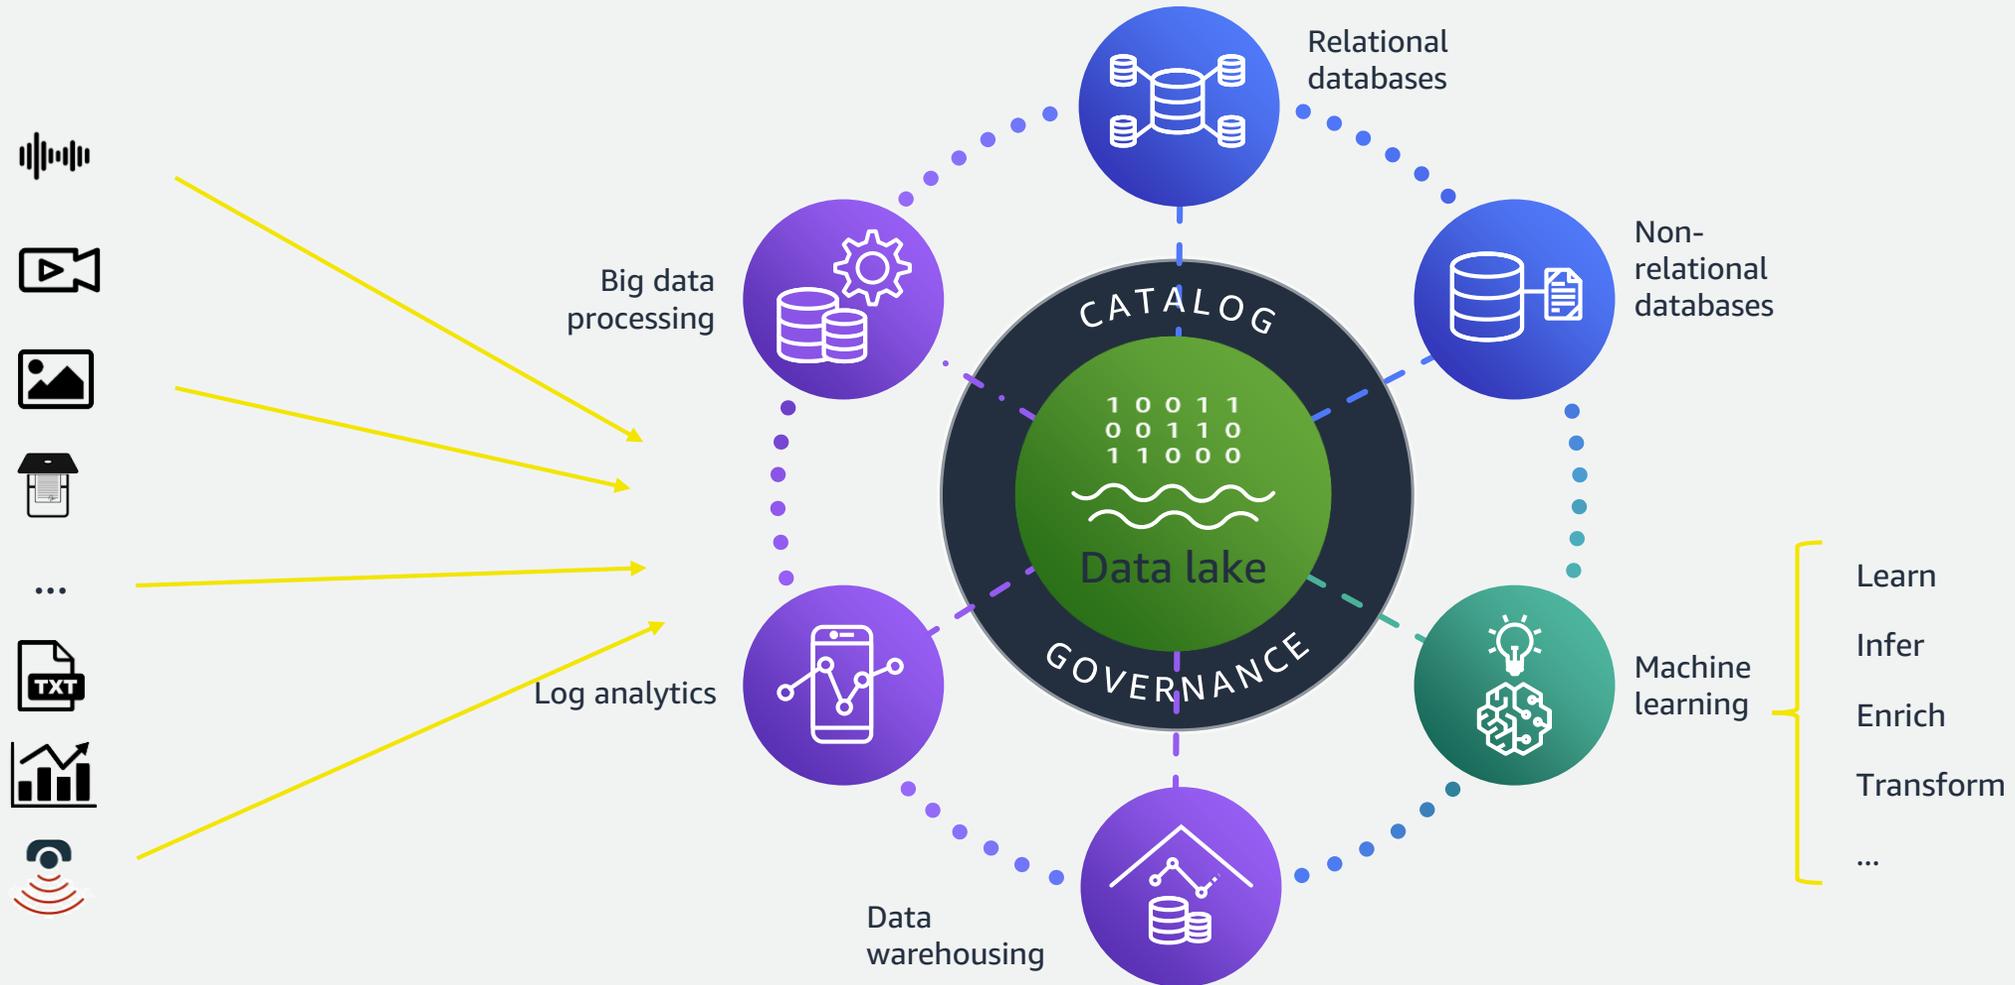# Come valorizzare i propri dati con l'Intelligenza artificiale e machine learning

Federico D'Alessio – dalesf@amazon.it

AWS WWPS – Senior Solutions Architect

Overview degli usecase e servizi di AI/ML di AWS

Data store and information extraction

Relational databases

Non-relational databases

Big data processing

CATALOG

Data lake

GOVERNANCE

Machine learning

Learn

Infer

Enrich

Transform

...

Log analytics

Data warehousing

Overview of AWS AI/ML services

AI services

Vision
Amazon Rekognition

Chatbots
Amazon Lex

Business tools
Amazon Personalize, Amazon Forecast, Amazon Fraud Detector, Amazon Lookout for Metrics

Search
Amazon Kendra

Healthcare
Amazon HealthLake, Amazon Comprehend Medical, Amazon Transcribe Medical

Speech
Amazon Polly
Amazon Transcribe

Text
Amazon Comprehend
Amazon Translate
Amazon Textract

Contact centers
Contact Lens for Amazon Connect
Amazon Connect Voice ID

Code & DevOps
Amazon CodeGuru
DevOps Guru

Industrial
AWS Panorama Appliance and SDK, Amazon Monitron, Amazon Lookout for Equipment, Amazon Lookout for Vision

Amazon SageMaker

SageMaker Studio IDE

Label data	Data collection prep	Store features	Detect bias and explain predictions	Visualize in notebooks	Pick algorithm	Train models faster	Tune parameters	Deploy in production	Manage & monitor	Manage edge devices

CI/CD

Ml frameworks & infrastructure

PyTorch, Apache MXNet, TensorFlow	AMIs & containers	GPUs	Inferentia	Elastic inference	FPGA

intel Gaudi

© 2022, Amazon Web Services, Inc. or its affiliates.

6

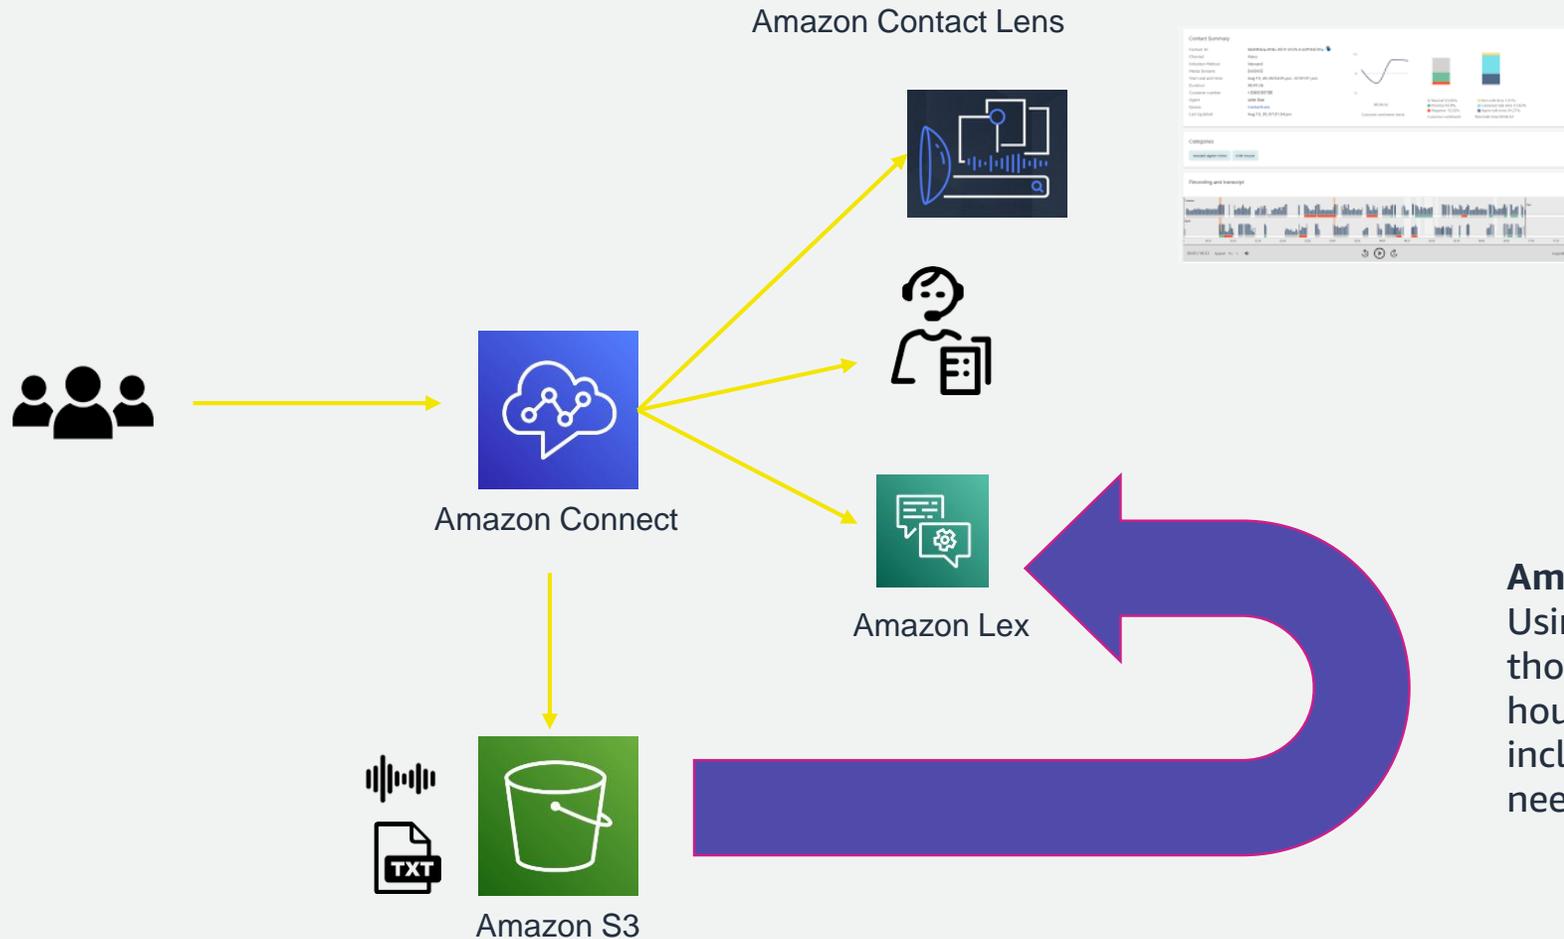## Now available: Gaudi-based Amazon EC2 DL1 Training Instances

"Best price performance for training deep learning models in the cloud" – AWS

- Featuring up to 8 Gaudi processors
- Up to 40% better price performance than latest GPU-based instances
- Habana SynapseAI® Software Platform integrates TensorFlow and PyTorch
- Launch DL1 instances via Amazon EC2 and EKS for containerized ML applications
- Get started easily using DLC, DL AMIs, or Amazon SageMaker (coming soon)

60% lower than EC2 P4d Instances (8 x A100 GPUs)

Instance	Gaudi	vCPU	Memory (GiB)	HBM	Networking	Local NVMe
DL1.24xlarge	8	96	768	32 GB per accelerator	400Gbps	4 TB

Contact Center – Chat/Voice Bots

Amazon Contact Lens

Amazon Connect

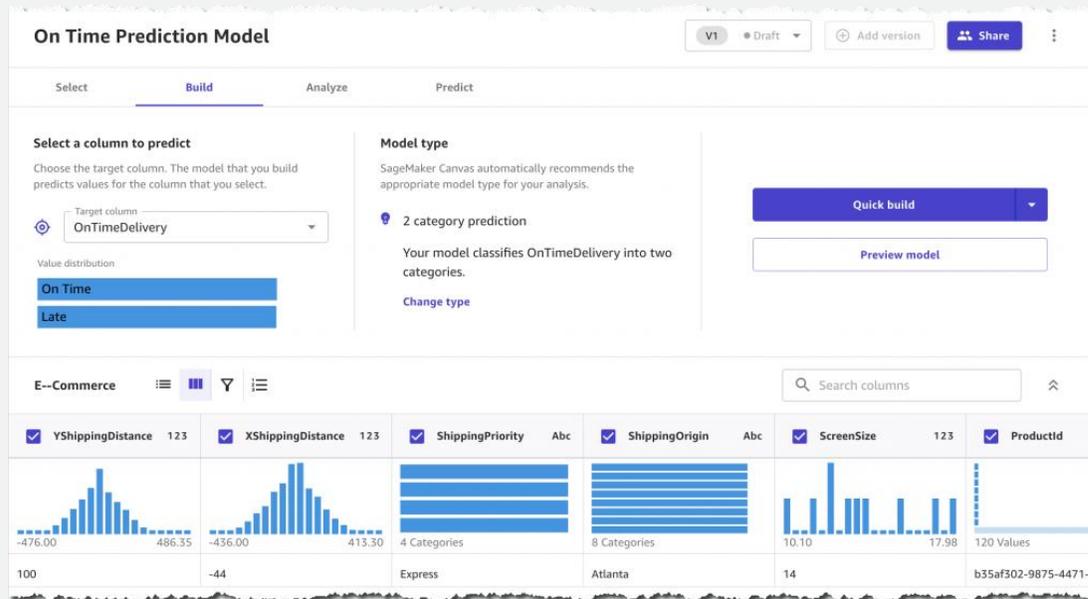Amazon Lex

Amazon S3

Sentiment Analysis

Automated Transcription

Automated document classification by using ML to identify PII data

Amazon Lex Automated Chatbot Designer
Using machine learning (ML), it can analyze thousands of lines of transcripts in a couple of hours and provide an initial bot design that includes common intents and the information needed to fulfill them.

Classification, Regression and Forecasting

Amazon SageMaker Canvas gives you the ability to use machine learning to generate predictions without needing to code.
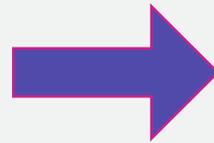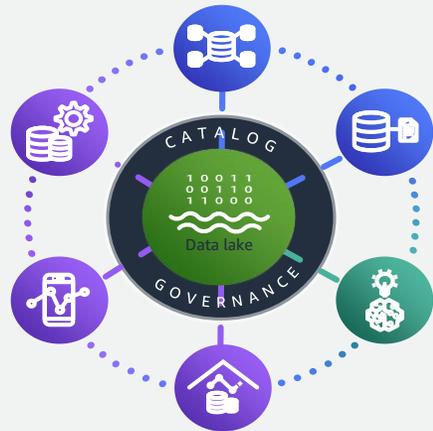

No Datascientist experience needed

Guided process

Output model can be transferred to datascientist for further improvement

churn, wharehouse optimization, detecting fraud, usage forecast, …

Machine learning on Edge devices

Amazon SageMaker

Train the model in cloud leveraging **fast CPU and accelerators**, leveraging **large data** available in the **Datalake.**

Amazon SageMaker Edge Manager

Optimize model for small devices

Orchestrate model deployment and updates

A2A SmartCity e Neosperience

Conclusioni e Q&A

Thank you!

Federico D'Alessio – dalesf@amazon.it

Senior Solutions Architect

Amazon Web Services - WWPS