

# IT PRO | DEV Connections

12-13.12.2020 | on your screen

IIIIII00100



## Sponsors

PLATINUM



GOLD



SILVER



SUPPORTERS



CONTRIBUTORS

COMMUNITIES



MEDIA SPONSORS



# Azure AI on-premises by using Docker!

#itprodevgreece



Georgia Kalyva  
Microsoft AI MVP, MCP, MSP  
Web Applications Developer @ ITT Inc.

# About me

IT PRO | DEV connections 



Microsoft MVP in AI

Web Applications Developer @ ITT Inc.

Microsoft Learn Ambassador

Blogger

Traveler

Speaker



[www.codestories.gr](http://www.codestories.gr)



[fb.com/CodeStoriesgr](https://fb.com/CodeStoriesgr)



[georgiakalyva](https://in.linkedin.com/in/georgiakalyva)



[GeorgiaKalyva](https://github.com/GeorgiaKalyva)




[@GeorgiaKalyva](https://twitter.com/GeorgiaKalyva)

# Agenda

- ✓ Microsoft Cognitive Services
- ✓ Docker
- ✓ Azure Cognitive Services containers
- ✓ Live Demo
- ✓ Q&A


# Definition

- Artificial Intelligence is technology that appears to emulate human performance typically by learning, coming to it's own conclusions, appearing to understand complex content, engaging in natural dialogs with people, enhancing human cognitive performance or replacing or supporting people on execution of tasks.



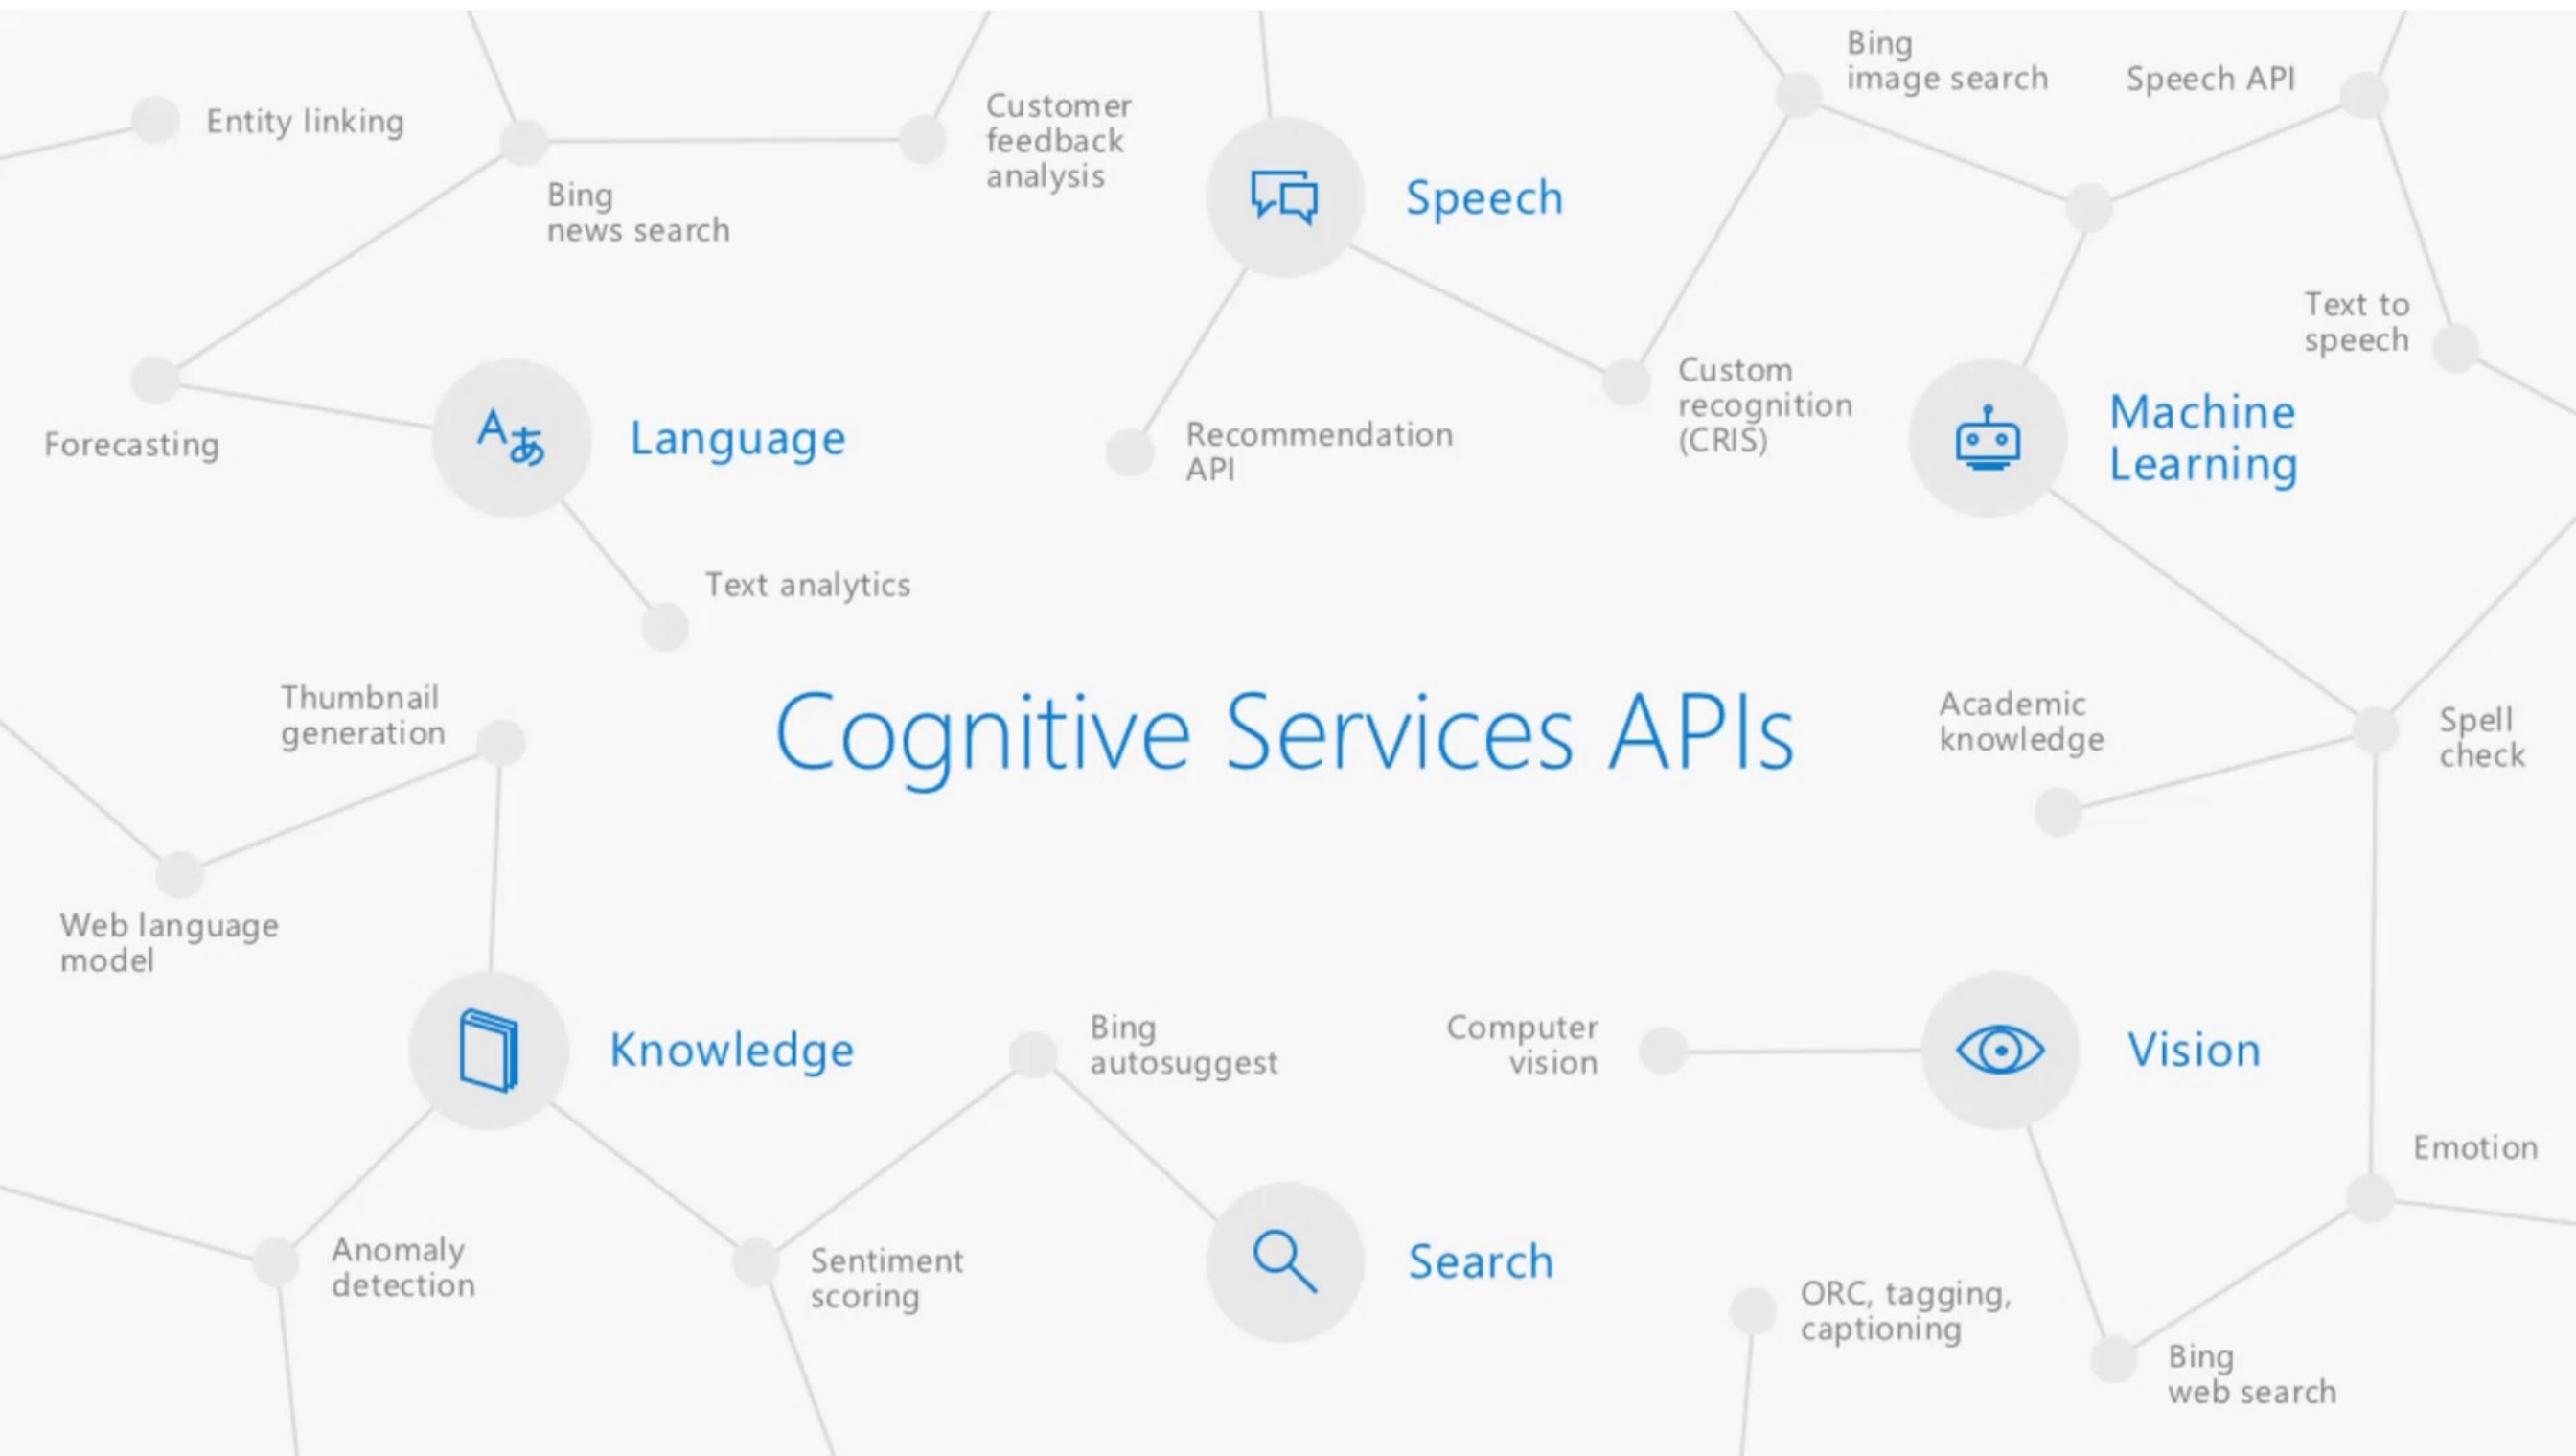
# Microsoft Cognitive Services

Give your apps a human side




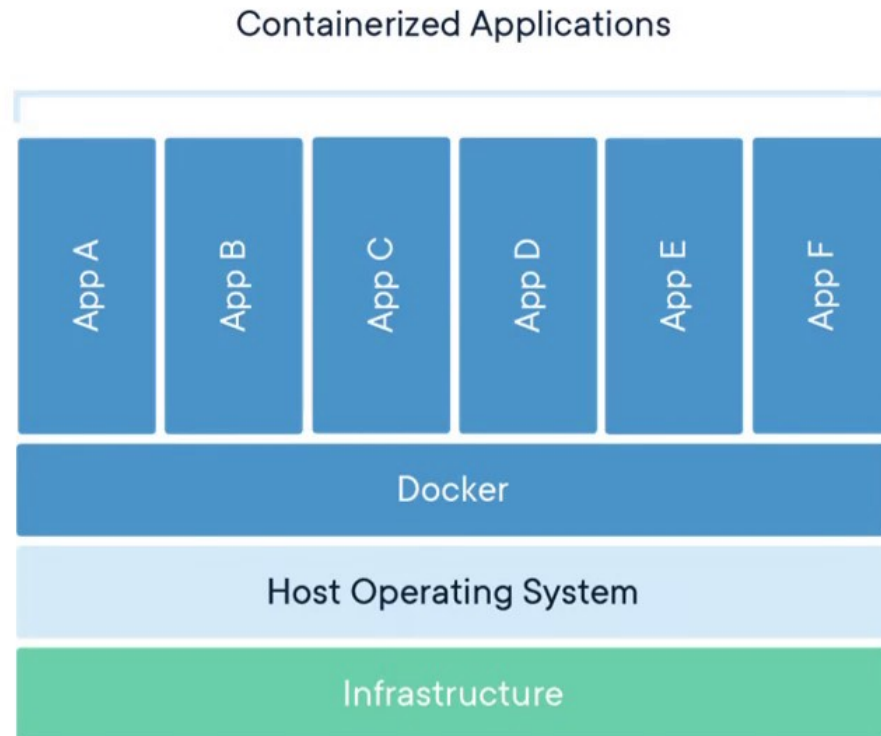
A set of ready to use APIs for every developer to use without requiring machine-learning expertise

# Cognitive Services APIs

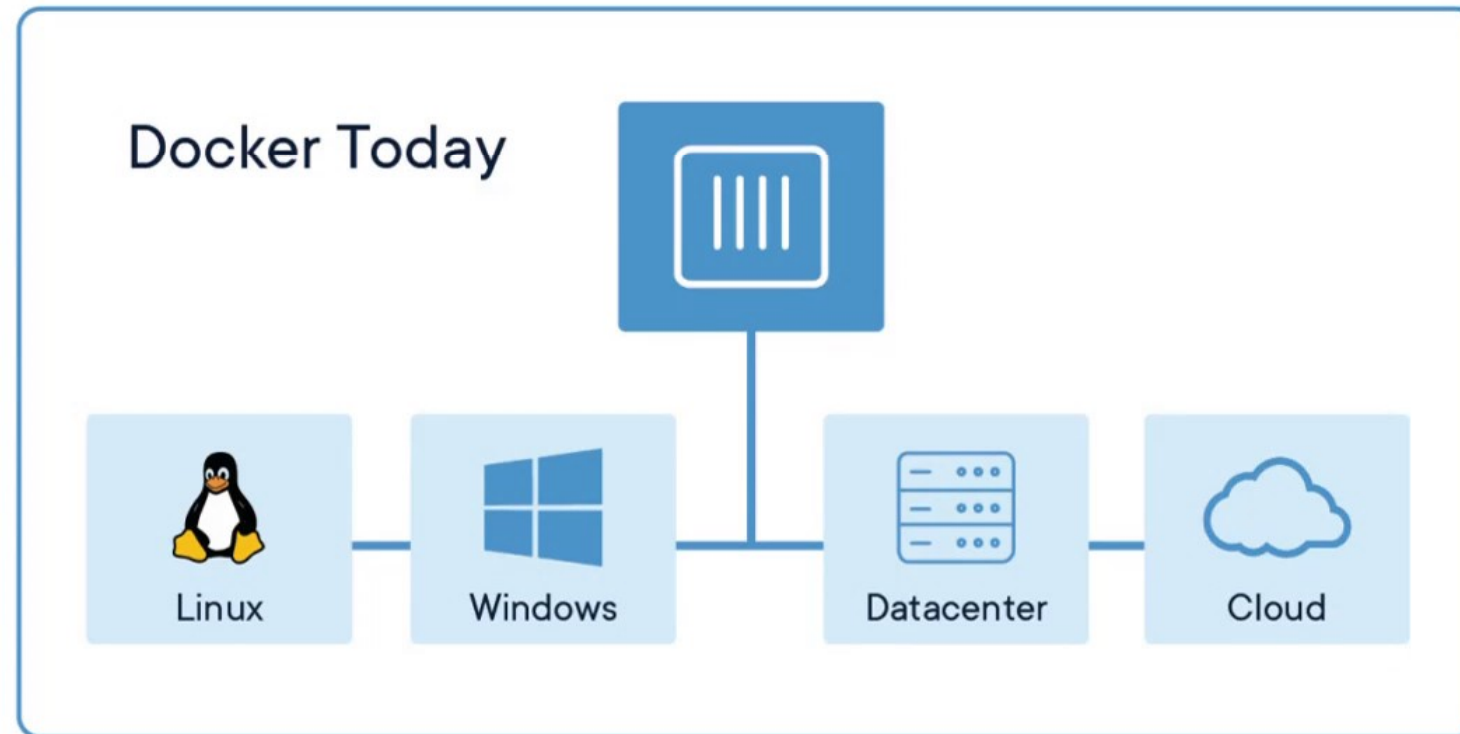




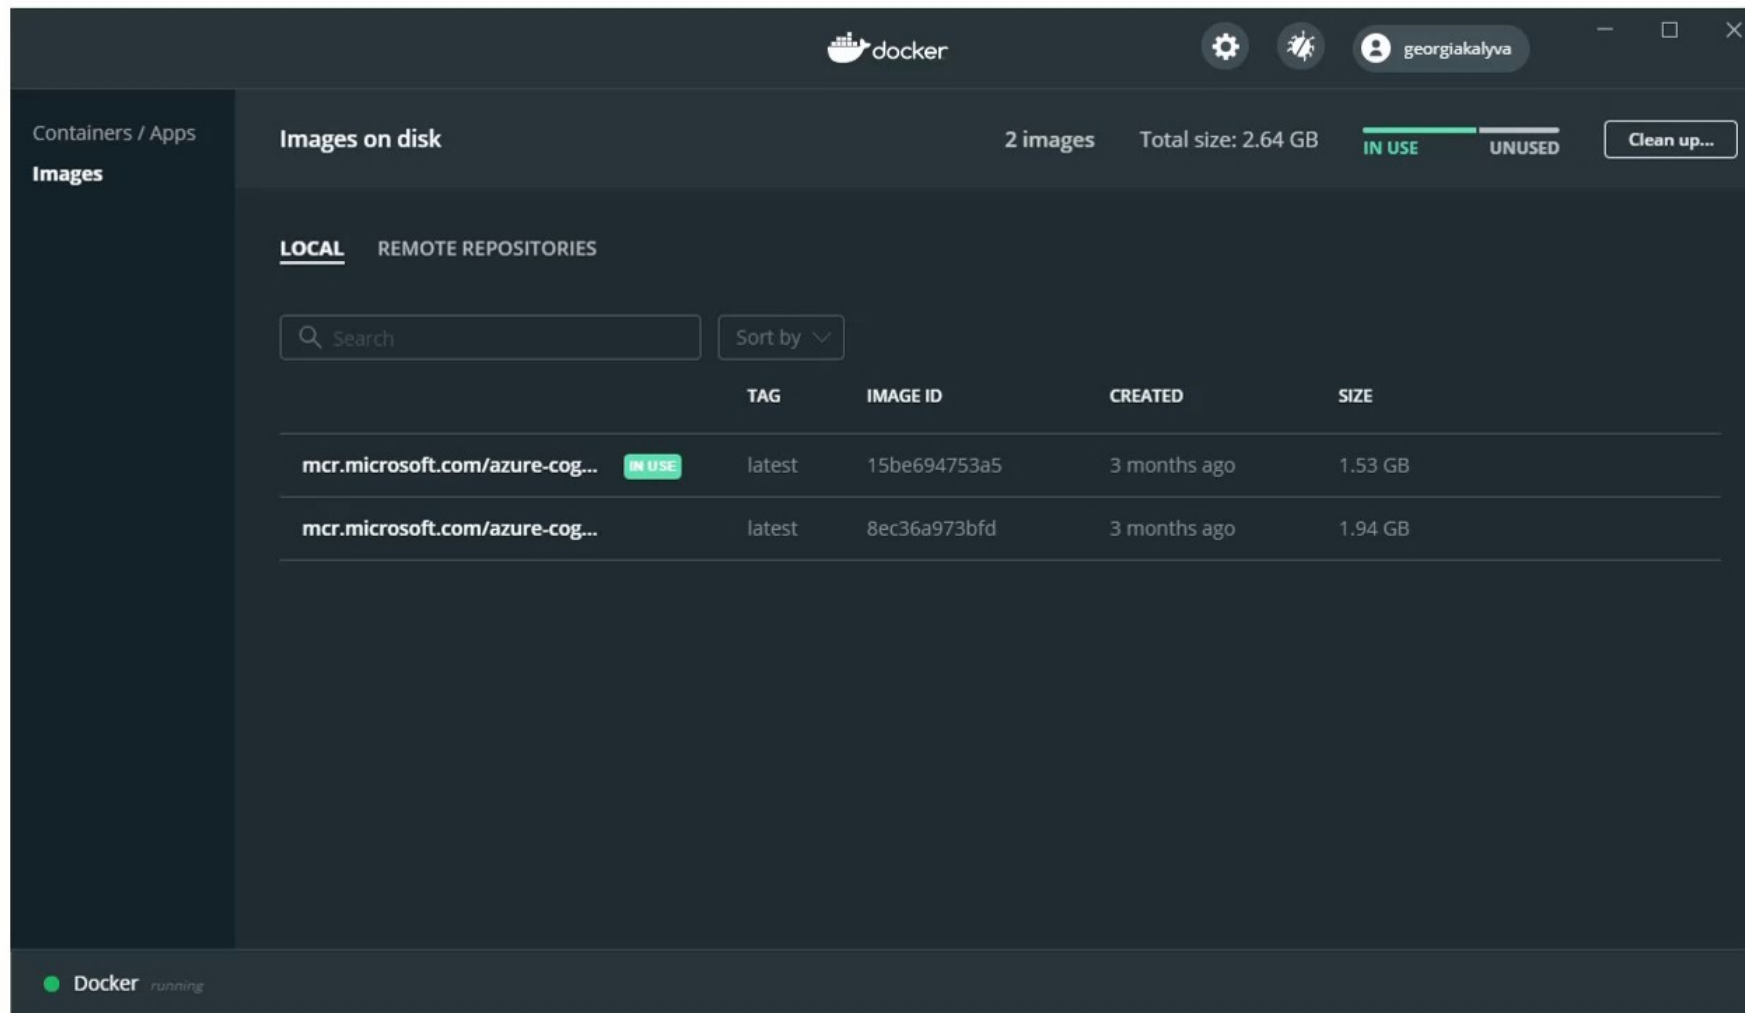
# Docker Containerized Applications



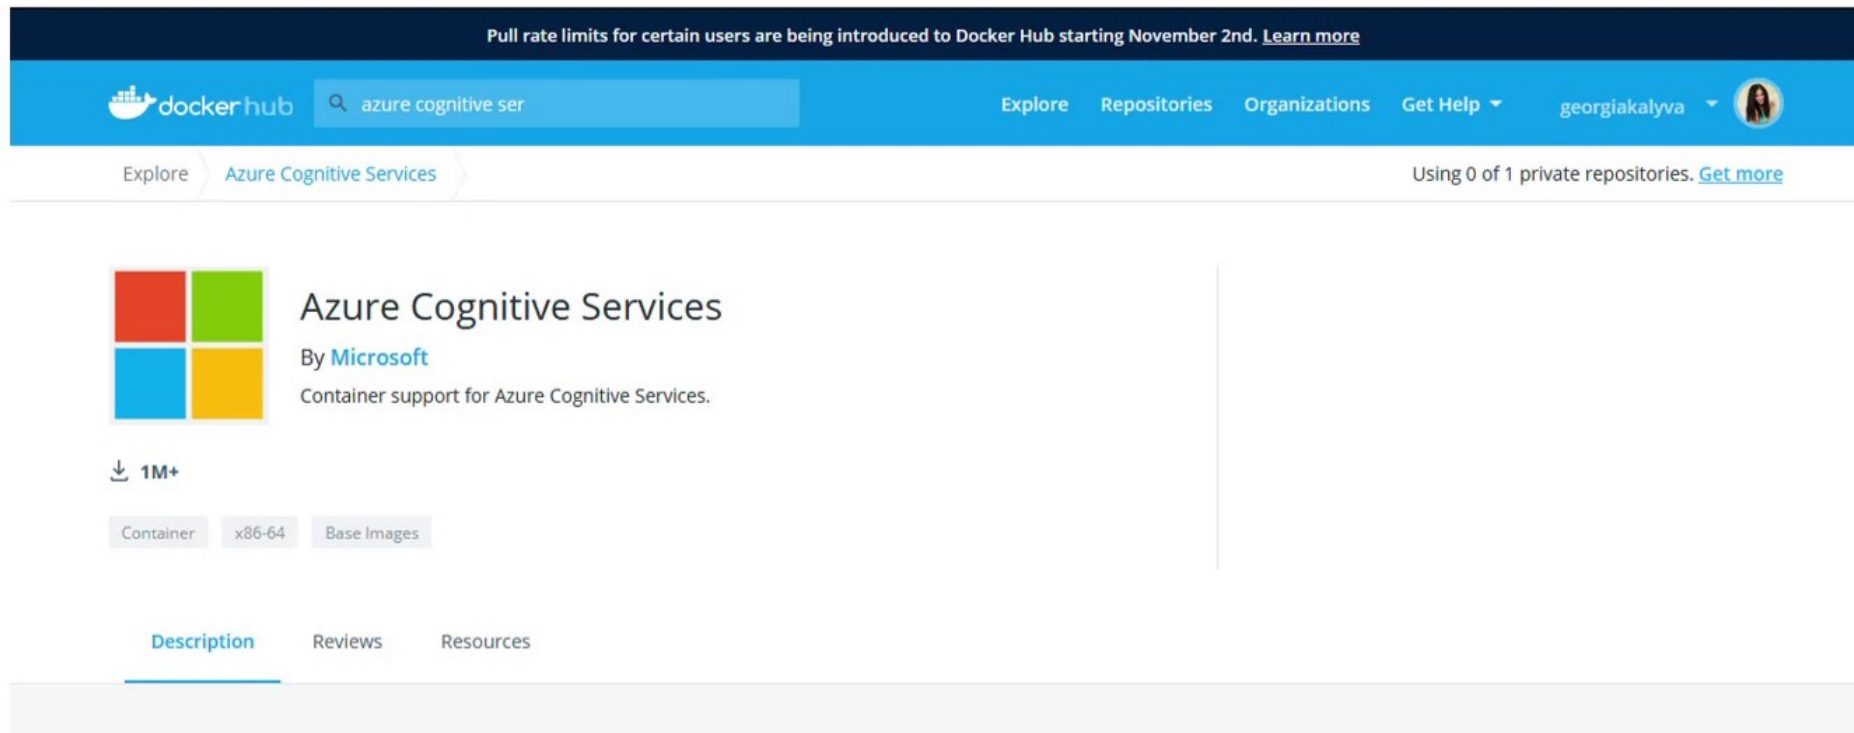
# Docker Containerized Applications



# Docker Desktop



# Docker Hub



The screenshot shows the Docker Hub interface. At the top, a dark blue banner contains the text: "Pull rate limits for certain users are being introduced to Docker Hub starting November 2nd. [Learn more](#)". Below this is a blue navigation bar with the Docker Hub logo, a search bar containing "azure cognitive ser", and navigation links for "Explore", "Repositories", "Organizations", and "Get Help". A user profile for "georgiakalyva" is visible on the right. Below the navigation bar, a breadcrumb trail shows "Explore" > "Azure Cognitive Services". On the right side of the navigation bar, it says "Using 0 of 1 private repositories. [Get more](#)".

The main content area displays the search results for "Azure Cognitive Services". It features a logo consisting of four colored squares (red, green, blue, yellow) and the text "Azure Cognitive Services" followed by "By Microsoft" and "Container support for Azure Cognitive Services." Below the logo, it indicates "1M+" downloads. There are three tabs: "Container", "x86-64", and "Base Images". At the bottom, there are three tabs: "Description", "Reviews", and "Resources".

Docker Hub is a Software-as-a-Service (SaaS) Docker container registry.



Azure Cognitive Services containers

# Azure Cognitive Services containers

- Anomaly Detector
- Read OCR (Optical Character Recognition)
- Spatial analysis
- Face
- Form Recognizer
- Language Understanding (LUIS)
- Speech Service API
- Text Analytics



# Benefits

- Immutable infrastructure
- Control over data
- Control over model updates
- Portable architecture
- High throughput / low latency
- Scalability



# Prerequisites



Active Azure  
Subscription with  
Cognitive Services  
Keys



Docker Engine  
configured to  
connect to Azure



Basic knowledge  
with Microsoft  
Container Registry  
and Docker



# Getting the images

- Use docker pull to get the container image

```
docker pull mcr.microsoft.com/azure-cognitive-services/textanalytics/language:latest
```

- Use docker run to run the container

```
docker run --rm -it -p 5000:5000 --memory 4g --cpus 1 \  
mcr.microsoft.com/azure-cognitive-services/textanalytics/language \  
Eula=accept \  
Billing={ENDPOINT_URI} \  
ApiKey={API_KEY}
```


# Mandatory Configuration Parameters

- API Key
- Billing Endpoint
- EULA Agreement

```
docker run --rm -it -p 5000:5000 --memory 4g --cpus 1 \  
mcr.microsoft.com/azure-cognitive-services/textanalytics/language \  
Eula=accept \  
Billing={ENDPOINT_URI} \  
ApiKey={API_KEY}
```


# Other Configuration Parameters

- ApiKey\*
- Billing\*
- Eula\*
- ApplicationInsights
- Fluentd
- HTTP Proxy
- Logging
- Mounts



# Other Configuration Parameters

- ApiKey\*
- Billing\*
- Eula\*
- ApplicationInsights
- Fluentd
- HTTP Proxy
- Logging
- Mounts



# Container networking

- Allow list port 443
- Access to domain \*.cognitive.microsoft.com
- Access to domain \*.cognitiveservices.azure.com



# Container requirements

	Minimum host specs	Recommended host specs	Minimum TPS	Maximum TPS
Language detection, key phrase extraction	1 core, 2GB memory	1 core, 4GB memory	15	30
Sentiment Analysis v3	1 core, 2GB memory	4 cores, 8GB memory	15	30
Text Analytics for health - 1 document/request	4 core, 10GB memory	6 core, 12GB memory	15	30
Text Analytics for health - 10 documents/request	6 core, 16GB memory	8 core, 20GB memory	15	30

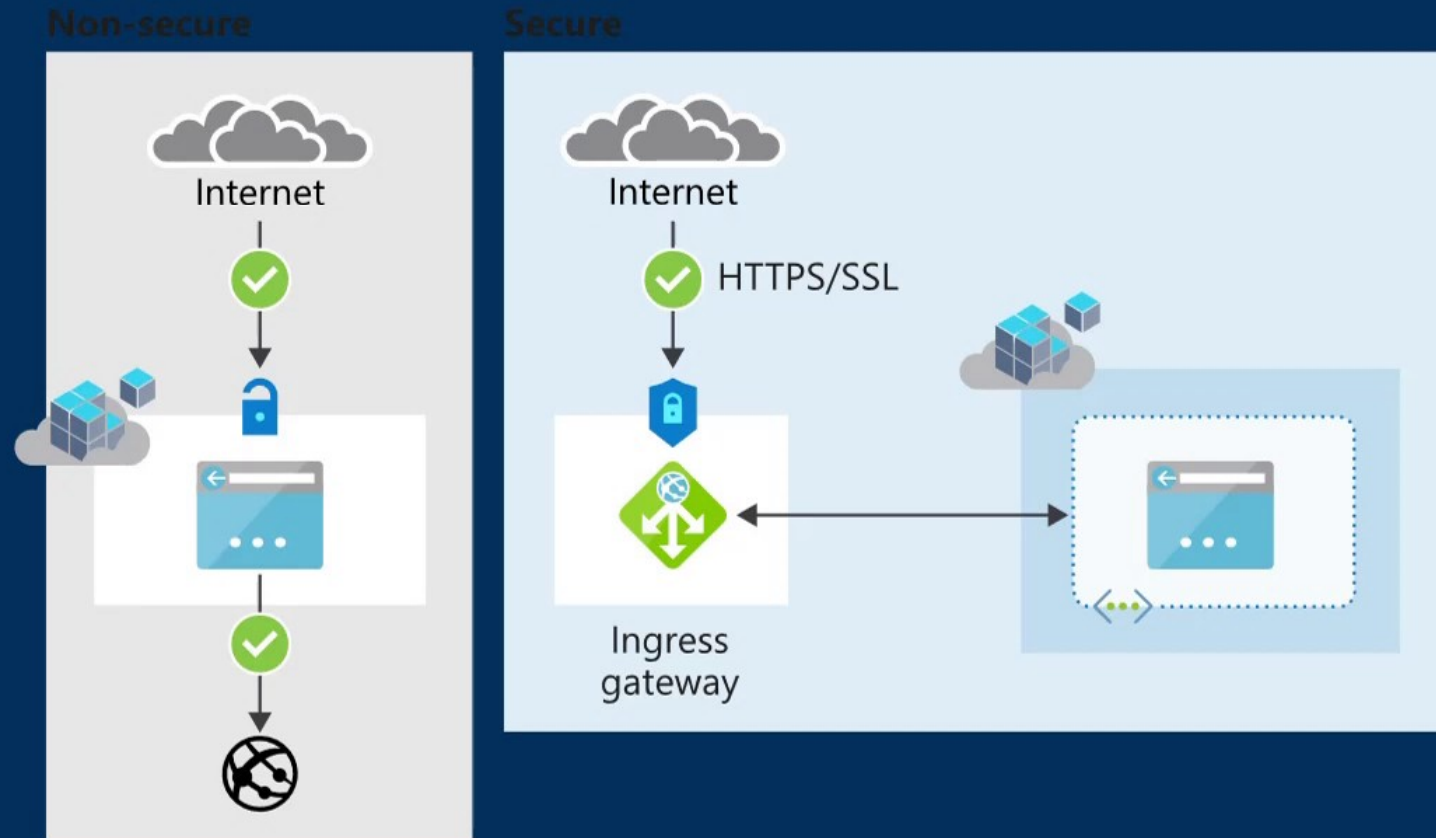
```
docker run --rm -it -p 5000:5000 --memory 4g --cpus 1 \  
mcr.microsoft.com/azure-cognitive-services/textanalytics/language \  
Eula=accept \  
Billing={ENDPOINT_URI} \  
ApiKey={API_KEY}
```

# Container configuration for reuse

- Create and build image from a dockerfile

```
FROM mcr.microsoft.com/azure-cognitive-services/language
ARG key
ARG billing_endpoint
ENV apikey=$key
ENV billing=$billing_endpoint
ENV eula=accept
```

# Security





# Container deployment options

- On-premise server with Docker engine
- Azure Container Instances
- Azure Kubernetes Service
- Publish to Azure Container Registry



IT PRO | DEV connections

IIIIII00100



Start  
LIVE DEMO



# Questions?



[www.codestories.gr](http://www.codestories.gr)



[fb.com/CodeStoriesgr](https://fb.com/CodeStoriesgr)



[georgiakalyva](https://www.linkedin.com/in/georgiakalyva)



[GeorgiaKalyva](https://github.com/GeorgiaKalyva)



[@GeorgiaKalyva](https://twitter.com/GeorgiaKalyva)