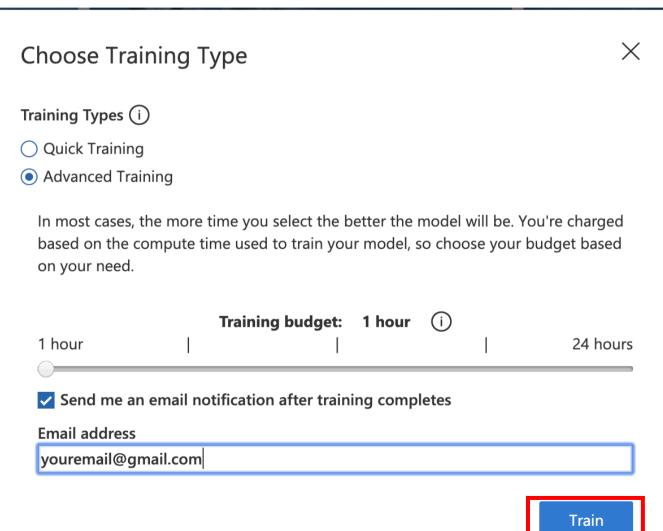


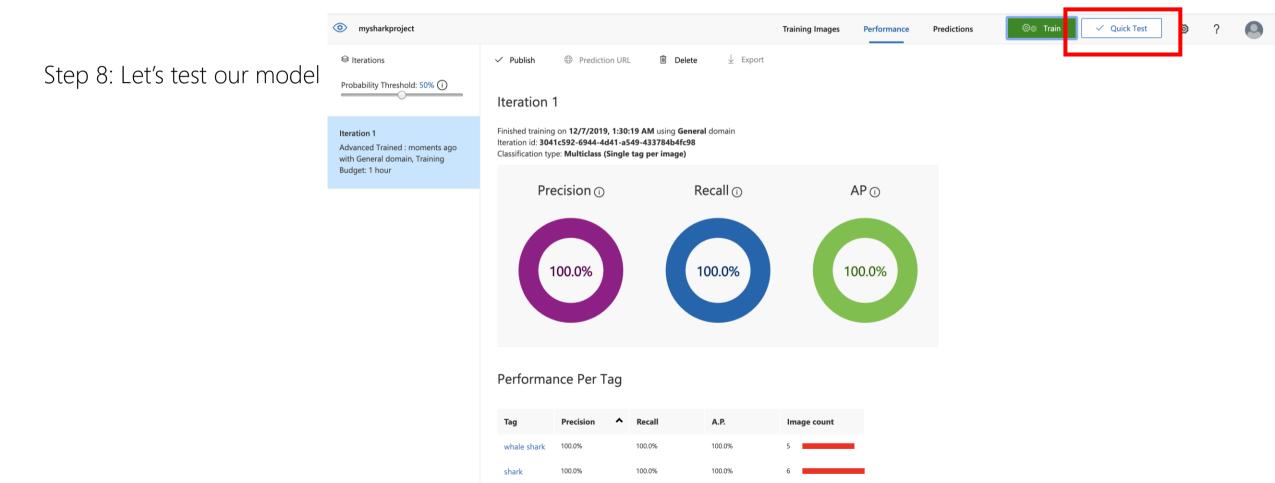
Step 6: Train the model



Step 7: Choose your training For this workshop we will choose advance training.

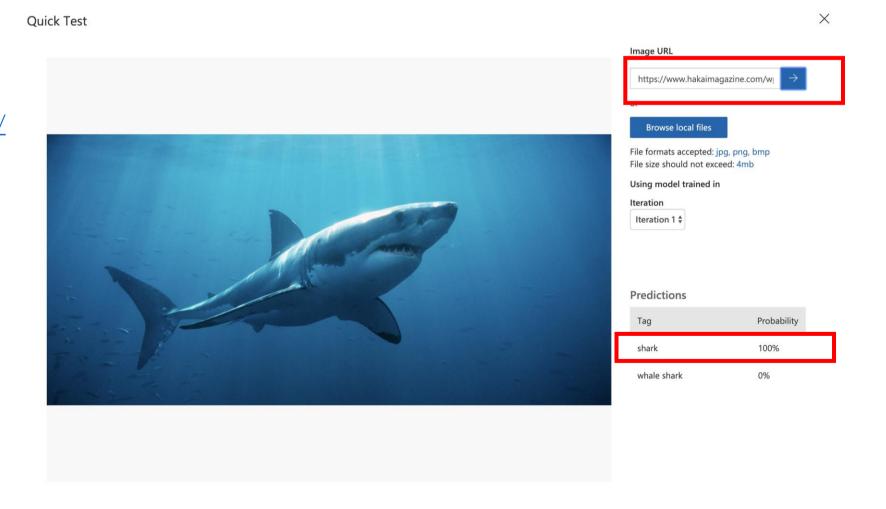
Don't forget to enter your email





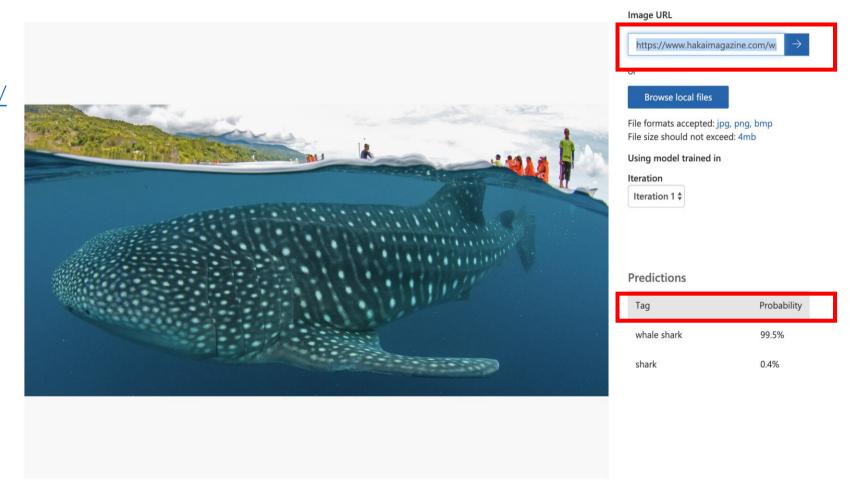
Get started

Step 9: Paste this url and test Your model – https://www.hakaimagazine.com/ wp-content/uploads/headershark-eddies.jpg



Quick Test

Step 9: Paste this url and test Your model – https://www.hakaimagazine.com/ wp-content/uploads/headertourist-shark-feeding.jpg



Hands on lab - Identify faces and expressions by using the Computer Vision API in Azure Cognitive Services

Go to this link - https://docs.microsoft.com/en-us/learn/modules/identify-faces-with-computer-vision/1-introduction

Hands on lab - Process images with the Computer Vision service

Go to this link - https://docs.microsoft.com/en-us/learn/modules/create-computer-vision-service-to-classify-images/1-introduction

Speech Service

- Speech technologies used in products like Cortana and Microsoft Office
- The Speech service unites the Azure speech features previously available via the <u>Bing Speech</u> <u>API</u>, <u>Translator Speech</u>, <u>Custom Speech</u>, and <u>Custom Voice</u> services

Main Speech service functions

- Speech to text (=speech recognition or transcription)
 - · In real-time, batch
 - End of speech detection, automatic text formating, profanity masking
 - · Can call on Language Understanding (LUIS) to derive user intent from transcribed speech
 - · You can use your own data to train the models (ex: Accoustic model, Langage model..)

Text to speech

· Voice font (give your app a voice by training the model on samples of human voice

Speech translation

- · In real-time, can process recorded speech
- Result as text or synthetized speech

Speech scenarios

- Create voice-triggered apps
 - If your app is used in a specific domain, for example, medicine or IT, you can create a <u>language model</u>
 - If your app is used in a noisy environment, such as a factory, you can create a custom <u>acoustic model</u>
- Transcribe call center recording
 - You can easily index the text for <u>full-text search</u> or apply <u>Text Analytics</u> to detect sentiment, language, and key phrases.
 - · You can also use language mode and acoustic model
- Implement voice bots

Translator Text

Microsoft



Microsoft Translator API is an ISO and HIPAA compliant neural machine translation (NMT) service that developers can easily integrate into their applications, websites, tools, or any solution requiring multi-language support such as company websites, e-commerce sites, customer support applications, messaging applications, internal communication, and more.

Extend the reach of your applications

Translate text to and from 60+ supported languages through the open REST interface of Translator API.

Transliterate into different alphabets

Display text in different alphabets to make it easier to read - translate from Chinese characters to PinYin, display any of the supported transliteration languages in the Latin alphabet, and even show words written in the Latin alphabet in non-Latin characters such as Arabic, Hindi or Japanese.

Bilingual dictionary

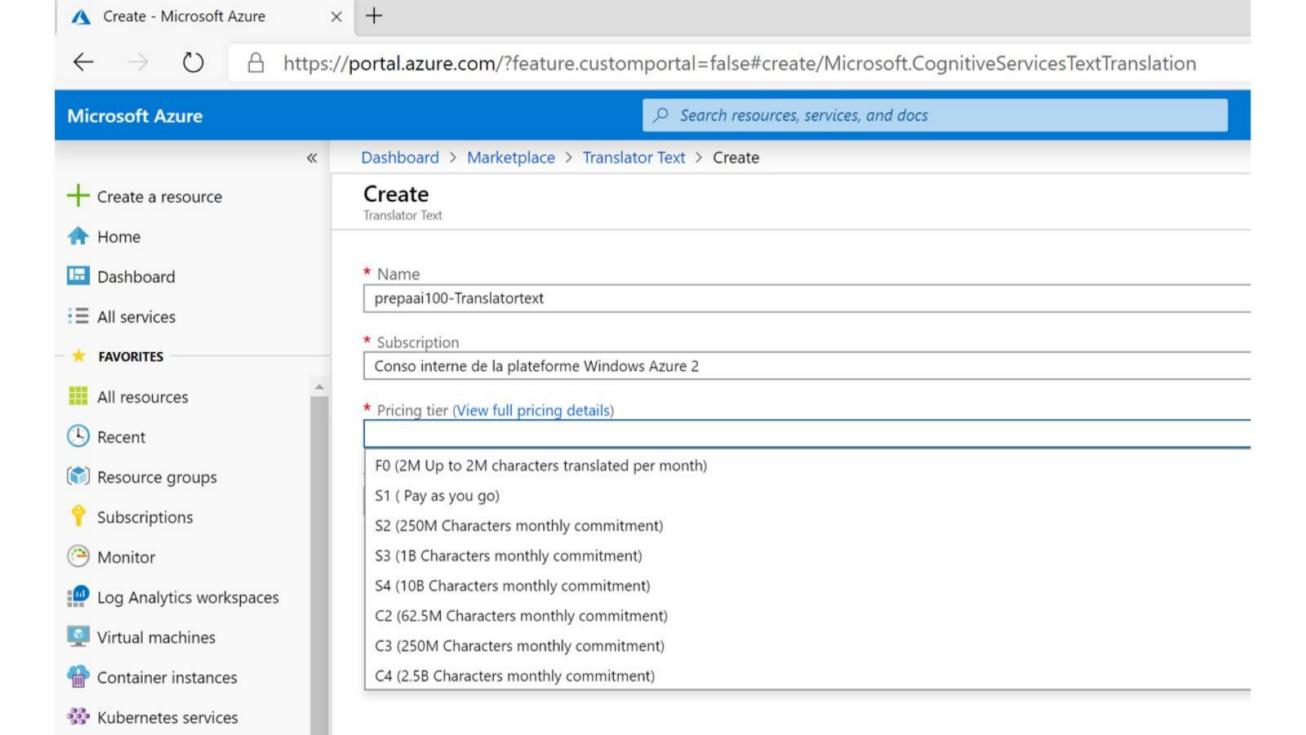
Find alternative translations for words from or to English and the context for alternate word translations by giving examples of human translated sentences.

Add both online and offline neural machine translation into your Android app

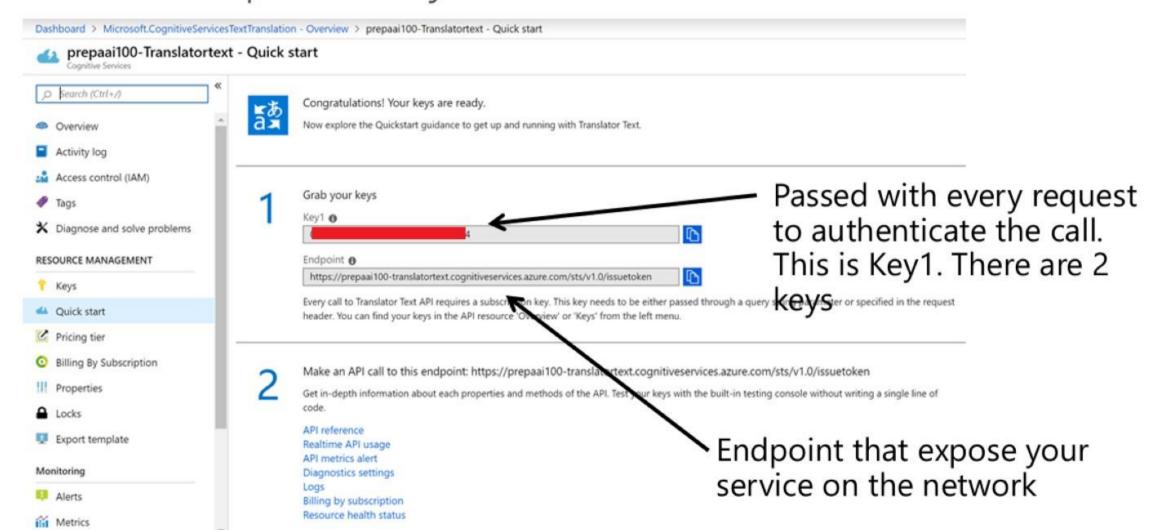
With the Microsoft Translator app local feature (Preview) on Android, easily add online and offline neural translation to your apps. Call the Translator app to integrate text translation capabilities into your own app. The Translator app will return the translations by using the Translator cloud service when the device is online or the offline packs when it is offline. It will use your existing Translator Text API subscription for billing.

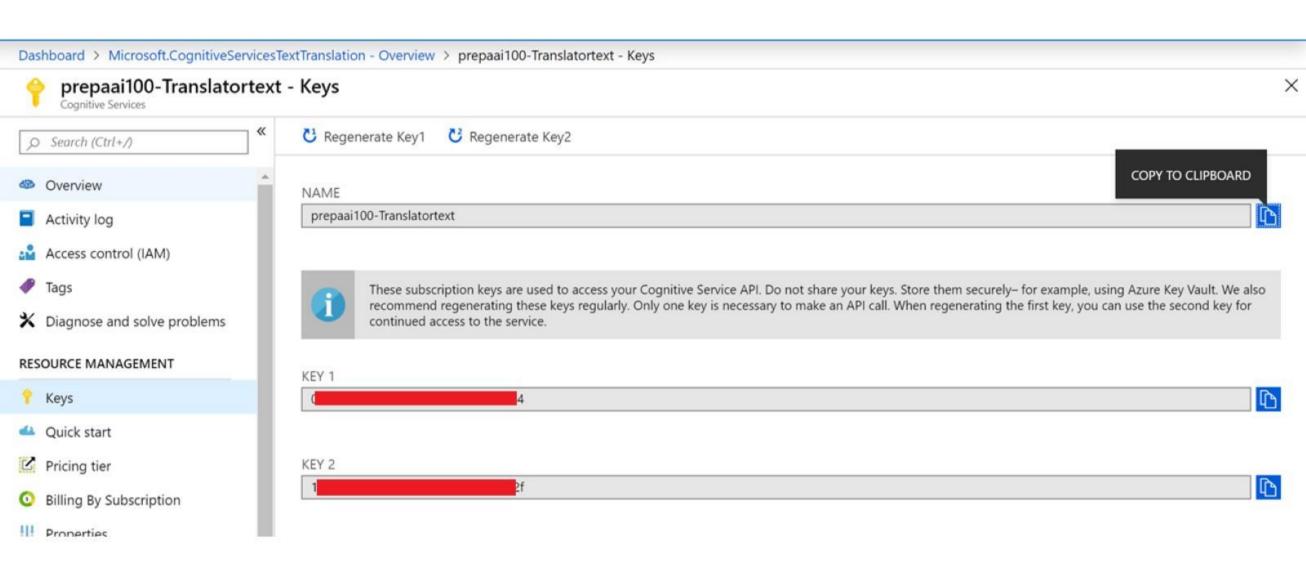
Automatically detect languages

Easily and accurately detect the language of any text string to decide which one needs to be translated and into which language(s).



 With Speech Translation API subscription created your are now able to access your API endpoint and subscription key





Develop solutions by using intelligent algorithms related natural language processing

- Custom acoustic model
- Custom language model

Custom acoustic model

- Creating a custom acoustic model is helpful if your application is designed for use in an environment, such as a noisy factory, or by a particular user population.
- Walkthrough
 - Prepare the data: a set of audio files (.WAV) of speech data and a text file of transcriptions of each audio file.
 - Data should be representative
 - 2. Import the acoustic data set
 - 3. Create the custom acoustic model based on an existing model
 - 1. Microsoft Search & Dictation AM or Microsoft Conversational model

Custom language model

- Create a custom language model for text queries or utterances you expect users to say or type in an application.
 - You can then use this custom language model in conjunction with existing stateof-the-art speech models from Microsoft to add voice interaction to your application

Walkthrough

- Prepare the data: provide a list of example utterances to the system (=specific vocabulary link to usage ex: medical, ...) in plain text
- 2. Import the language data set
- 3. Create the custom language model

Language APIs

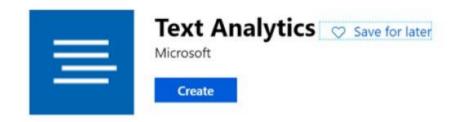
- Bing Spell Check
 - · Contextual grammar and spell checking
- LUIS (Language Understanding Intelligent Service)
 - Custom machine learning intelligence to a user's conversational, natural language text to predict overall meaning and pull out relevant detailed information
- Translator Text
 - · Translator (more than 60 languages) & Custom translator
- Text Analytics
 - Help you to extract information from text. To identify: language, discover sentiment, extract key phrases, detect well-known entities. Score between 0 (negative) to 1 (positive)

Test Text Analytics API

Dashboard > New > Text Analytics

Text Analytics

Microsoft



Use a few lines of code to easily analyze sentiment, extract key phrases, recognize entities and detect language for any kind of text.

Sentiment analysis

Find out what users think about your products or services. You can score the overall sentiment of text.

Key phrase extraction

Automatically extract key phrases from text to quickly identify the main points.

Language detection

Determine what language a piece of text is written in-from 120 supported languages.

Named Entity Recognition

Identify and categorize entities in your text as people, locations, organizations, date/time, currencies and more. Well-known entities are also recognized and linked to information available on the web.

Useful Links

More about Text Analytics API

Bing Search

- The Bing Web Search API is a RESTfull API that provides an experience similar to Bing.com/Search by returning search results that Bing determines are relevant to the user's query.
 - The results may include Web pages, images, videos, news, and entities, along with related search queries, spelling corrections, time zones, unit conversion, translations, and calculations.
 - · Search response : a **JSON** file with the following format :

```
"_type": "SearchResponse",
    "queryContext": {...},
    "webPages": {...},
    "images": {...},
    "relatedSearches": {...},
    "videos": {...},
    "news": {...},
    "spellSuggestion": {...},
    "computation": {...},
    "timeZone": {...},
    "rankingResponse": {...}
}, ...
```

Bing Search

 If you're interested in specific types of content such as images, videos, and news, you can request only those answers by using the responseFilter query parameter

```
" type": "SearchResponse",
"queryContext": {
  "originalQuery": "Microsoft Cognitive Services"
"webPages": {
  "webSearchUrl": "https://www.bing.com/search?q=Microsoft+cognitive+services",
  "totalEstimatedMatches": 22300000,
  "value": [
     "id": "https://api.cognitive.microsoft.com/api/v7/#WebPages.0",
      "name": "Microsoft Cognitive Services",
      "url": "https://www.microsoft.com/cognitive-services",
      "displayUrl": "https://www.microsoft.com/cognitive-services",
      "snippet": "Knock down barriers between you and your ideas. Enable natural and conte
      "deepLinks": [
          "name": "Face API",
          "url": "https://azure.microsoft.com/services/cognitive-services/face/",
          "snippet": "Add facial recognition to your applications to detect, identify, and
          "name": "Text Analytics",
          "url": "https://azure.microsoft.com/services/cognitive-services/text-analytics/"
          "snippet": "Cognitive Services; Text Analytics API; Text Analytics API . Detect
```

```
GET https://api.cognitive.microsoft.com/bing/v7.0/search?
q=sailing+dinghies&responseFilter=images%2Cvideos%2Cnews&mkt=en-us HTTP/1.1
Ocp-Apim-Subscription-Key: 123456789ABCDE
User-Agent: Mozilla/5.0 (compatible; MSIE 10.0; Windows Phone 8.0; Trident/6.0; IEMobile/10.0;
ARM; Touch; NOKIA; Lumia 822)
X-Search-ClientIP: 999.999.999
X-Search-Location: 47.60357;long:-122.3295;re:100
```

Hands on lab - Classify and moderate text with Azure Content Moderator

Go to this link - https://docs.microsoft.com/en-us/learn/modules/classify-and-moderate-text-with-azure-content-moderator/1-introduction

Bing Search

- Bing News Search
- Bing Web Search
- Bing Video Search
- Bing Autosuggest
 - Need to send a request to the API each time a user types a new characters in search box
- Bing Custom Search
- Bing Entities Search
 - Returns entities and place

Bing Search

- Bing Images Search
- Bing Visual Search
 - · Upload and image and API identify variety of details about it
- Bing Local Business Search

Recommendations and decision making Custom Decision Service

- In a typical web or mobile application, a front page links to several articles or other types of content. As the front page loads, it could request the Custom Decision Service to rank articles included on that front page
 - So, when a user chooses an article by clicking on it, a second request could be sent to the Custom Decision Service, which will log the outcome of that user decision
- Custom Decision Service is easy to use, as it requires only an RSS feed for your content and a few lines of JavaScript to be added into your application

Custom Decision Service

- Custom Decision Service converts your content into features for machine learning
- Some common-use cases for Custom Decision Service include:
 - · Personalizing articles on a news website
 - · Personalizing video content on a media portal
 - · Optimizing ad placements or web pages that the ad directs to
 - · Ranking recommended items on a shopping website
 - Anomaly detector

Custom Decision Service

 Custom Decision Service can be applied to both webpages and mobile apps. The APIs can be called from either a browser or an app Part 2

Map security requirements to tools, technologies, and processes



Azure Bot Service

- · Is now compliant with the following
 - · ISO 27001:2013
 - · ISO 27018:2014
 - PCI DSS
 - Microsoft HIPAA BAA
 - Health Insurance Portability and Accountability Act Business Associate Agreement
 - · SOC 1 & SOC 2

Enable robust GDPR compliance

- New offerings include the general availability of the Azure GDPR Data Subject Request (DSR) portal, Azure Policy, Compliance Manager for GDPR, Data Log Export, and the Azure Security and Compliance Blueprint for GDPR
 - · Azure Data Subject Request (DSR) portal enables you to fulfill GDPR requests
 - · Azure Policy enables you to set policies to conform to the GDPR
 - Extend Azure Policies for the GDPR into Azure Security Center
 - The Azure Security and Compliance GDPR Blueprint accelerates your GDPR deployment
 - · Compliance Manager for Azure helps you assess and manage GDPR compliance

Part 3

Select the software, services, and storage required to support a solution



Al reference Architecture

<u>https://docs.microsoft.com/en-us/azure/architecture/browse/#ai--machine-learning</u>

When you need Custom AI?

- Data restrictions
- Algorithm limitations
- Network contraints
- Domain specific needs

→ Azure Machine Learning Service is Microsoft offering for Custom Al

Machine Learning at scale

- Before deciding which ML services to use in training and operationalization, consider whether you need to train a model at all, or if a prebuilt model can meet your requirements. In many cases, using a prebuilt model is just a matter of calling a web service or using an ML library to load an existing model.
- Some options include:
 - · Use the web services provided by Azure Cognitive Services.
 - · Use the pretrained neural network models provided by the Cognitive Toolkit.
 - · Embed the serialized models provided by Core ML for an iOS app

Machine Learning at scale

- · If a prebuilt model does not fit your data or your scenario, options in Azure include Azure Machine Learning, HDInsight with Spark MLlib and MMLSpark, Azure Databricks, Cognitive Toolkit, and SQL Machine Learning Services.
- If you decide to use a custom model, you must design a pipeline that includes model training and operationalization

Machine Learning at scale

- Machine learning at scale addresses two different scalability concerns:
 - The first is training a model against large data sets that require the scale-out capabilities of a cluster to train
 - The second centers on operationalizing the learned model so it can scale to meet the demands of the applications that consume it. Typically this is accomplished by deploying the predictive capabilities as a web service that can then be scaled out

Challenges with Machine learning at scale

- You typically need a lot of data to train a model, especially for deep learning models
- You need to prepare these big data sets before you can even begin training your model
- The model training phase must access the big data stores. It's common to perform the model training using the same big data cluster, such as Spark, that is used for data preparation
- For scenarios such as deep learning, not only will you need a cluster that can provide you scale-out on CPUs, but your cluster will need to consist of GPU-enabled nodes

Choosing a natural language processing technology in Azure

- In Azure, the following services provide natural language processing (NLP) capabilities:
 - Azure HDInsight with Spark NLP and Spark Mllib
 - · If you need to train custom models against a large corpus of text data
 - If you need low-level NLP capabilities like tokenization, stemming, lemmatization, and term frequency/inverse document frequency (TF/IDF)
 - Azure Databricks
 - Microsoft Cognitive Services :
 - · if you want prebuilt models
 - If you need simple, high-level NLP capabilities like entity and intent identification, topic detection, spell check, or sentiment analysis

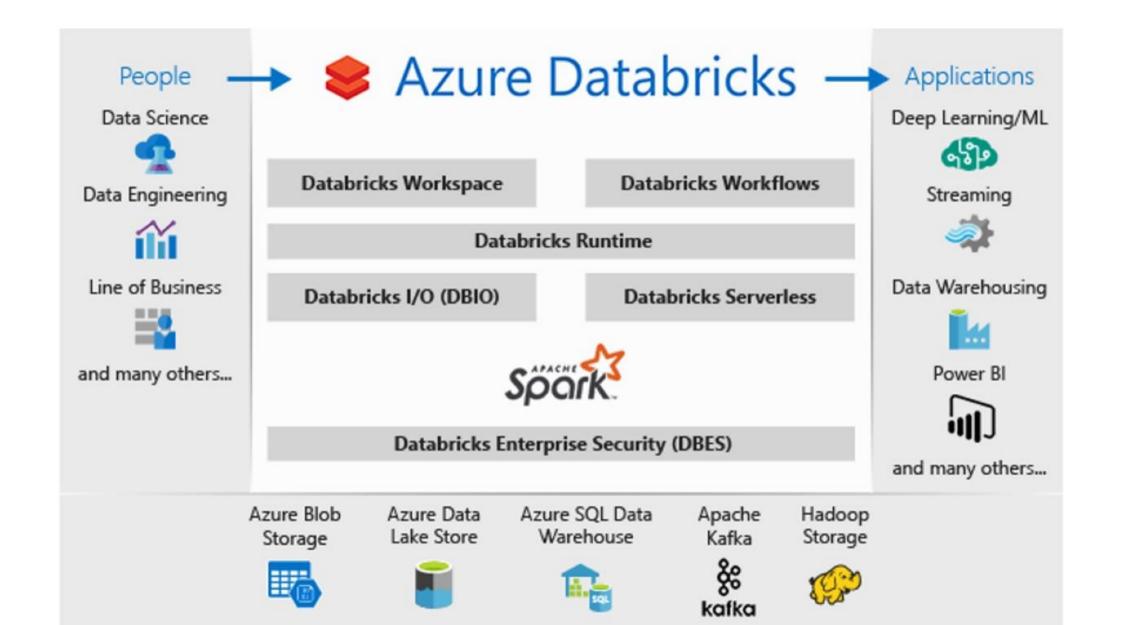
Azure HDInsight vs Cognitive Services

Capability	Azure HDInsight	Microsoft Cognitive Services
Provides pretrained models as a service	No	Yes
REST API	Yes	Yes
Programmability	Python, Scala, Java	C#, Java, Node.js, Python, PHP, Ruby
Support processing of big data sets and large documents	Yes	No

Azure Databricks

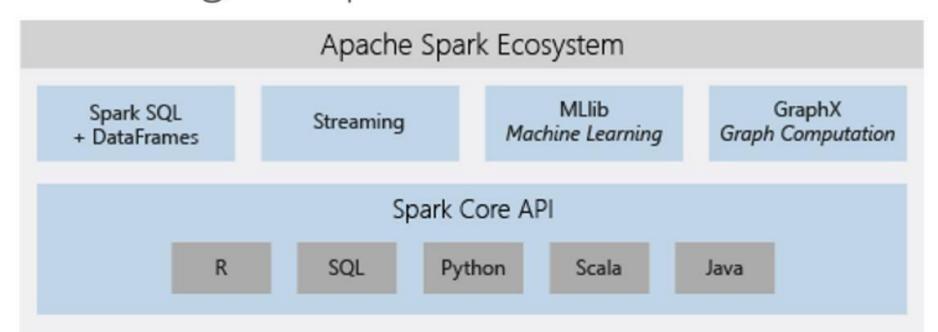
- Azure Databricks is an Apache Spark-based analytics platform optimized for the Microsoft Azure cloud services platform
 - Designed with the founders of Apache Spark, Databricks is integrated with Azure to provide one-click setup, streamlined workflows, and an interactive workspace that enables collaboration between data scientists, data engineers, and business analysts

Azure Databricks



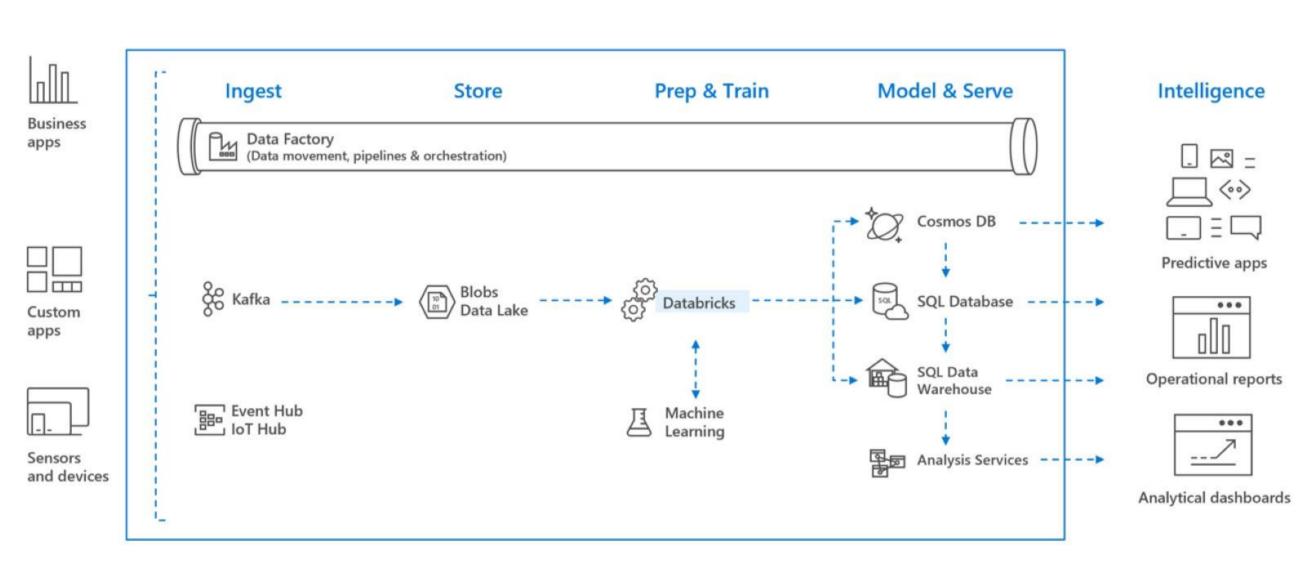
Apache Spark-based analytics platform

 Azure Databricks comprises the complete opensource Apache Spark cluster technologies and capabilities. Spark in Azure Databricks includes the following components



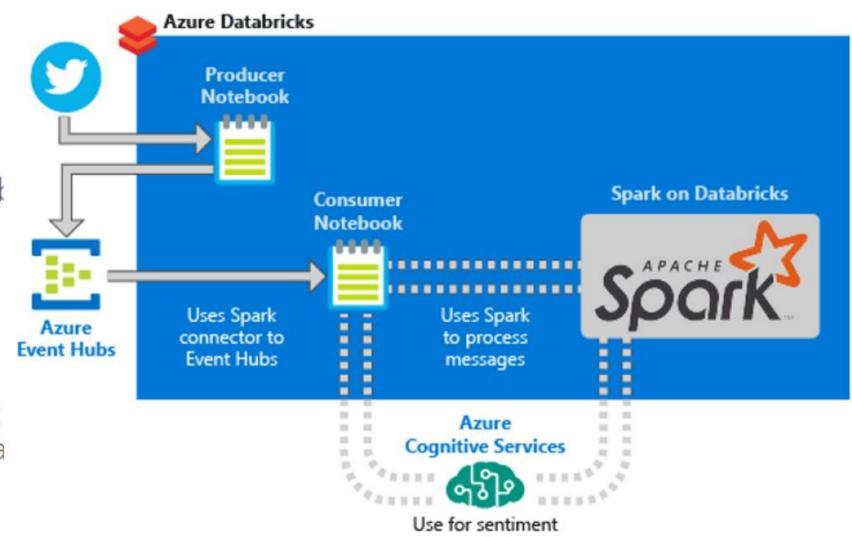
Azure Databricks

- Azure Databricks is a fast, easy, and collaborative Apache Spark-based analytics service
- For a big data pipeline, the data (raw or structured) is ingested into Azure through Azure Data Factory in batches, or streamed near real-time using Kafka, Event Hub, or IoT Hub
- This data lands in a data lake for long term persisted storage, in Azure Blob Storage or Azure Data Lake Storage
- As part of your analytics workflow, use Azure Databricks to read data from multiple data sources such as <u>Azure Blob</u> <u>Storage</u>, <u>Azure Data Lake Storage</u>, <u>Azure Cosmos DB</u>, or <u>Azure SQL Data Warehouse</u> and turn it into breakthrough insights using Spark



Example: Databrick + cognitives services

- Run sentiment analysis on a stream of data using Azure Databricks in near real time
- You set up data ingestion system using Azure Event Huł
- You consume the messages from Event Hubs into Azure Databricks using the Spark Event Hubs connector
- Finally, you use Cognitive Service APIs to run sentiment analysis on the streamed data



analysis

https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-sentiment-analysis-cognitive-services

Training of Python scikit-learn and deep learning models on Azure

 This reference architecture shows recommended practices for tuning the hyperparameters (training parameters) of python models. Two scenarios are covered: hyperparameter tuning of scikit-learn models and deep learning models with GPUs

Training and test data and script

Training results

Microsoft Data Science Virtual Machine (DSVM)

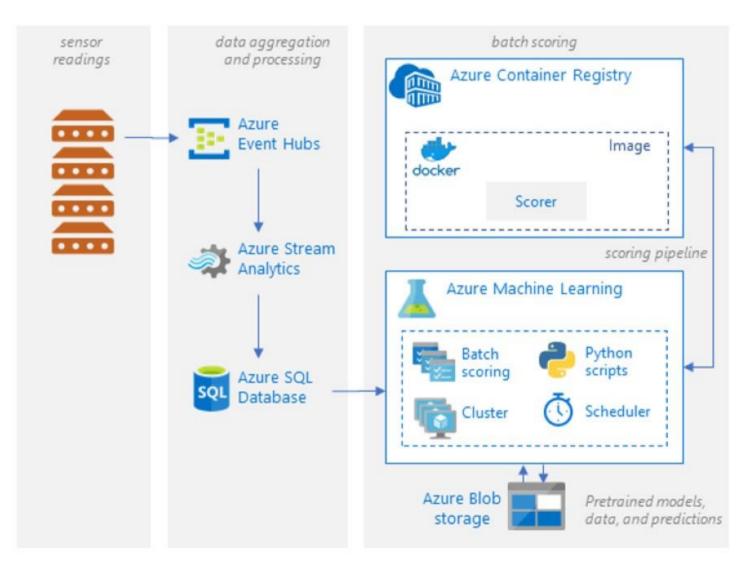
Azure Storage Training Blobs results Script execution commands Status Azure Machine Azure Machine Learning service Learning compute Docker container

Azure Container Registry

https://docs.microsoft.com/enus/azure/architecture/referencearchitectures/ai/training-python-models

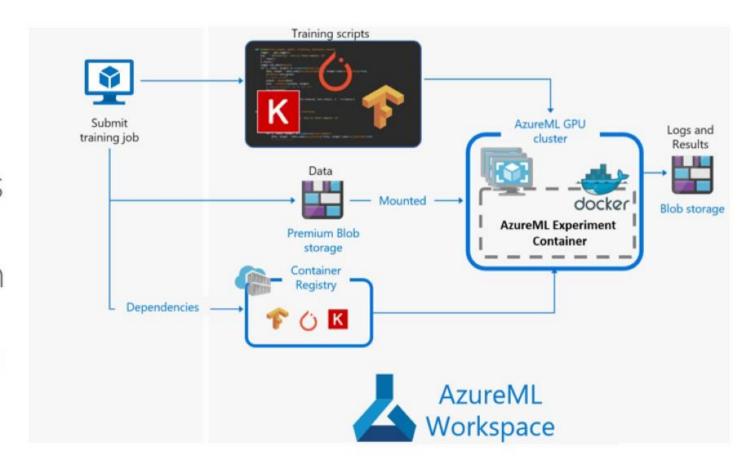
Batch scoring of Python machine learning models on Azure

 This reference architecture shows how to build a scalable solution for batch scoring many models on a schedule in parallel using Azure Machine Learning Service.



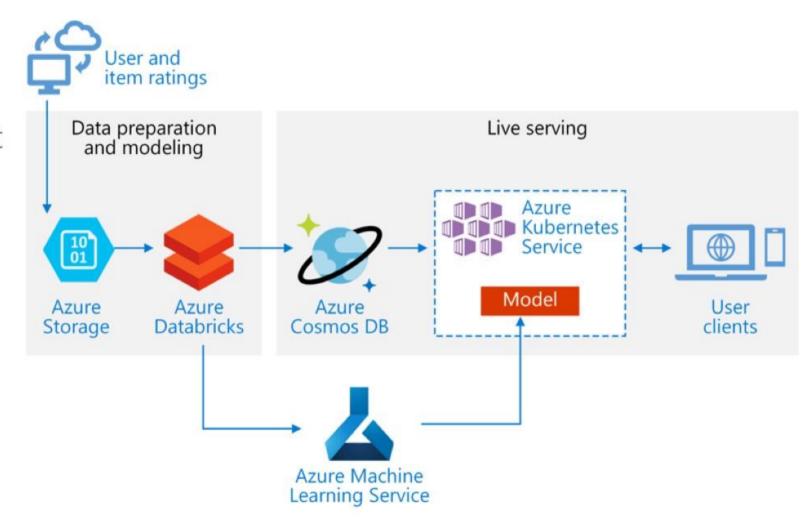
Distributed training of deep learning models on Azure

- This reference architecture shows how to conduct distributed training of deep learning models across clusters of GPU-enabled VMs
- The scenario is image classification, but the solution can be generalized for other deep learning scenarios such as segmentation and object detection

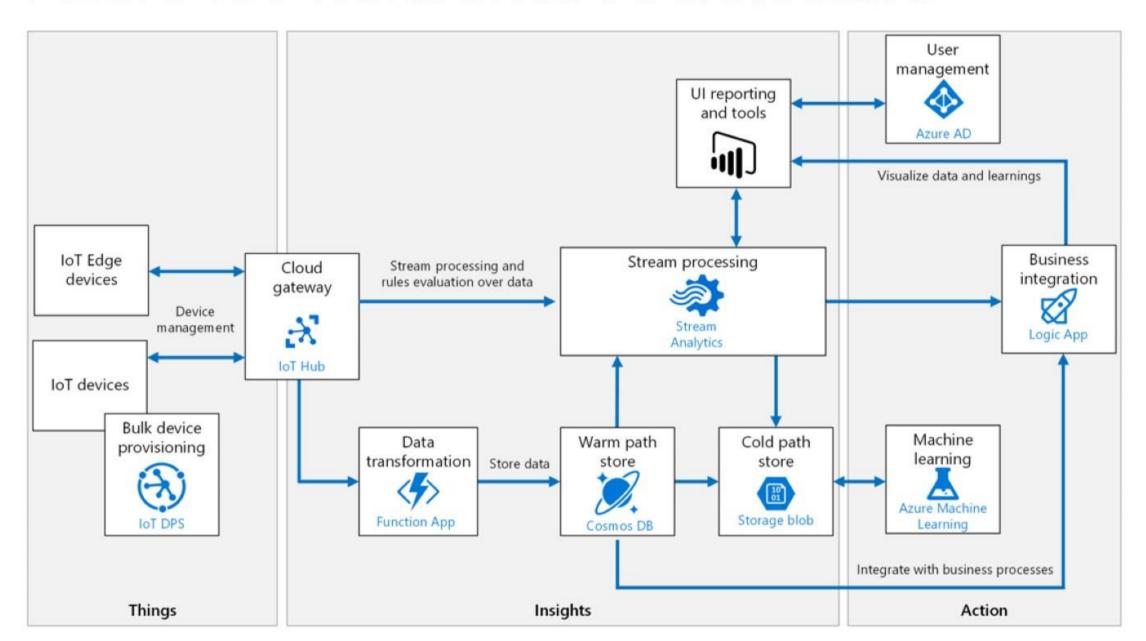


Build a real-time recommendation API on Azure

- This reference architecture shows how to train a recommendation model using Azure Databricks and deploy it as an API by using Azure Cosmos DB, Azure Machine Learning, and Azure Kubernetes Service (AKS)
- This architecture can be generalized for most recommendation engine scenarios, including recommendations for products, movies, and news



Azure IoT Reference Architecture

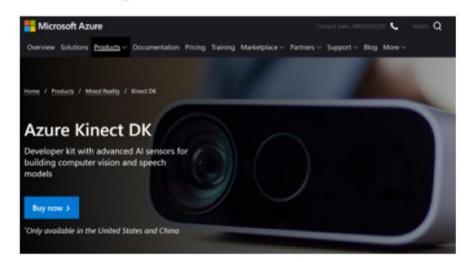


Process IoT telemetry data

- There are two ways to process telemetry data, hot path and cold path. The difference has to do with requirements for latency and data access
 - The **hot path** analyzes data in near-real-time, as it arrives. In the hot path, telemetry must be processed with very low latency. The hot path is typically implemented using a stream processing engine. The output may trigger an alert or be written to a structured format that can be queried using analytical tools
 - · For warm path storage, Microsoft recommend Cosmo DB
 - The cold path performs batch processing at longer intervals (hourly or daily). The
 cold path typically operates over large volumes of data, but the results don't need
 to be as timely as the hot path. In the cold path, raw telemetry is captured and
 then fed into a batch process

What is Edge computing?

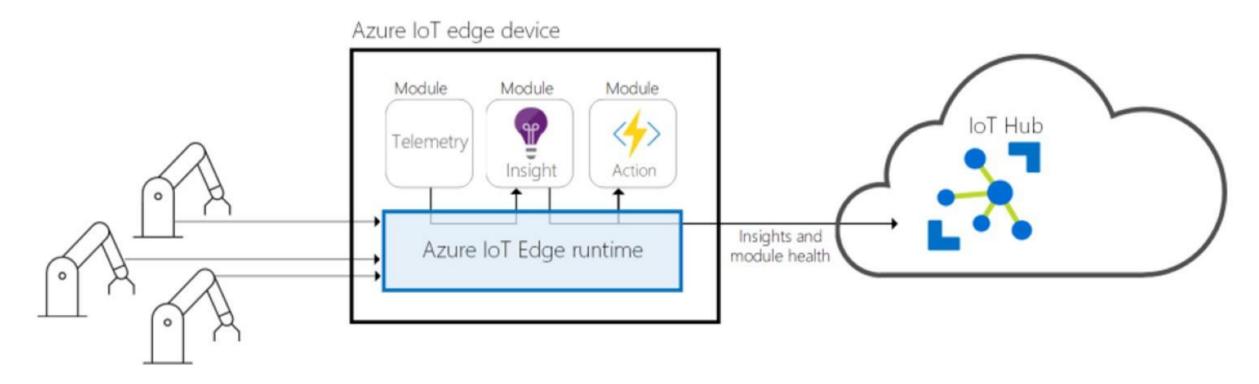
- Set of connected sytems and devices that gather and analyze/process data-close to the users/devices
- Edge computing solves problems of latency, bandwidth, offline support, connectivity, resource utilization, Intellectual Property Security
- Examples in Microsoft portfolio
 - Azure Databox Edge
 - Azure Stack
 - Azure IoT Edge
 - Azure Kinect DK



Edge Tiers

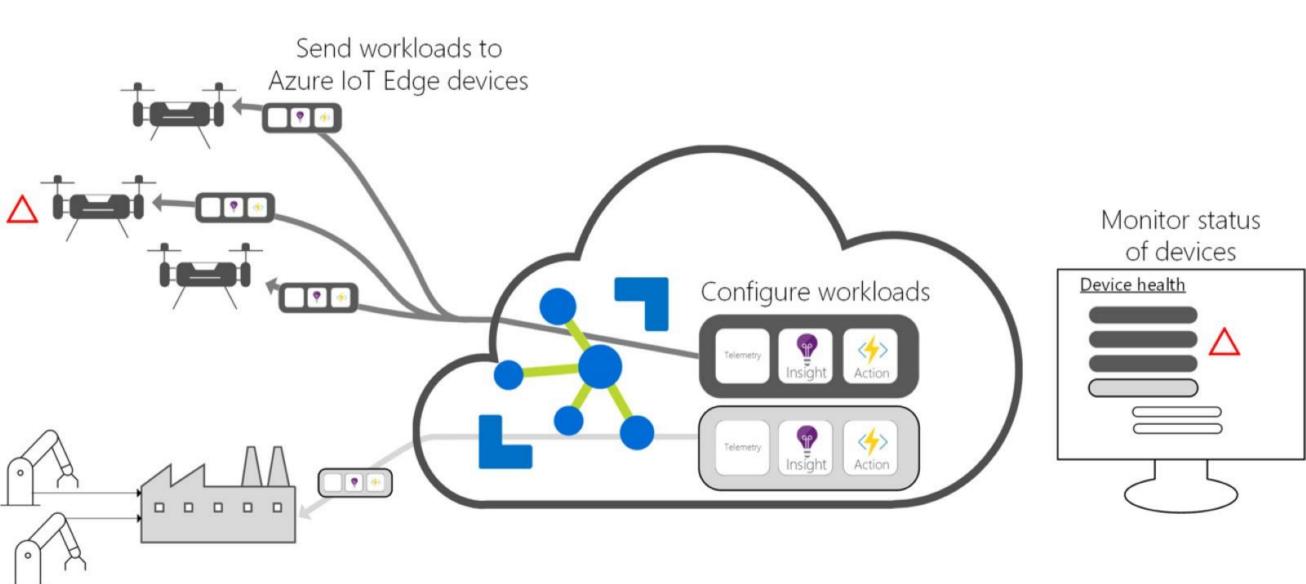
- Datacenter
- Server
- Gateway (ex: protocol gateway)
- Device

Azure IoT Edge

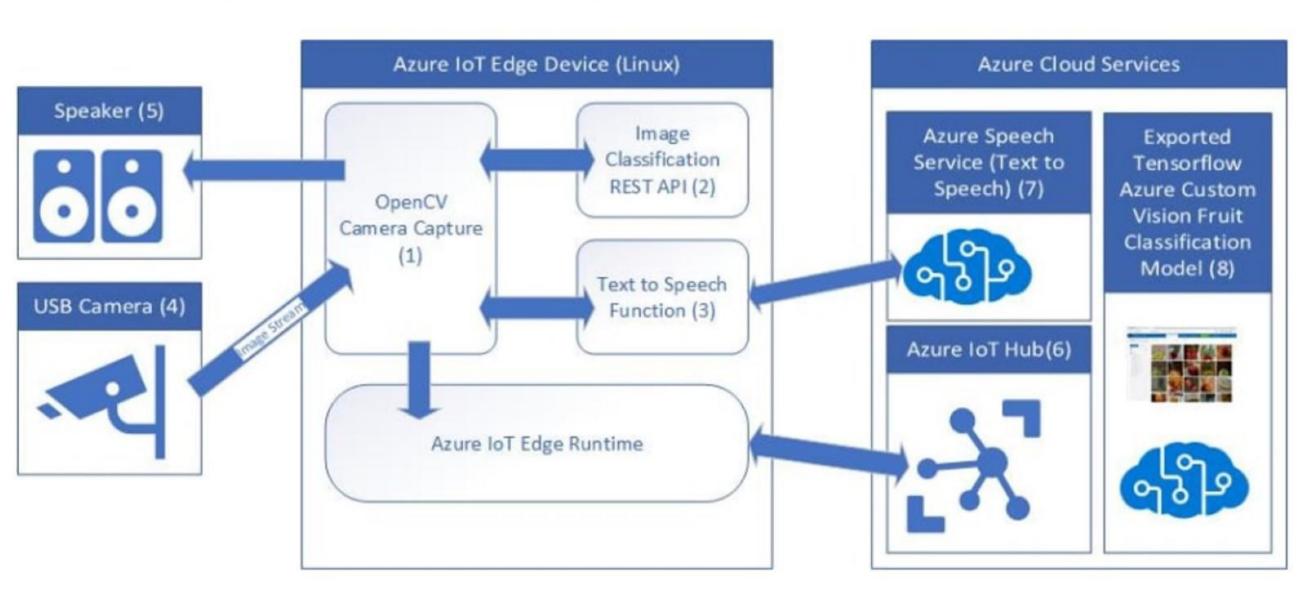


https://docs.microsoft.com/en-us/azure/iot-edge/about-iot-edge

Example of IoT Edge solution

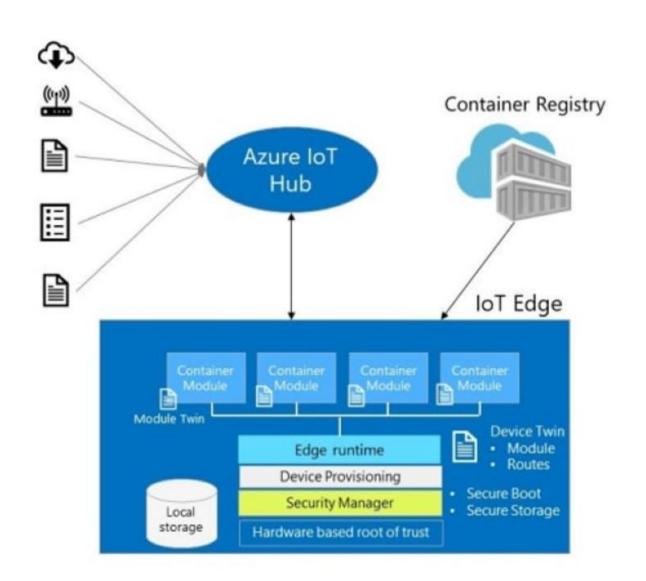


Example of IoT Edge solution



Azure Edge in action

- Develop your Container Workloads
- 2. Upload Modules to a Container Registry
- 3. Provision IoT Edge Device Agents
- 4. Select Device/Deployment Set
- 5. Define Modules
- 6. Define Message Routes
- 7. Define Module Parameters



Home > Resource groups	> RG-Prepa-AI-100 >	Marketplace >	IoT Hub > IoT hub
------------------------	---------------------	---------------	-------------------

IoT hub

Microsoft

Basics Size and scale Review + create

Create an IoT Hub to help you connect, monitor, and manage billions of your IoT assets. Learn More

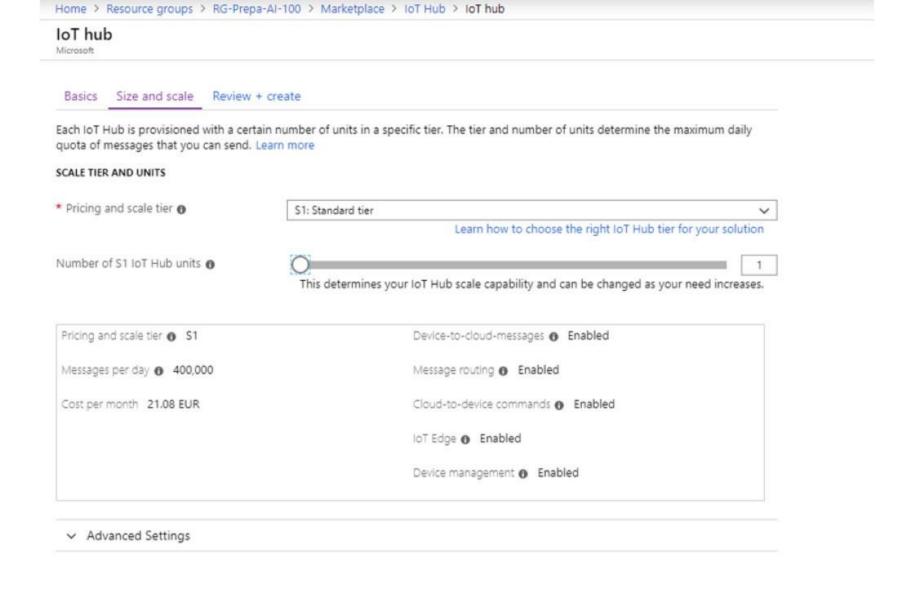
PrepaAl100

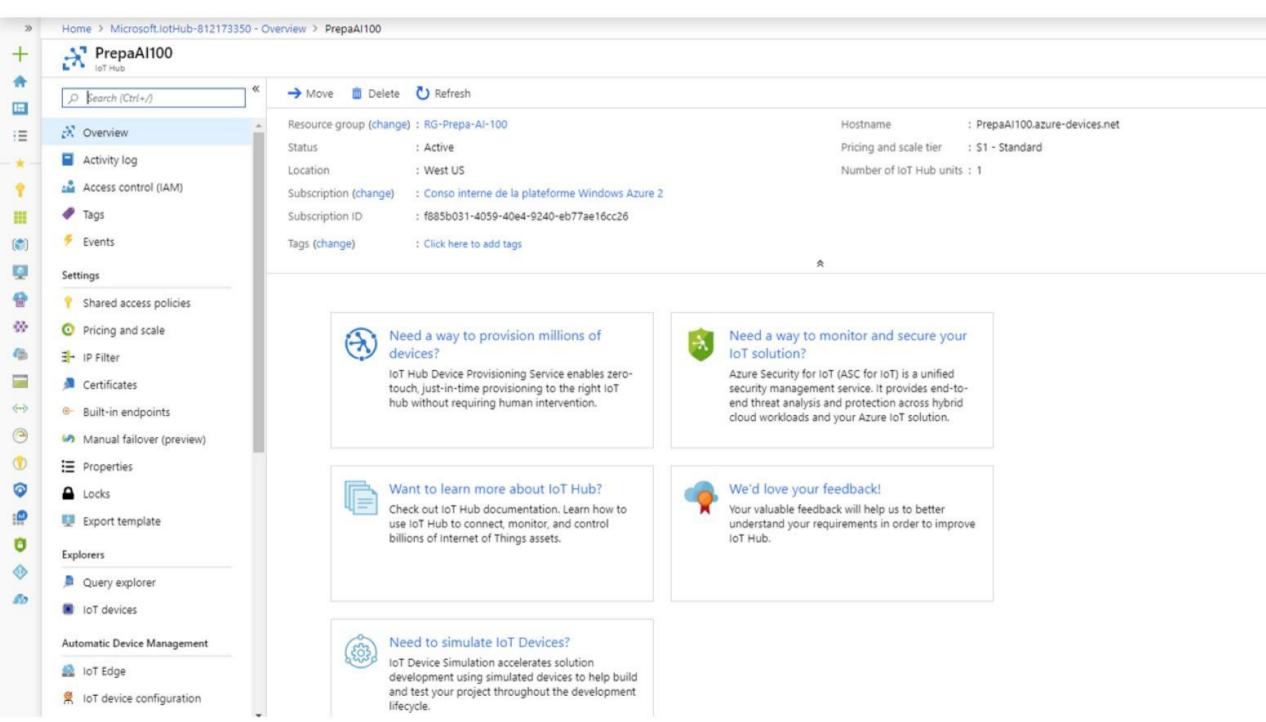
PROJECT DETAILS

* IoT Hub Name 6

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

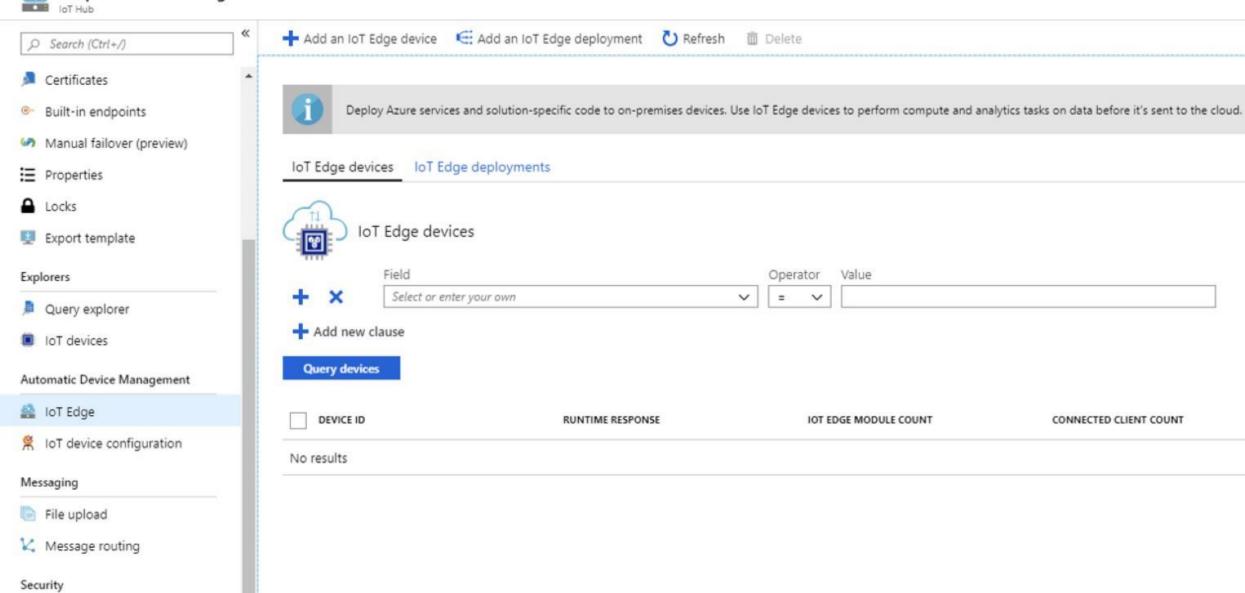
* Subscription •	Conso interne de la plateforme Windows Azure 2	~
* Resource Group 1	RG-Prepa-Al-100	~
	Create new	
* Region 🚯	West US	~

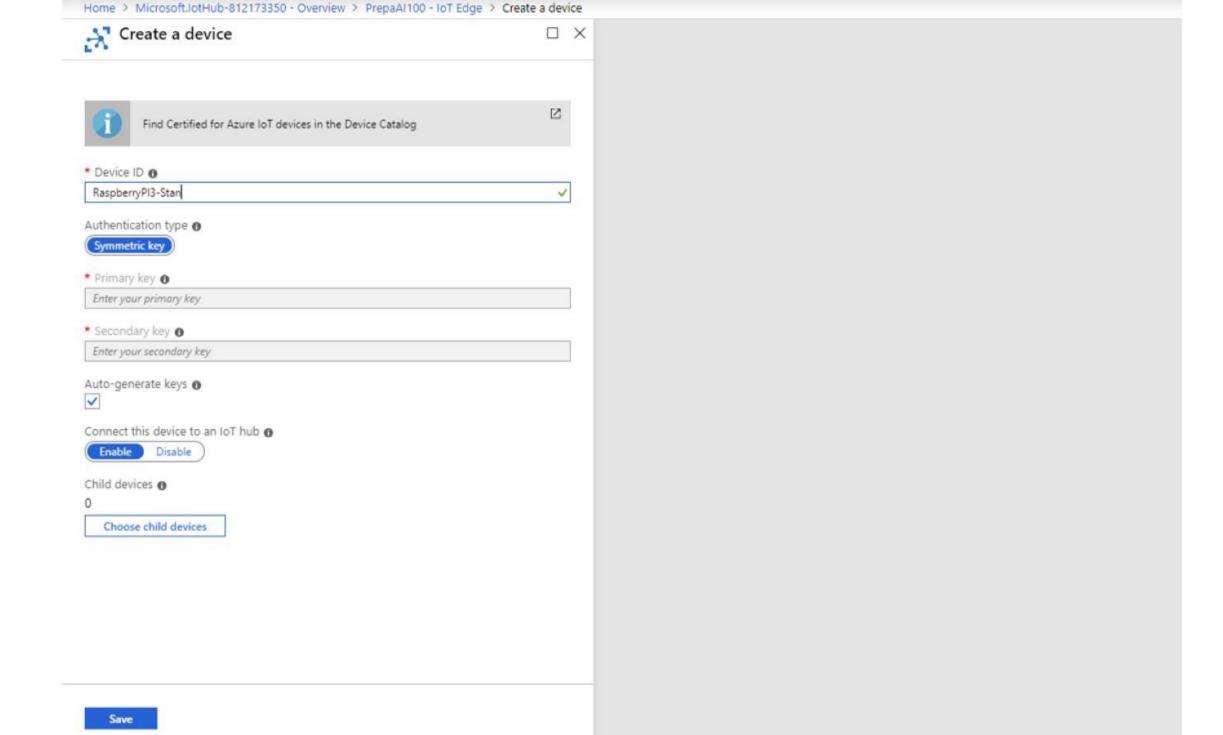


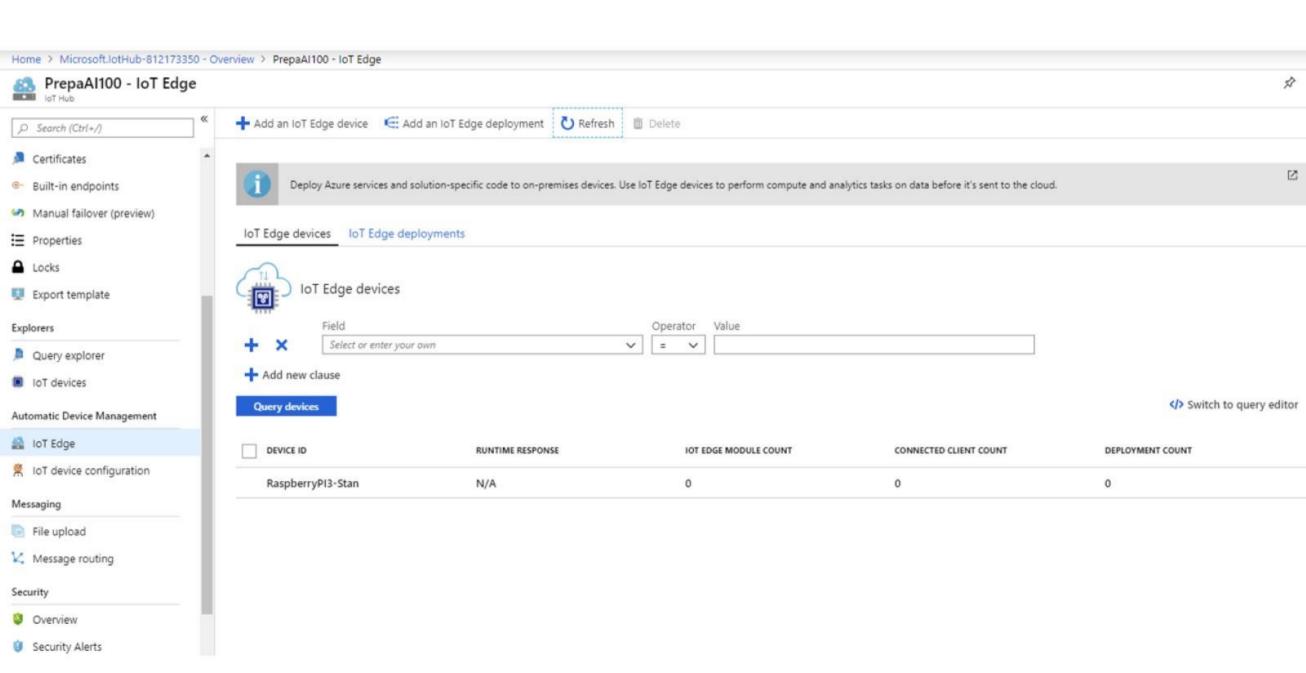


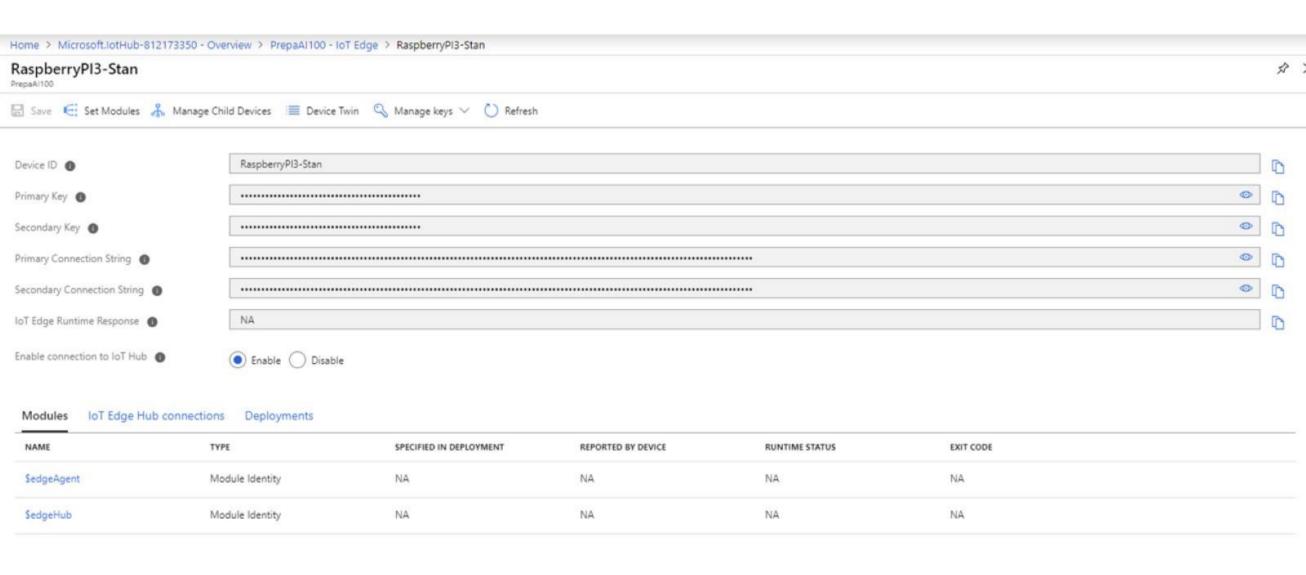


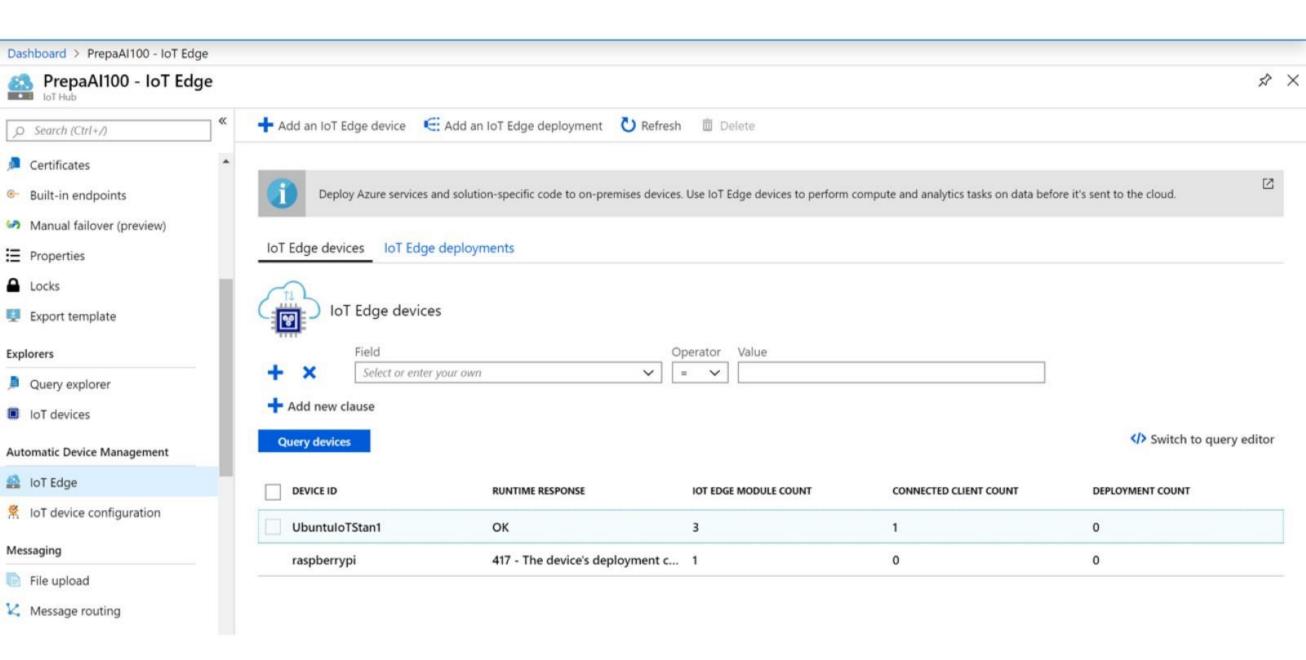
PrepaAl100 - IoT Edge











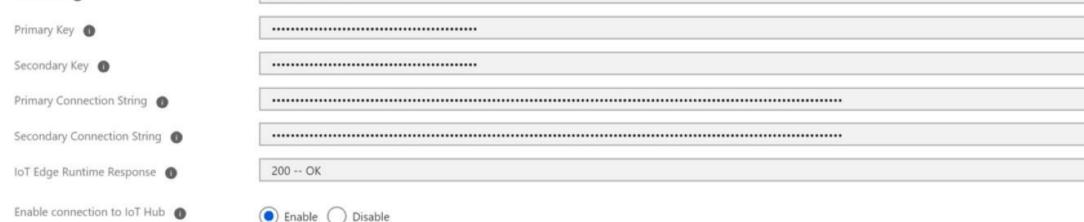
UbuntuloTStan1

PrepaAl100

Device ID 0

-	Set Modules	age.	Manage Critic Devices	=	Device IWIN	40	ivianage keys	0	Kerresn

UbuntuloTStan1



Modules	IoT Edge Hub connections	Deployments
---------	--------------------------	-------------

NAME	TYPE	SPECIFIED IN DEPLOYMENT REPORTED BY DEVICE		RUNTIME STATUS	EXIT CODE	
SedgeAgent	IoT Edge System Module	✓ Yes	✓ Yes	running	0	
SedgeHub	IoT Edge System Module	✓ Yes	✓ Yes	running	0	
SimulatedTemperatureSensor	IoT Edge Custom Module	✓ Yes	✓ Yes	running	0	

stan@UbuntuIoTStan1:~\$ sudo iotedge list SimulatedTemperatureSensor edgeAgent

edgeHub

STATUS DESCRIPTION running Up 2 hours running Up 2 hours running Up 2 hours

CONFIG

mcr.microsoft.com/azureiotedge-simulated-temperature-sensor:1.0 mcr.microsoft.com/azureiotedge-agent:1.0 mcr.microsoft.com/azureiotedge-hub:1.0

EdgeAgent, EdgeHub

· EdgeHub :

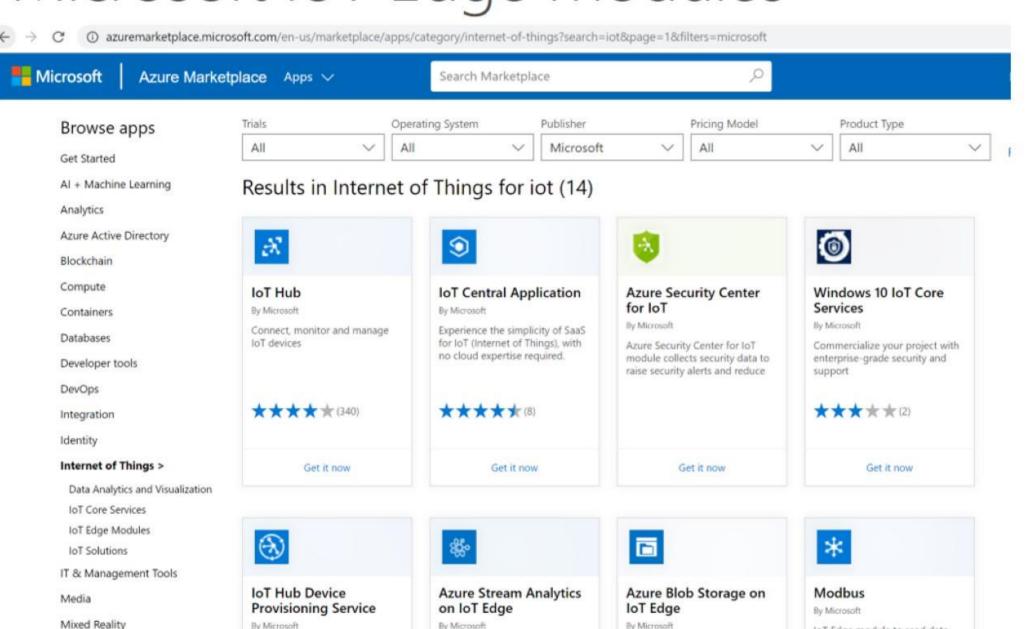
- local IoT Hub proxy
- Support AMQP and MQTT, not HTTP
- Optimizes connections to the Cloud
- Automatic local message storage
- Facilitate module communication

EdgeAgent

- · Instanciates module
- Ensures modules run
- Monitors modules status

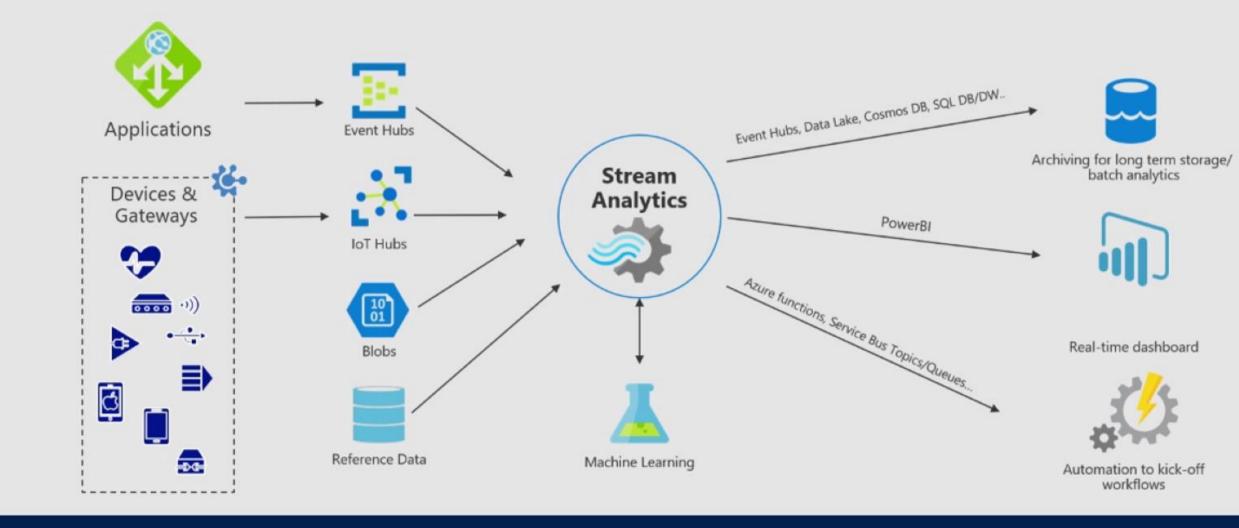
```
STATUS
                                                    CONFIG
                                      DESCRIPTION
SimulatedTemperatureSensor
                                      Up 2 hours
                                                    mcr.microsoft.com
                       running
edgeAgent
                       running
                                      Up 2 hours
                                                    mcr.microsoft.com
edgeHub
                       running
                                      Up 2 hours
                                                    mcr.microsoft.com
```

Microsoft IoT Edge modules

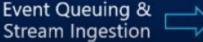


IoT Edge module to read data

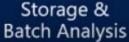
Stream Analytics: Canonical hot-path pipeline design













Azure Stream Analytics

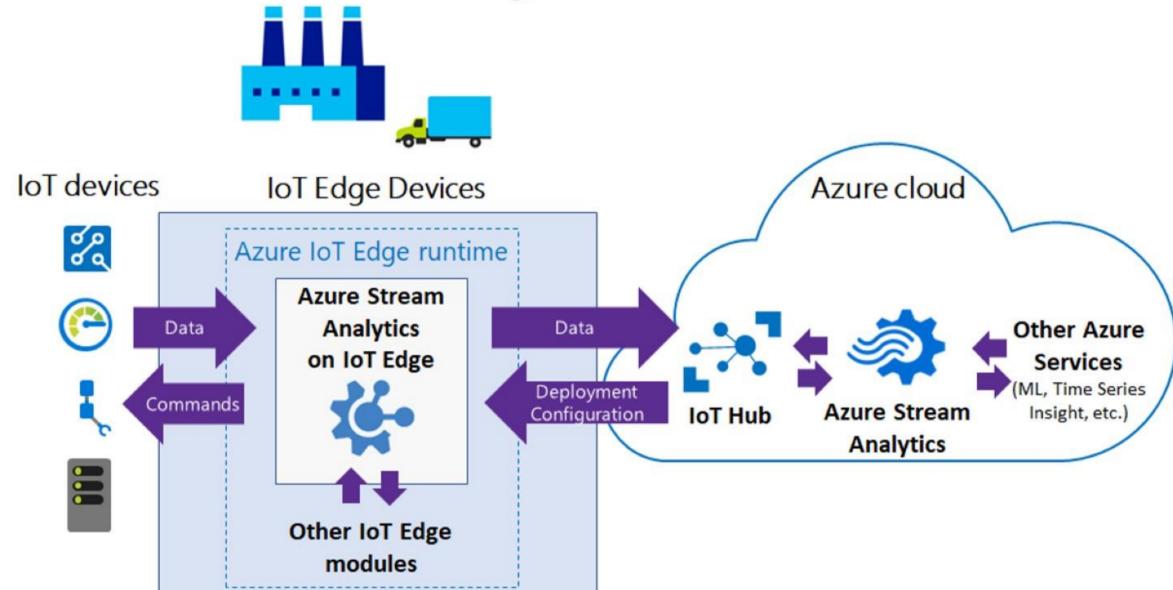
Unlocking Real-Time Insight

- · When insights are perishable and time to decision is critical
- · When you need to query data while it is still in motion

Easy to start

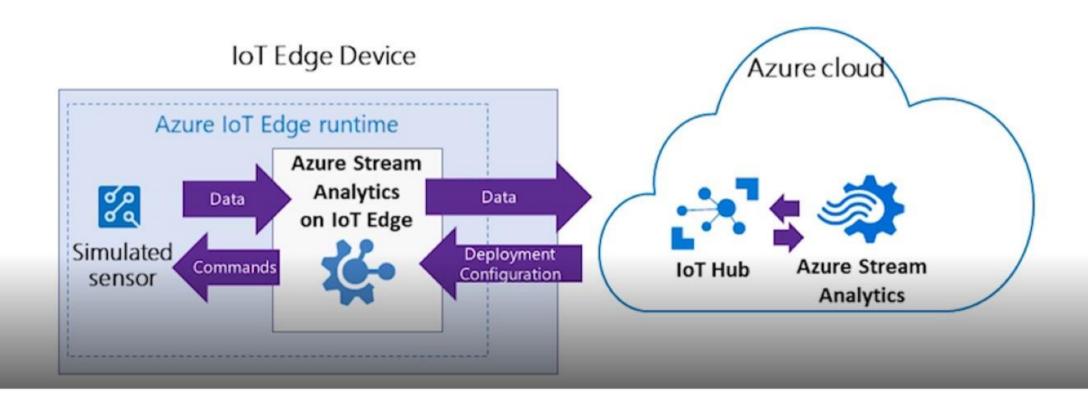
- Simple SQL Like Language
 - Inputs -> Query -> Outputs
- Integrate with the rest of the Azure ecosystem
- Can run on the Edge with Azure IoT Edge Runtime
 - But still managed from the Cloud
 - Same language for both Cloud and Edge « jobs »
 - Cross Platform: runs on Docker Container (Linux/Windows, x86/x64/ARM)

Azure Stream Analytics



Demo Architecture

- · <u>Simulated sensor</u>: sends temperature and humidity
- · ASA: triggers an alert if the temperature is above 25, and reset sensor



ASAEdgejob2

Query





```
▼ 🖰 Outputs (1)
```



Need help with your query? Check out some of the most common Stream Analytics query patterns here.

```
1 SELECT
2    'reset' AS command
3 INTO
4    alert
5 FROM
6    temperature TIMESTAMP BY timeCreated
7 GROUP BY TumblingWindow(second,2)
8 HAVING Avg(machine.temperature) > 25
9
```

Your query could be put in logs that are in a potentially different geography.

Missing some language constructs? Let us know! (Powered by UserVoice - Privacy Policy)

```
{"temperature":45.... {"temperature":21.... "2017-11-10T07:2... "2017-11-10T07:2... 0

{"temperature":45.... {"temperature":20.... "2017-11-10T07:2... "2017-11-10T07:2... 0

{"temperature":47.... {"temperature":20.... "2017-11-10T07:2... "2017-11-10T07:2... 0
```

Main edge Scenarios for Azure Stream Analytics

- Low-latency command and control
 - Azure Stream Analytics can analyze sensor data in near real time and issu commands when you detect anomalies
- Limited connectivity to the cloud
 - · Your Streaming logic runs independently of the network connectivity
- Limited bandwidth
 - With Azure Stream Analytics you can filter or aggregate data before sending it to the cloud
- Compliance
 - With Azure Stream Analytics you can anonymize or aggregate data before sending it to the cloud

Built-in ML based Anomaly Detection

Easily invoked with simple function calls

Un-supervised learning models

Support in-line learning and real-time scoring

2 broad categories of anomalies detected:

Temporary Anomalies: Spikes, Dips

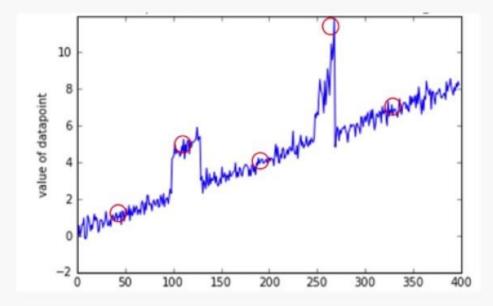
Persistent Anomalies: Slow +/- trends, Bi-level changes

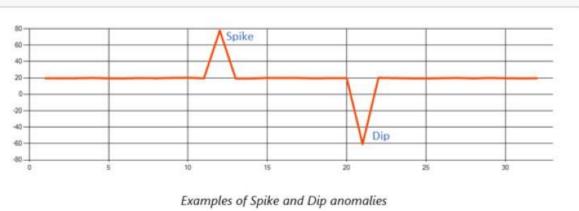
Customer benefits

Time to Value

Lower costs

Better performance than 'Bring Your Own Model'





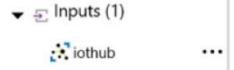
Dashboard > adASAMxChip > adASAMxChip

adASAMxChip

Query







```
idq [
```

Need help with your query? Check out some of the most common Stream Analytics query patterns here.

```
With ADStep as
 2
     SELECT
 4
         EVENTPROCESSEDUTCTIME as time,
 5
         AnomalyDetection_SpikeAndDip(z, 9%, 500, 'spikesanddips') OVER(LIMIT DURATION(ss, 120)) as scores
 6
     FROM
 8
         iothub
 9
10
11
     SELECT
12
         time,
13
         Z,
         CAST(GetRecordPropertyValue(scores, 'IsAnomaly') as bigint) as IsAnomaly,
14
15
         CAST(GetRecordPropertyValue(scores, 'Score') as float) as Score
16
     INTO
```

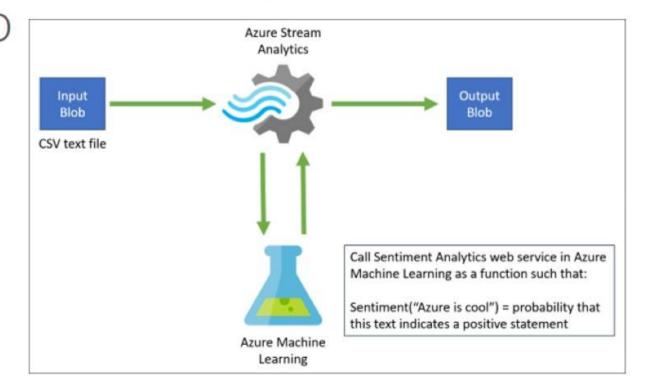
Your query could be put in logs that are in a potentially different geography.

Missing some language constructs? Let us know! (Powered by UserVoice - Privacy Policy)

Performing sentiment analysis by using Azure Stream Analytics and Azure Machine Learning Studio

Azure Stream Analytics job that integrates Azure

Machine Learning Studio



Stream analytics with Spark or Storm

- Another option for stream analytics is Apache Spark in Azure Databrick, or Apache Storm
- Along with the other projects of Apache such as Hadoop and Spark, Storm is one of the star performers in the field of data analysis

Situation	Spark	Storm
Stream processing	Batch processing	Micro-batch processing
Latency	Latency of a few seconds	Latency of milliseconds
Multi-language support	Lesser language support	Multiple language support

Identify storage required to store logging, bot state data, and Cognitive Services

Security and governance Azure AD Key Vault Identity and Keys and secrets access control authenticate. token, keys authorize Data ETL Raw data Bot cognition and intelligence Bot logic and UX Azure Functions Web App Custom serverless **Bot Service** request 4 queries compute Custom model for Input channels. Structured Intents and entities data data unstructured Q&A authentication CRM, SQL etc. Data Factory Scheduled ETL OnA Maker Bot Logic Azure Search response 5 results pipelines Knowledge base Custom bot code Search index Unstructured Logic Apps FAQs, PDFs. improved bot intelligence Data connectors DOCs. etc. end-to-end testing. Quality assurance and enhancements conversations, new features enhanced data feedback, logs Azure DevOps, VS Code collection Retraining, testing, logs debugging, Logging feedback analysis Bot testing utility Application Insights Web and system Monitoring and reporting **Bot logging utility** Cosmos DB Application Insights Power BI emails, alerts Dashboards and Conversations. Dashboards and logs Bot feedback feedback analytics analytics User initiated utility System initiated Azure Storage SendGrid Web App DevOps Conversations, Best-practice Custom dashboard Email service

utility sample

feedback, logs

Choose the right data store

- Relational database management options
 - Azure SQL DB
 - Azure Database for PostgreSQL
 - Azure Database for MySQL
 - Azure Database for MariaDB
 - SQL Server Big Data Clusters
- Key Value Stores
 - CosmosDB
 - Azure cache for Redis
- Document database
 - CosmosDB

Choose the right data store

- Graph database
 - CosmosDB
- Column family databases
 - CosmosDB
 - Azure cache for Redis
- Data analytics
 - SQL Data Warehouse (now called Synapse Analytics)
 - SQL Server Big Data Clusters
 - Azure Data Lake
 - Azure Data Explorer
- Search engine databases
 - Azure Search

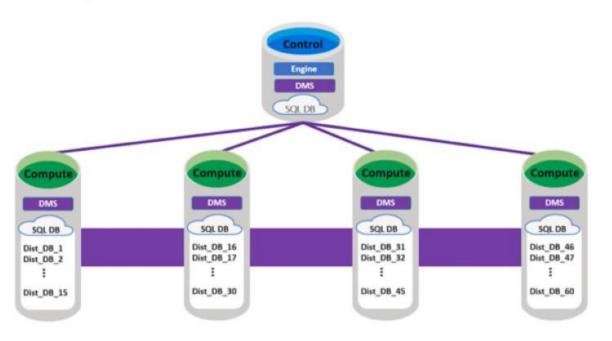
Choose the right data store

- Time Series Databases
 - Time Series Insights
- Object Storage
 - Blob Storage
- Shared files
 - File Storage

Azure SQL Data Warehouse

- SQL Data Warehouse is a cloud-based Enterprise Data Warehouse (EDW) that uses Massively Parallel Processing (MPP) to quickly run complex queries across petabytes of data
- Combines SQL relational databases with massively parallel processing
- Use SQL Data Warehouse as a key component of a big data solution. Import big data into SQL Data Warehouse with simple <u>PolyBase</u> T-SQL queries, and then use the power of MPP to run high-performance analytics

SQL DW Fundamentals



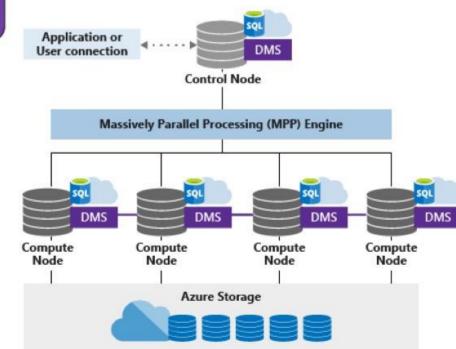
Control

Connection and tool endpoint. Coordinates compute activity.

Compute

Handles query processing, ability to scale up/down

DMS: Data Movement Services
Coordinates data movement
between nodes



Azure Data Lake Store

- Azure Data Lake Storage Gen2 is a set of capabilities dedicated to big data analytics, built on <u>Azure Blob</u> storage
 - Data Lake Storage Gen2 is the result of converging the capabilities of our two existing storage services, Azure Blob storage and Azure Data Lake Storage Gen1
 - Data Lake Storage Gen2 allows you to easily manage massive amounts of data
 - Data Lake Storage Gen2 allows you to manage and access data just as you would with a Hadoop Distributed File System (HDFS). The new ABFS driver is available within all Apache Hadoop environments including Azure HDInsight, Azure Databricks and SQL Data Warehouse

Azure Data Lake Analytics

- Azure Data Lake Analytics is an on-demand analytics job service that simplifies big data
- Instead of deploying, configuring, and tuning hardware, you write queries to transform your data and extract valuable insights
- The analytics service can handle jobs of any scale instantly by setting the dial for how much power you need. You only pay for your job when it is running, making it cost-effective