



Advanced Analytics & AI on Azure

Nishant Thacker

Technical Product Manager
Analytics & AI



Advanced Analytics & AI represents a growing opportunity

Global business value derived from **AI in 2022** will reach



Adopting Advanced Analytics and AI in your company



Where do you see **your company today?**



Where do you see your **top competitors today?**

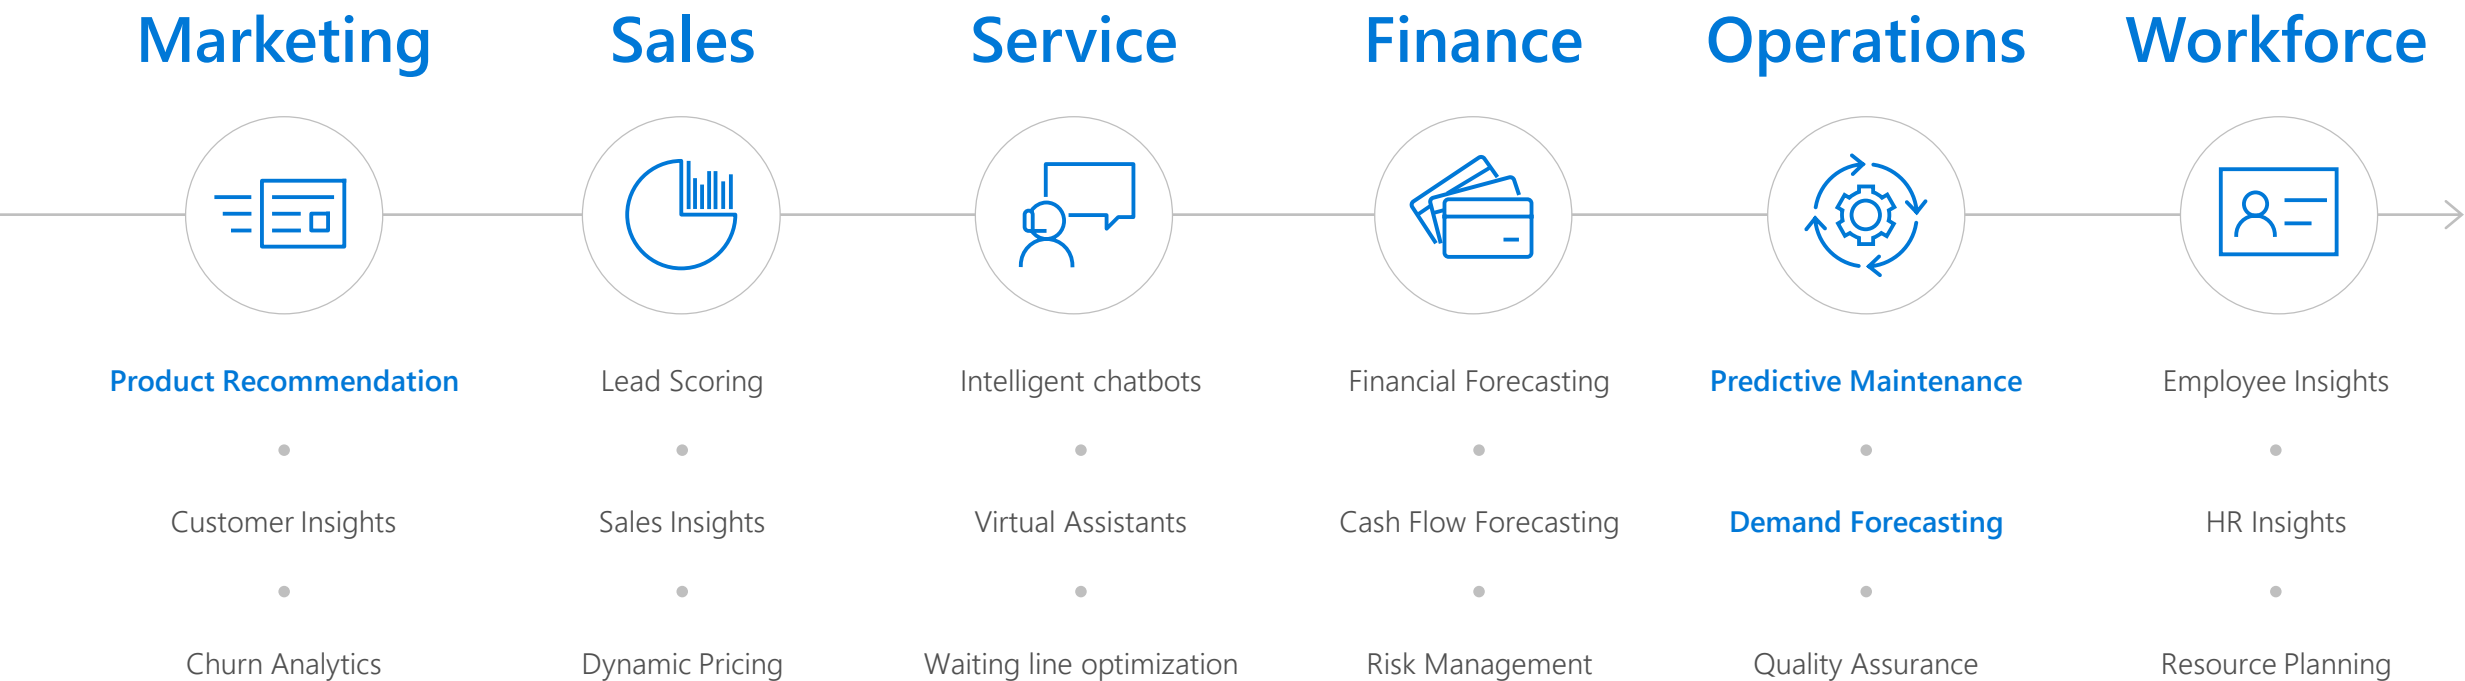


Where do you see **your company in the future?**

Analytics Capabilities

← BASIC ————— ADVANCED →

Helping you innovate across your business



Leading to transformational changes

Product recommendation



The average size of a single cart has decreased



Provide personalized digital content to shoppers



Increase cart size



ASOS
discover fashion online

ASOS delivers 15.4 million personalized experiences with 33 orders per second

Predictive maintenance



Unplanned downtime results in cost overruns



Predict when maintenance should be performed



Minimize downtime



CARNIVAL
MARITIME

Hybrid solution predicts onboard water usage, saving \$200k/ship/year

Demand forecasting



Solar energy production is inconsistent



Align energy supply with the optimal markets



Maximize revenue




renewablesAI
DATA SCIENCE & ASSET MANAGEMENT

Distributed power generation increases revenue by over €100 million

How companies are transforming



Serving business users and end users with **intelligent** and **dynamic** applications




Build a unified and usable
data pipeline

Train ML and DL models to
derive insights


Operationalize models and
distribute insights at scale

But there's a lot to consider




Complexity of solutions

Many options in the marketplace



Data silos

Incongruent data types



Difficult to scale effectively

Performance constraints

Microsoft has a comprehensive solution for Data & AI

Azure Data + AI Solution Areas

Data



Data Modernization on-premises



Data modernization to Azure



Globally distributed data



Cloud Scale Analytics

+

AI



AI apps & agents



Knowledge mining



Machine learning

Azure Data + AI Solution Areas

Data



Data Modernization on-premises



Data modernization to Azure



Globally distributed data



Cloud Scale Analytics

+

AI



AI apps & agents



Knowledge mining



Machine learning

Overview of Microsoft's Data Science & Machine Learning Portfolio

Machine Learning on Azure

Domain specific pretrained models

To reduce time to market



Vision



Speech



Language



Search

Familiar Data Science tools

To simplify model development



PyCharm



Jupyter



Visual Studio Code



Command line

Popular frameworks

To build advanced deep learning solutions



Pytorch



TensorFlow



Scikit-Learn



Onnx

Productive services

To empower data science and development teams



Azure
Databricks



Azure Machine
Learning



Machine
Learning VMs

Powerful infrastructure

To accelerate deep learning



CPU



GPU



FPGA



From the Intelligent Cloud to the Intelligent Edge



Sophisticated pretrained models

Infuse apps with powerful, pre-trained AI models

Customize easily and tailor to your needs



Vision



Computer Vision | Video Indexer | Face | Content Moderator



Language

I had a wonderful trip in Seattle, I enjoyed the Space Needle and Pike Place Market.


Analyze

Analyzed text		JSON
1 LANGUAGES:	English (confidence: 100 %)	
1 KEY PHRASES:		
1 SENTIMENT:	73 %	
1 LINKED ENTITIES (PREVIEW):	a	

Text Analytics | Spell Check | Language Understanding | Text Translation | QnA Maker




Bing Search



Big Web Search | Video Search | Image Search | Visual Search | Entity Search | News Search | Autosuggest



Speech



Speech to Text | Text to Speech | Speech Translation | Speaker Recognition

Productive services

Empower data science and development teams



Integrated data science & data engineering teams

Desktop solutions not adequate

Need a unified big data & machine learning solution



Azure Databricks

+



Azure Machine Learning
service



Individual data scientists

Desktop solutions adequate

Need cloud for sporadic compute needs



Machine Learning VMs



Azure Databricks

Fast, easy, and collaborative Apache Spark™-based analytics platform




Increase productivity



Build on a secure, trusted cloud



Scale without limits



Built with your needs in mind

- Optimized Apache Spark environment
- Collaborative workspace
- Integration with Azure data services
- Autoscale and autoterminate
- Optimized for distributed processing
- Support for multiple languages and libraries

Seamlessly integrated with the Azure Portfolio




Azure Machine Learning service

Bring AI to everyone with an end-to-end, scalable, trusted platform



Boost your data science productivity



Increase your rate of experimentation



Deploy and manage your models everywhere



Built with your needs in mind

- Automated machine learning
- Managed compute
- DevOps for machine learning
- Simple deployment
- Tool agnostic Python SDK
- Support for open source frameworks

Seamlessly integrated with the Azure Portfolio

Powerful infrastructure


Accelerate deep learning



CPUs

General purpose machine learning


D, F, L, M, H Series



GPUs

Deep learning

N Series



FPGAs

Specialized hardware accelerated deep learning

Project Brainwave

Optimized for flexibility

Optimized for performance


Flexible deployment

From the Intelligent Cloud to the Intelligent Edge

Train and deploy



Train and deploy



Track models in production
Capture model telemetry
Retrain models



Deploy

Machine Learning on Azure

Quickly and easily build, train, deploy and manage your models

Domain Specific Pretrained Models

To reduce time to market



Vision



Speech



Language



Search

Familiar Data Science tools

To simplify model development



PyCharm



Jupyter



Visual Studio Code



Command line

Popular Frameworks

To build machine learning and deep learning solutions



ONNX



Pytorch



TensorFlow



Scikit-Learn

Productive Services

To empower data science and development teams



Azure
Databricks



Azure Machine
Learning



Machine
Learning VMs

Powerful Infrastructure

To accelerate deep learning



CPU



GPU



FPGA

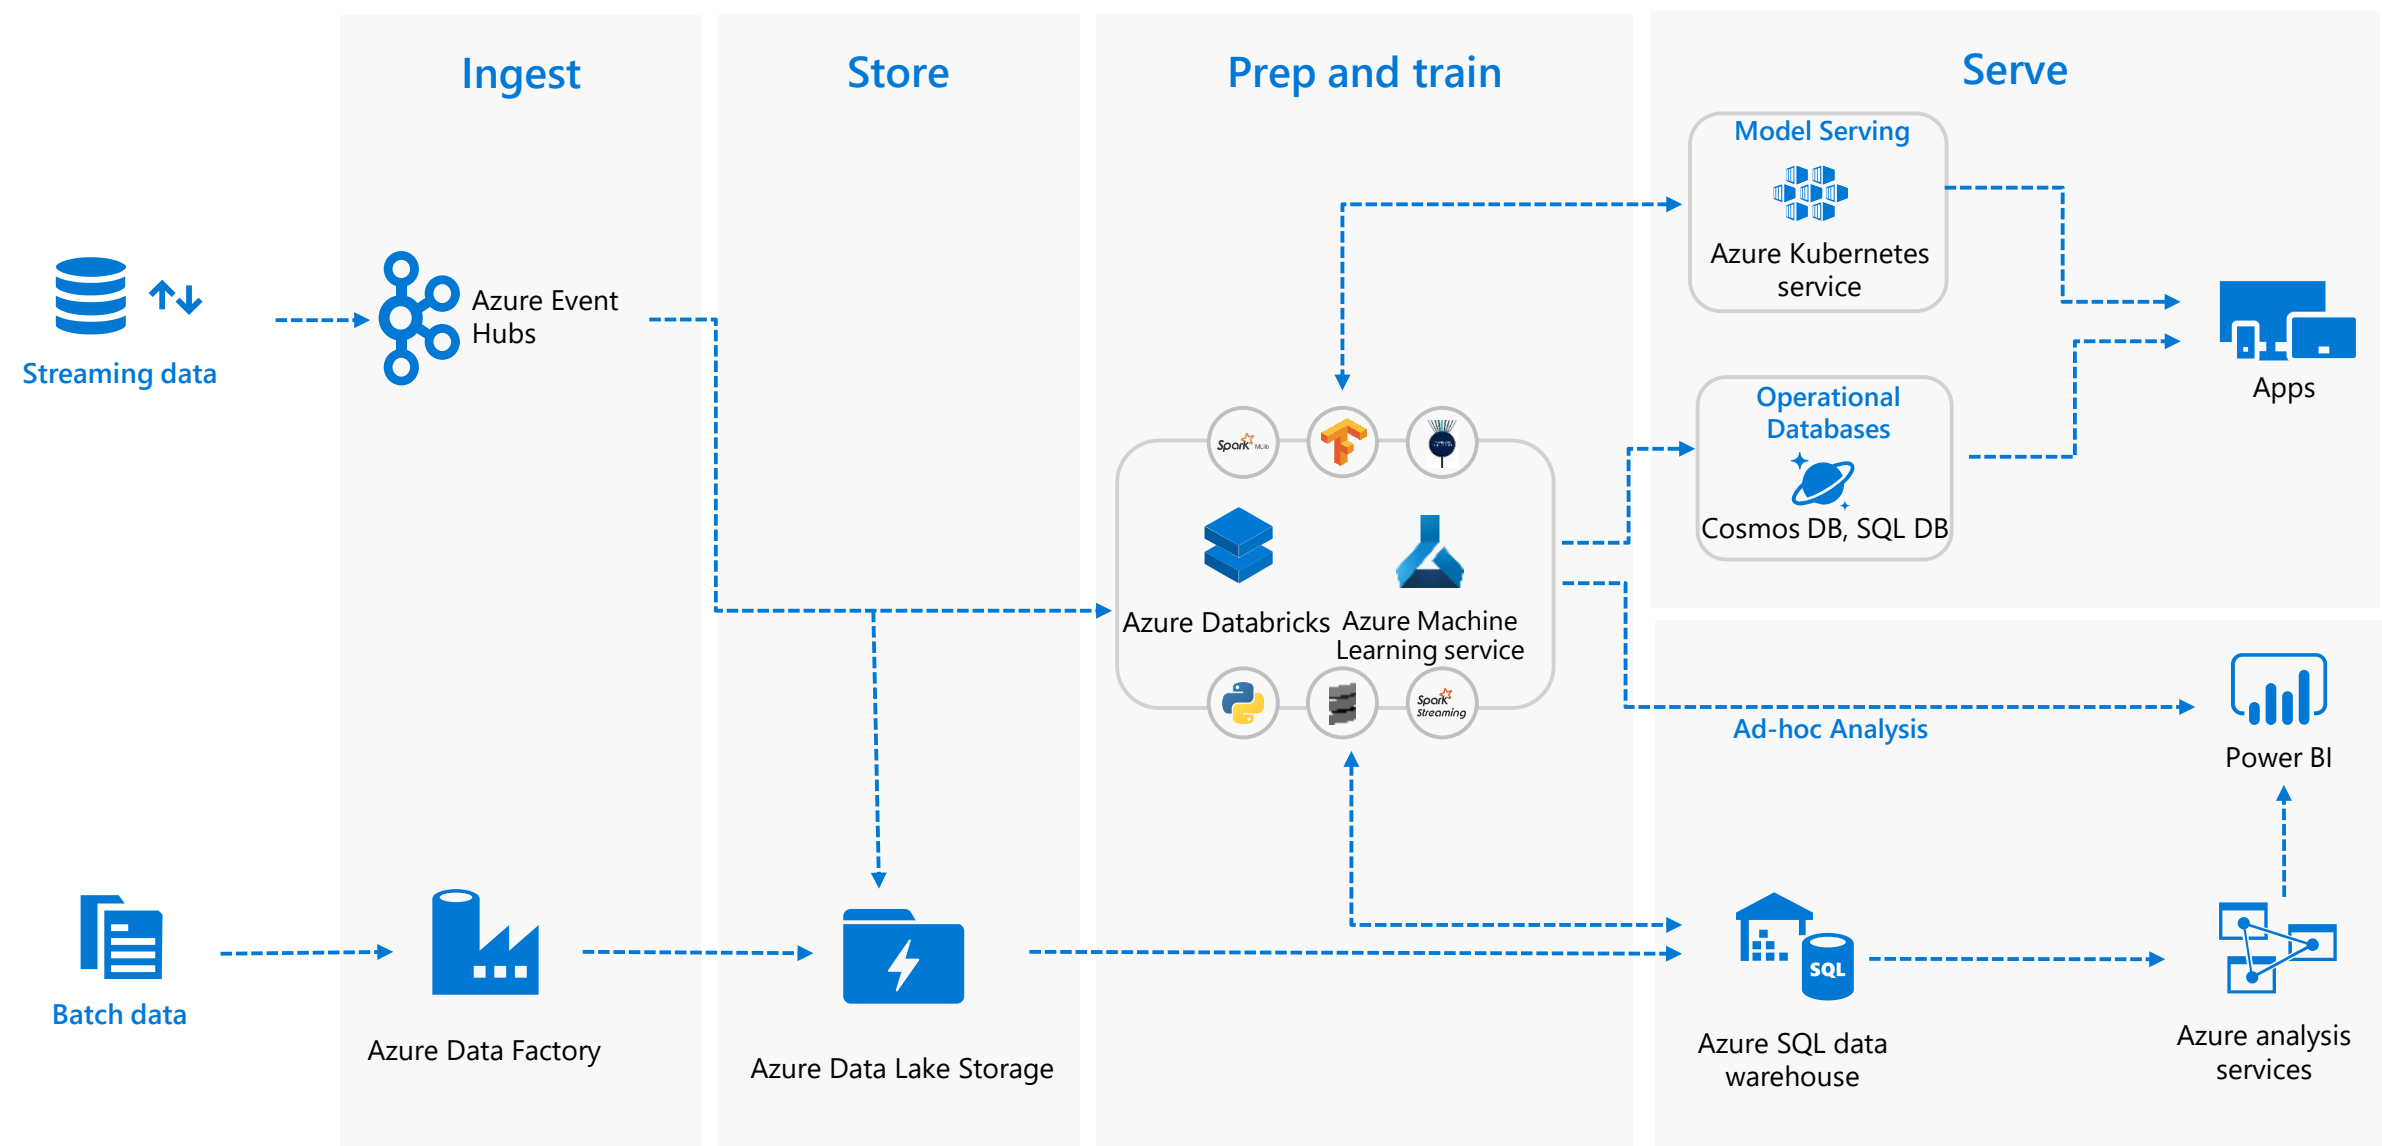


From the Intelligent Cloud to the Intelligent Edge




Understanding the Data Science Process on Azure


Recommended architecture to build e2e ML solutions




PREP & TRAIN



Collect and prepare data

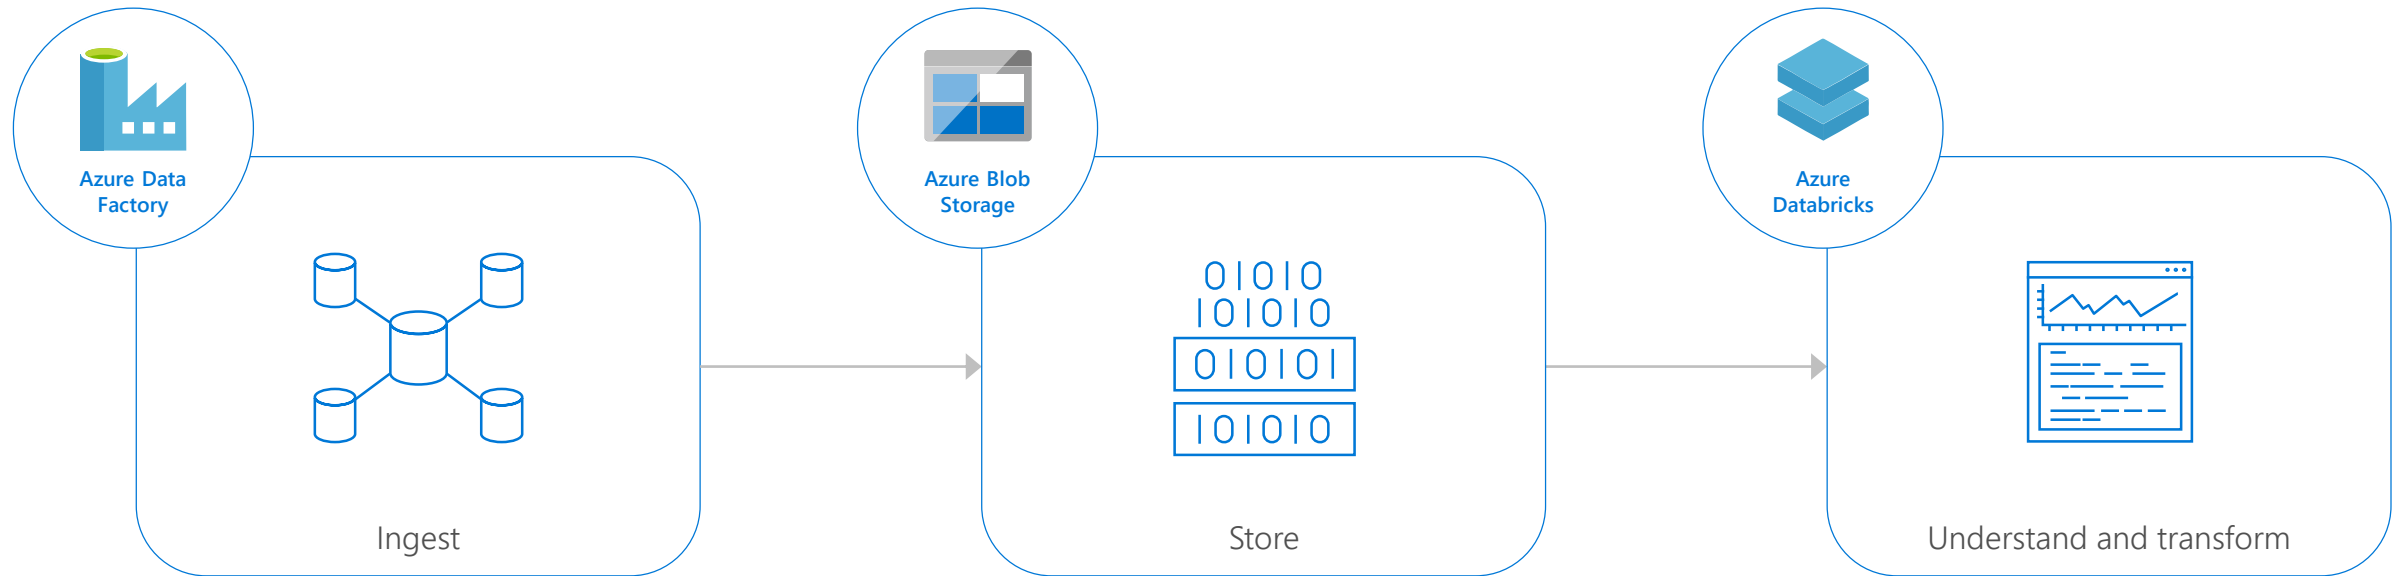


Train and evaluate model




Operationalize and manage

Collect and prepare all of your data at scale




Connect to data from any source

- Integrate with all of your data sources
- Create hybrid pipelines
- Orchestrate in a code-free environment




Leverage best-in-class analytics capabilities

- Leverage open source technologies
- Collaborate within teams
- Use ML on batch streams

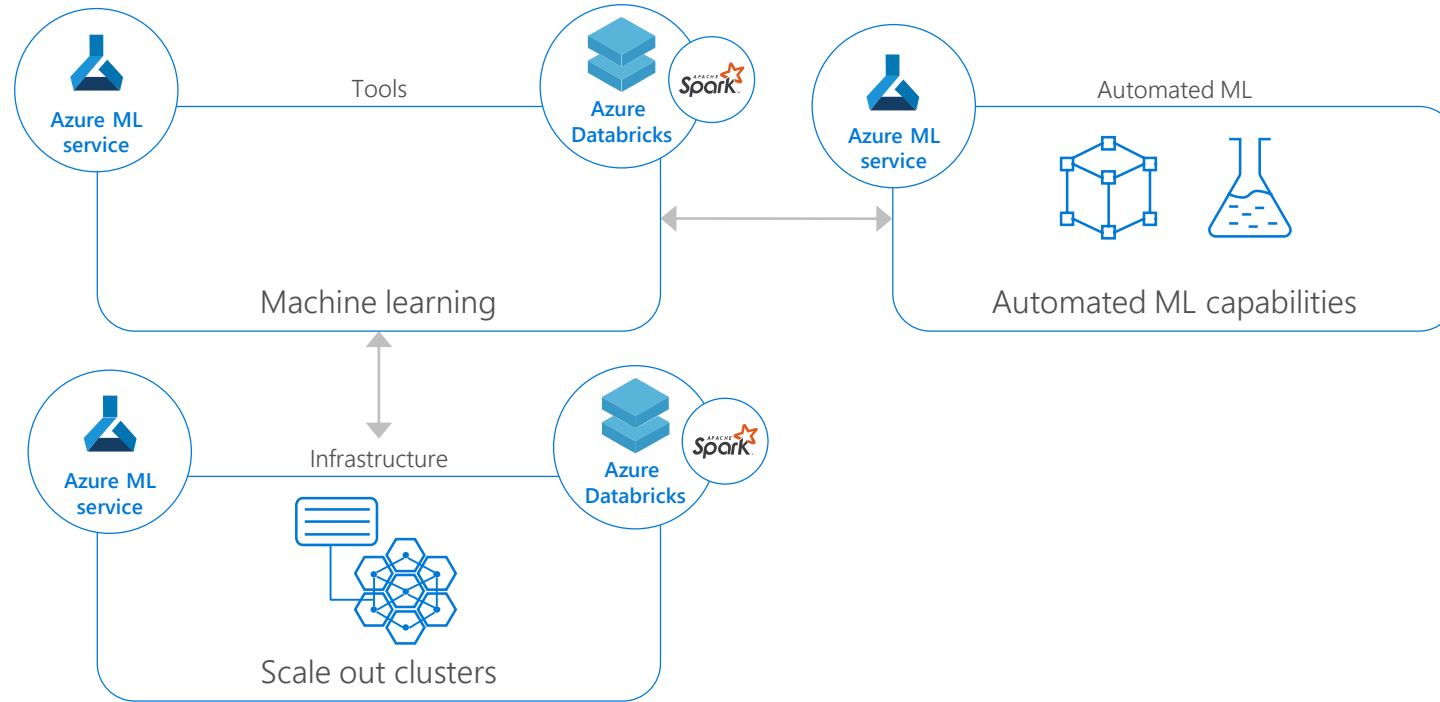


Scale without limits

- Build in the language of your choice
- Leverage scale out topology
- Scale compute and storage separately




Train and evaluate machine learning models



Simplify model development

- Collaborate in interactive workspaces
- Access a library of battle-tested models
- Automate job execution



Scale compute resources to meet your needs

- Easily scale up or scale out
- Autoscale on serverless infrastructure
- Leverage commodity hardware

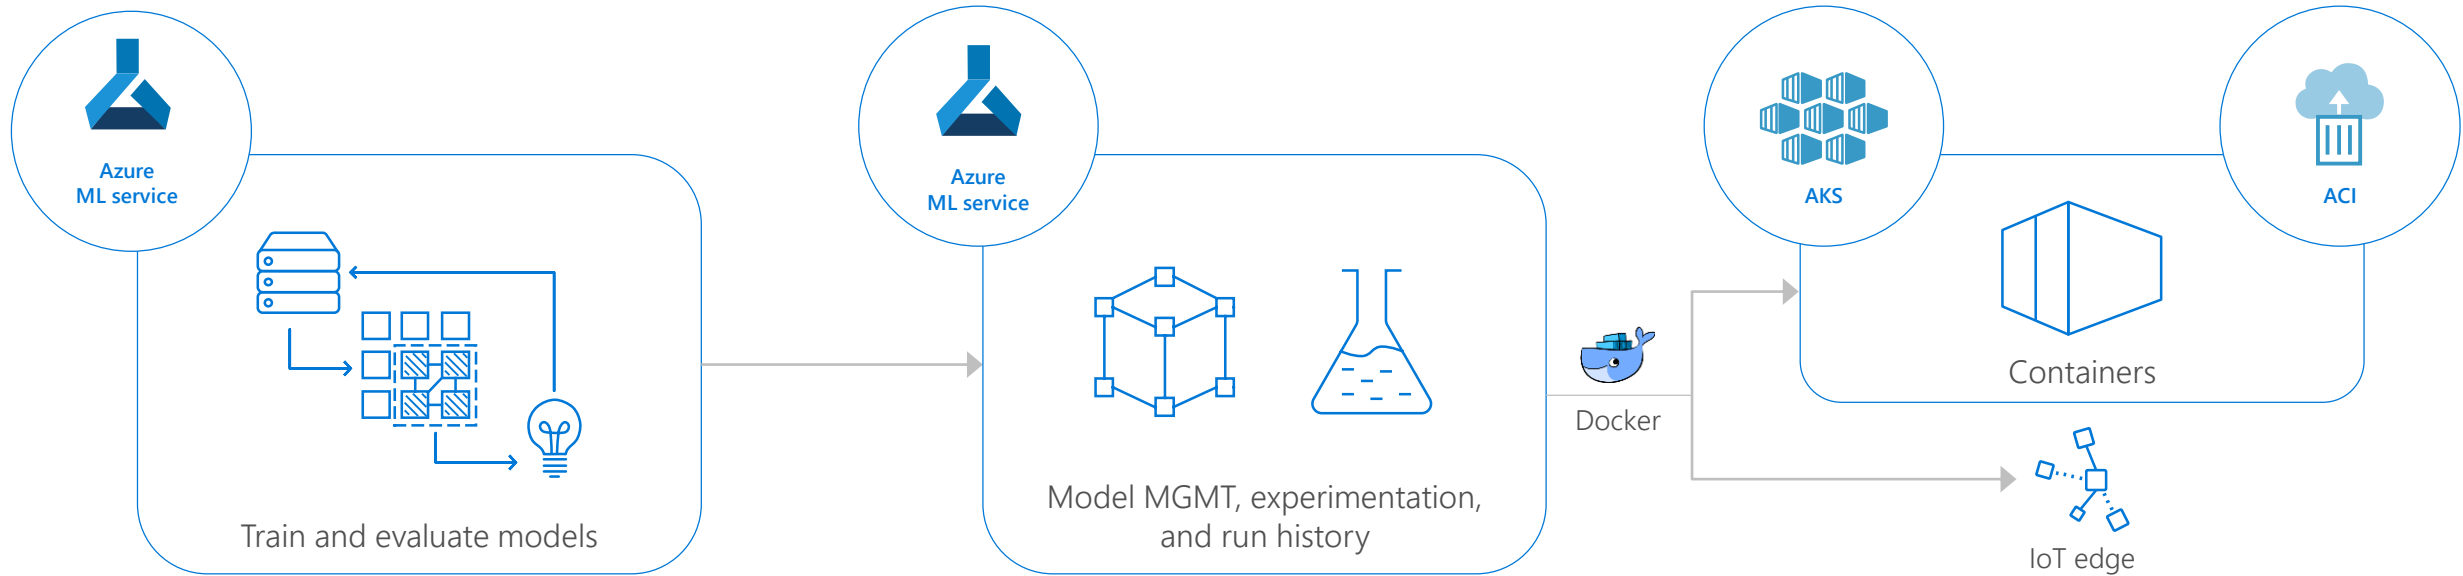


Quickly determine the right model for your data

- Determine the best algorithm
- Tune hyperparameters to optimize models
- Rapidly prototype in agile environments




Operationalize and manage models with ease



Bring models to life quickly

- Build and deploy models in minutes
- Iterate quickly on serverless infrastructure
- Easily change environments




Proactively manage model performance

- Identify and promote your best models
- Capture model telemetry
- Retrain models with APIs

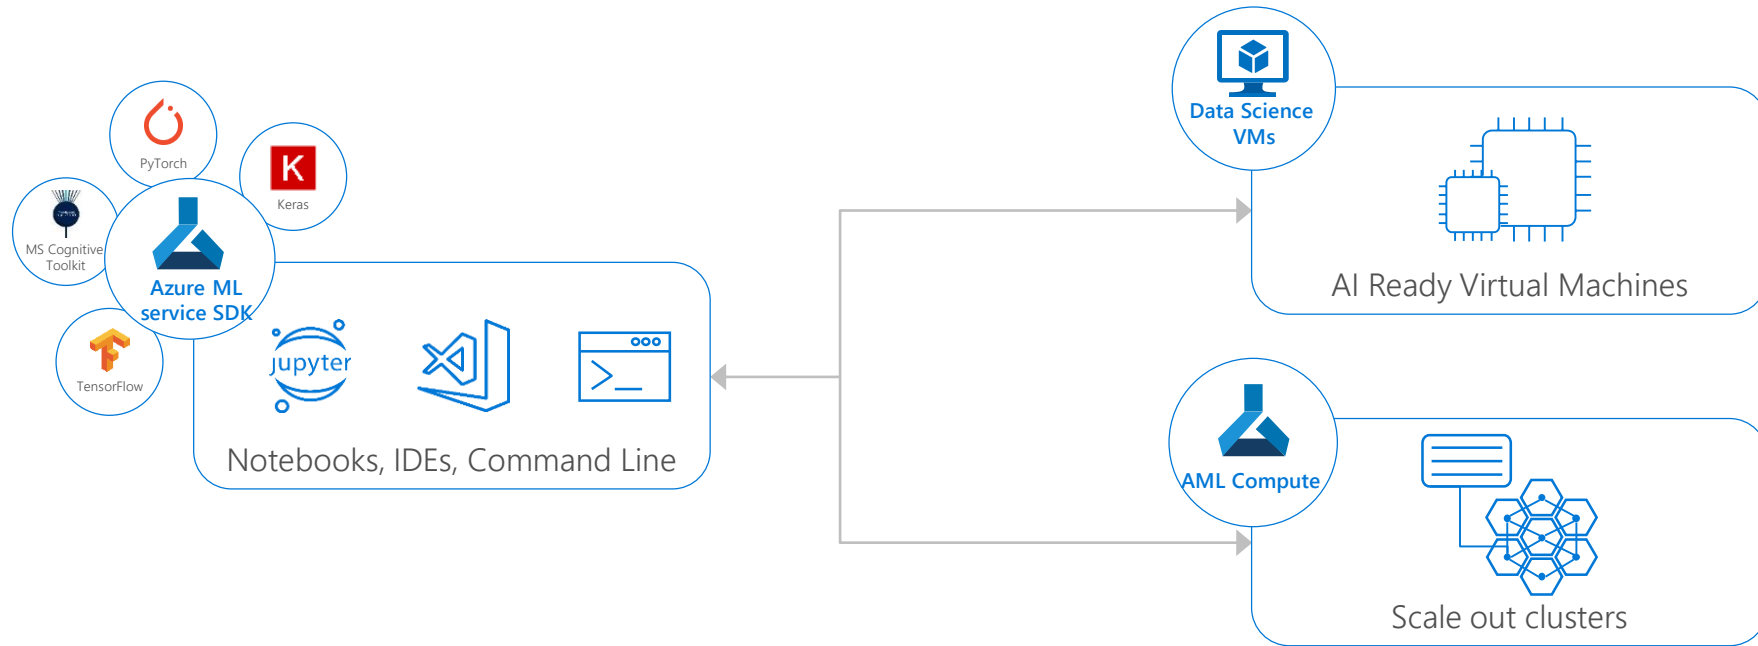


Deploy models closer to your data

- Deploy models anywhere
- Scale out to containers
- Infuse intelligence into the IoT edge



Build and deploy deep learning models




Streamline AI development efforts

- Leverage popular deep learning toolkits
- Develop your language of choice



Scale compute resources in any environment

- Choose VMs for your modeling needs
- Process video using GPU-based VMs



Quickly evaluate and identify the right model

- Run experiments in parallel
- Provision resources automatically



Azure Databricks for deep learning model

Ready-to-use clusters with Azure Databricks Runtime for ML



Tools

Use HorovodEstimator via a native runtime to enable build deep learning models with a few lines of code

Load images natively in Spark DataFrames to automatically decode them for manipulation at scale with distributed DNN training on Spark

Simultaneously collaborate within notebooks environments to streamline model development



Frameworks

Full Python and Scala support for transfer learning on images

Seamlessly use TensorFlow, Microsoft Cognitive Toolkit, Caffe2, Keras, and more

Use built-in hyperparameter tuning via Spark MLLib to quickly optimize the model



Infrastructure

Leverage powerful GPU-enabled VMs pre-configured for deep neural network training

Automatically store metadata in Azure Database with geo-replication for fault tolerance

Improve performance 10x-100x over traditional Spark deployments with an optimized environment

Machine Learning Product Investments

Machine Learning Investments

Enhance existing capabilities

Azure Databricks Remote support for other IDEs outside of native notebooks

MLFlow for better DevOps with Azure Databricks and other ML pipelines

Azure Machine Learning Python SDK support for popular IDEs & notebooks, including Azure Databricks

Azure Machine Learning managed compute capabilities

Introduce new models for FPGA scoring

Introduce new capabilities

Robust ONNX support - runtime engine in AML, model operationalization in SQL Server

Automated machine learning

Deploy and manage models to IoT edge

Extend Machine Learning services to SQL DB

Continue simplify machine learning

Leverage deep learning services and frameworks



AZURE ML SERVICE



Bring AI to the edge



Increase your rate of experimentation



Deploy and manage your models everywhere



AZURE DATABRICKS



Accelerate processing with the fastest Apache Spark engine



Integrate natively with Azure services



Access enterprise-grade Azure security

Leverage your favorite deep learning frameworks



TensorFlow



MS Cognitive Toolkit



PyTorch



Scikit-Learn



ONNX



Caffe2



MXNet

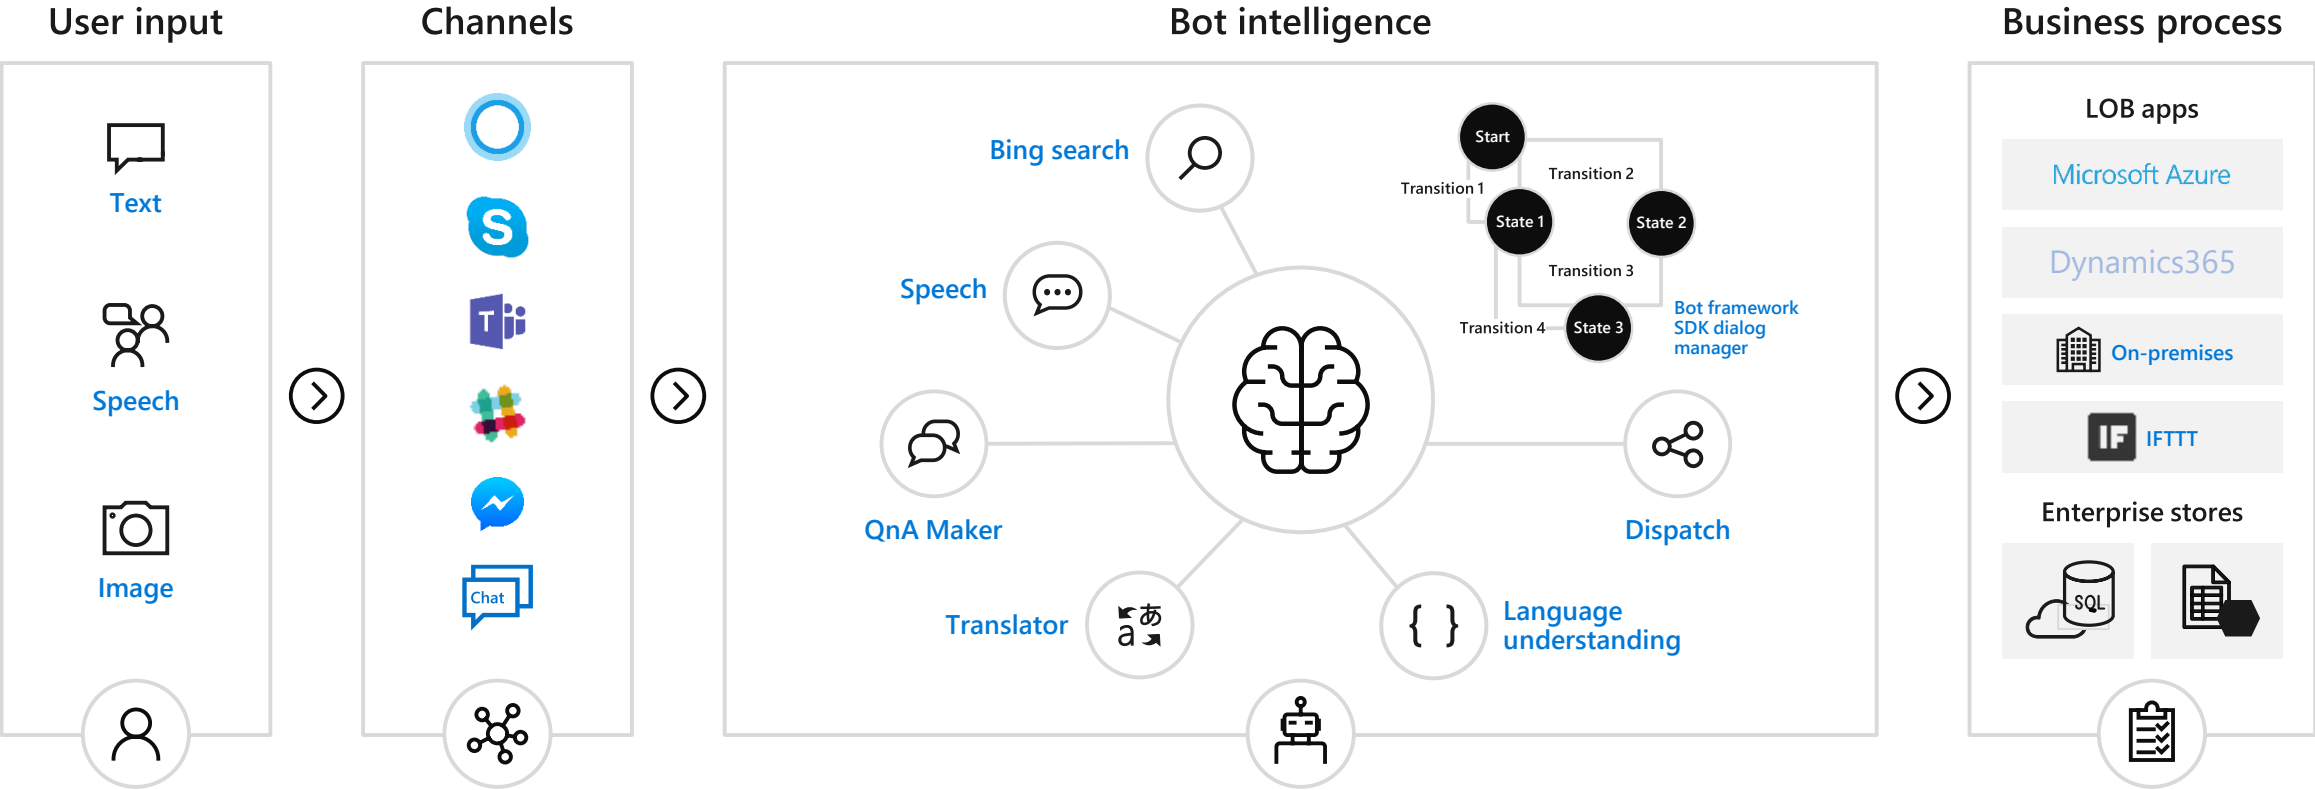


Chainer

Get started quickly with Developer AI

Conversational AI

Azure bot service + cognitive services



Azure Tools

Security

Logging

Auditing

Integration



Leverage out-of-the-box AI tools and services



Cognitive services



Use pre-built AI services to solve business problems



Map complex information and data



Allow your apps to process natural language



Azure search



Get up and running quickly



Reduce complexity with a fully-managed service



Use artificial intelligence to extract insights



Bot services



Speed development with a purpose-built environment for bot creation



Infuse intelligence into your bot using cognitive services




Integrate across multiple channels to reach more customers




Create a seamless developer experience across desktop, cloud, or at the edge using Visual Studio AI Tools

Enterprise scenarios for AI




Conversational agents

Transform your engagements with customers and employees




Of customer interactions powered by AI bots by 2025




Intelligent apps

Leverage AI to create the future of business applications




Applications to include AI by the end of this year



Business processes

Transform critical business processes with AI



Of enterprises using AI by 2020

AI-enabled devices



Custom AI

Drones

Smart kiosks



Enterprise AI

Smart factory

Smart kitchen

Smart bath



Developer AI

Voice-activated speakers

Smart cameras

Microsoft AI for accessibility

Empowering people with tools that amplify human capability



Employment

Facilitating development of professional skills



How Microsoft is helping the visually-impaired



Modern life

Delivering personalized experiences to improve independence



How Helpicto is helping non-verbal children



Human connection

Providing equal access to information and opportunity

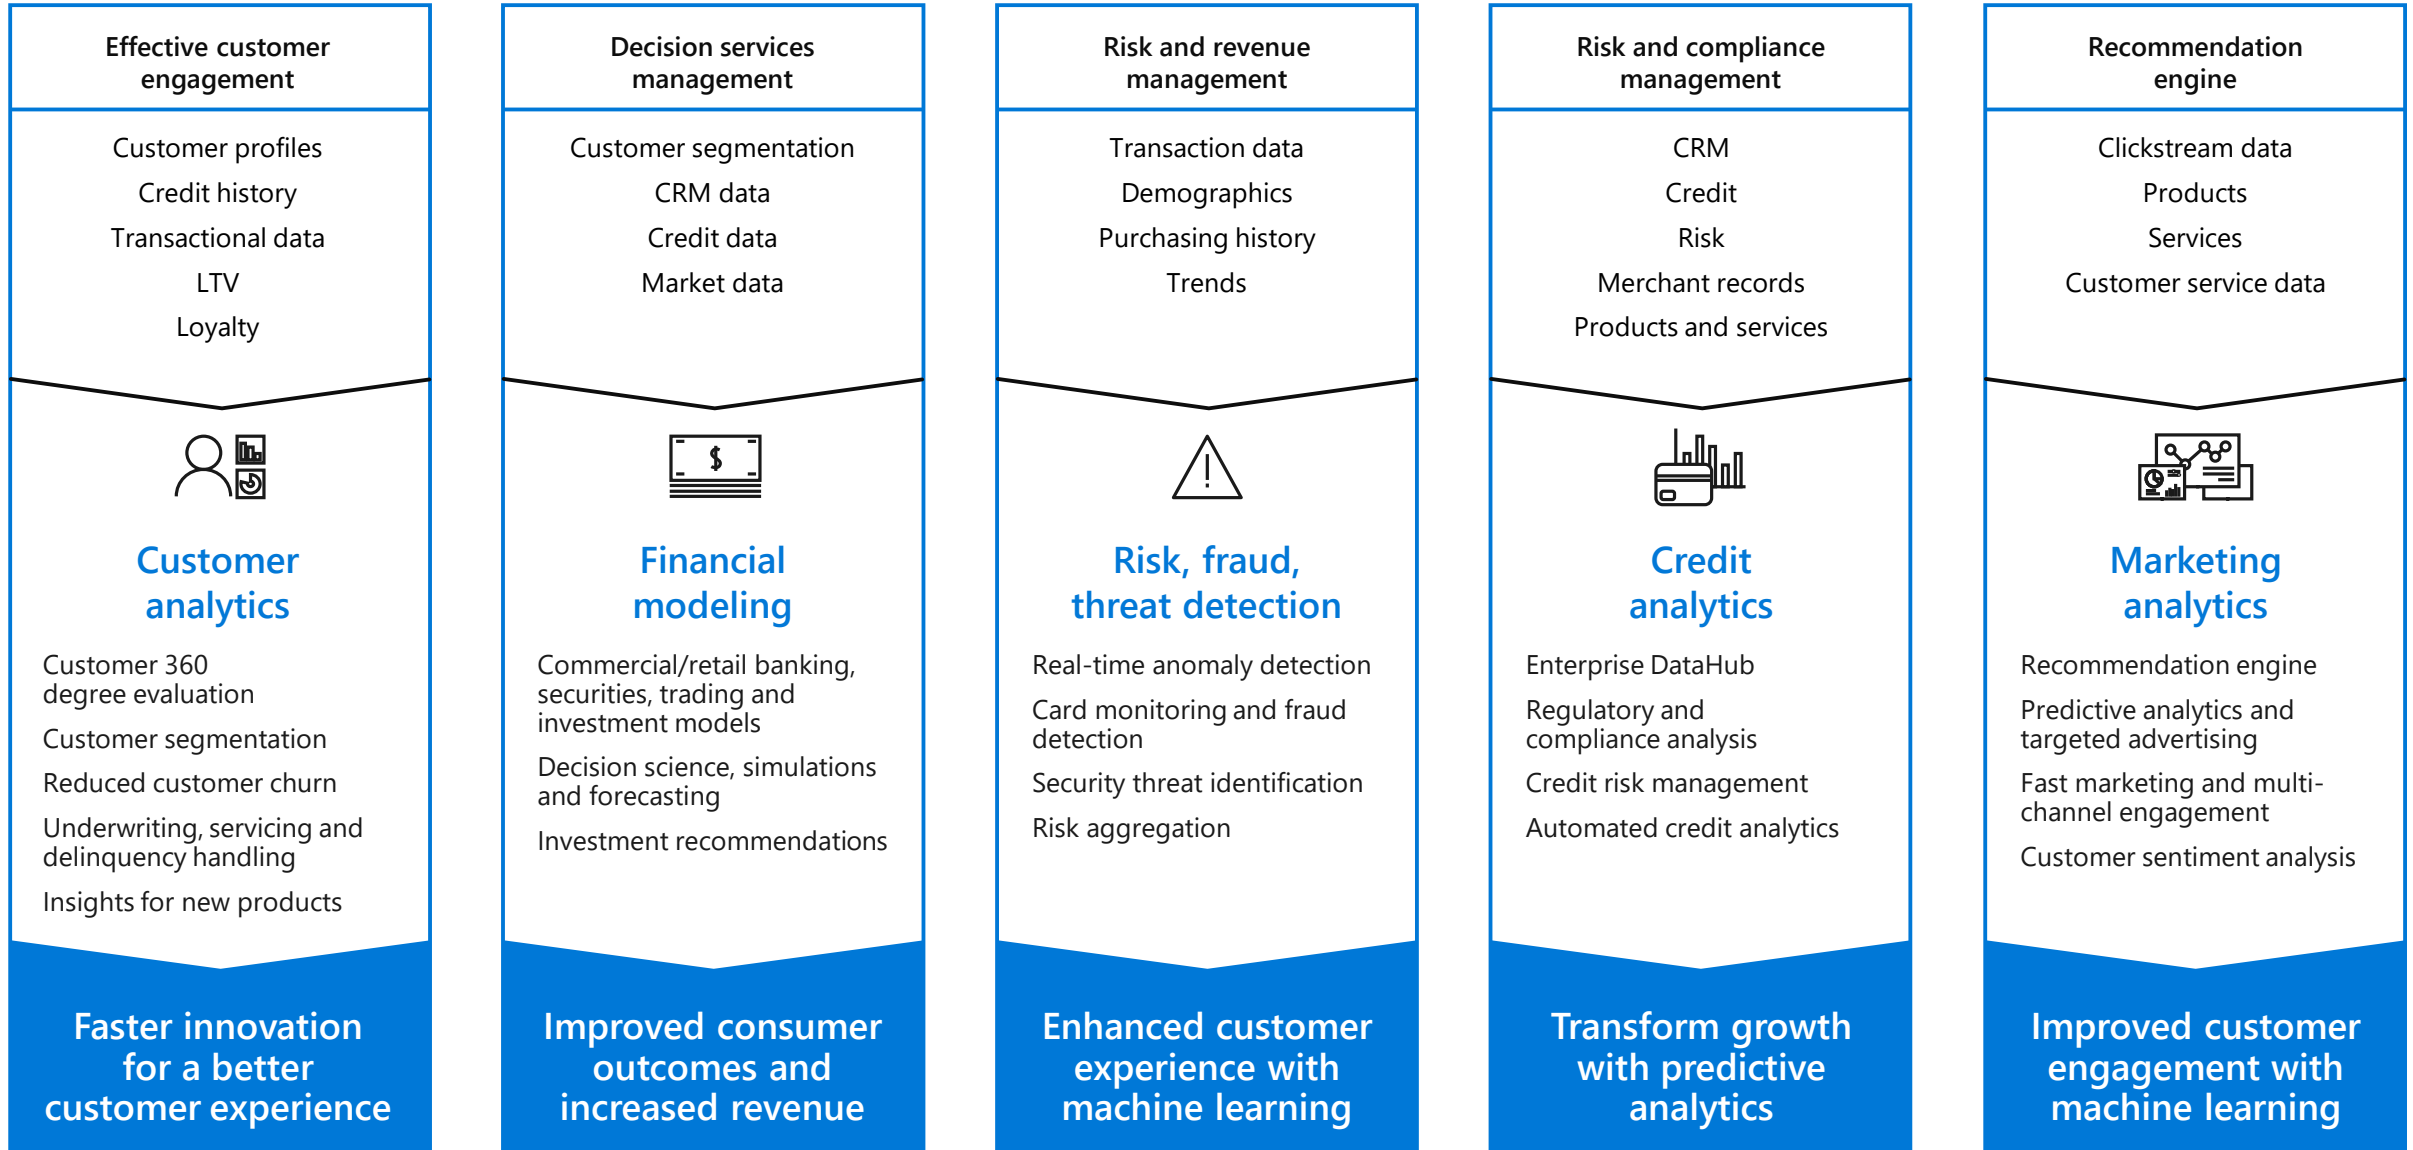


How RIT is helping the hearing-impaired



Data Science is for every industry

Financial services use cases



Health and life sciences use cases

DNA sequences

FAST-Q
BAM
SAM
VCF
Expression



Genomics and precision medicine

Single cell sequencing
Biomarker, genetic, variant and population analytics
ADAM and HAIL on Databricks

Faster innovation
for drug
development

Real world analytics

HL7/CCD
837
Pharmacy
Registry
EMR



Clinical and claims data

Claims data warehouse
Readmission predictions
Efficacy and comparative analytics
Prescription adherence
Market access analysis

Improved outcomes
and increased
revenue

Image deep learning

MRI
X-RAY
CT
Ultrasound



GPU image processing

Graphic intensive workloads
Deep learning using
Tensor Flow
Pattern recognition

Diagnostics
leveraging
machine learning

Sensor data

Readings
Time series
Event data



IoT device analytics

Aggregation of
streaming events
Predictive maintenance
Anomaly detection

Predictive analytics
transforms quality
of care

Social data listening

Social media
Adverse events
Unstructured



Social analytics

Real-time patient feedback
via topic modelling
Analytics across
publication data

Improved patient
communications
and feedback

Media and entertainment use cases

Personalized recommendations

Customer profiles
Viewing history
Online activity
Content sources
Channels



Content personalization

Personalized viewing and engagement experience
Click-path optimization
Next best content analysis
Improved real time ad targeting

**Faster innovation
for customer
experience**

Effective customer retention

Customer profiles
Online activity
Content distribution
Services data



Customer churn prevention

Quality of service and operational efficiency
Market basket analysis
Customer behavior analysis
Click-through analysis

**Improved consumer
outcomes and
increased revenue**

Information optimization

Consumption logs
Clickstream and devices
Marketing campaign responses



Recommendation engine

Ad effectiveness
Content monetization
Fraud detection
Information-as-a-service
High value user engagement

**Enhance user
experience with
machine learning**

Inventory allocation

Transactions
Subscriptions
Demographics
Credit data



Predictive analytics

Predict audience interests
Network performance and optimization
Pricing predictions
Nielsen ratings and projections
Mobile spatial analytics

**Predictive
analytics
transforms growth**

Consumer engagement analysis

Content metadata
Ratings
Comments
Social media activity



Sentiment analysis

Demand-elasticity
Social network analysis
Promotion events
time-series analysis
Multi-channel marketing attribution

**Improved consumer
engagement with
machine learning**

Retail use cases

Recommendation engine

Customer profiles
Shopping history
Online activity
Social network analysis



Next best and personalized offers

Customer 360/consumer personalization
Right product, promotion, at right time
Multi-channel promotion

Faster innovation for customer experience

Effective customer engagement

Shopping history
Online activity
Floor plans
App data



Store design and ergonomics

Path to purchase
In-store experience
Workforce and manpower optimization

Improved consumer outcomes and increased revenue

Inventory optimization

Demand plans
Forecasts
Sales history
Trends
Local events/weather patterns



Data-driven stock, inventory, ordering

Predict inventory positions and distribution
Fraud detection
Market basket analysis

Omni-channel shopping experience with machine learning

Inventory allocation

Demographics
Buyer perception
Consumer research
Market/competitive analysis



Assortment optimization

Economic modelling
Optimization for foot traffic, Online interactions
Flat and declining categories

Predictive analytics transforms growth

Consumer engagement

Historical sales data
Price scheduling
Segment level price changes



Real-time pricing optimization

Demand-elasticity
Personal pricing schemes
Promotion events
Multi-channel engagement

Improved consumer engagement with machine learning

Advertising and marketing tech use cases

Effective customer engagement

Customer profiles
Online history
Transaction data
Loyalty



Customer value analytics

Customer 360, segmentation aggregation and attribution
Audience modelling/index report
Reduce customer churn
Insights for new products
Historical bid opportunity as a service

Faster innovation for customer growth

Recommendation engine

Customer segmentation
CRM data
Credit data
Market data



Next best and personalized offers

Right product, promotion, at right time
Real time ad bidding platform
Personalized ad targeting
Ad performance reporting

Improved outcomes and increased revenue

Risk and fraud analysis

Transaction data
Demographics
Purchasing history
Trends



Risk and fraud management

Real-time anomaly detection
Fraud prevention
Customer spend and risk analysis
Data relationship maps

Risk management with machine learning

Campaign reporting analytics

CRM
Merchant records
Products
Services
Marketing data



Sales and campaign optimization

Optimizing return on ad spend and ad placement
Multi-channel promotion
Ideal customer traits
Optimized ad placement

Predictive analytics transforms growth

Brand promotion and customer experience

Social media
Online history
Customer service data



Sentiment analysis

Opinion mining/social media analysis
Deeper customer insights
Customer loyalty programs
Shopping cart analysis

Improved customer engagement with machine learning

Oil, gas, and energy use cases

Upstream optimization, maximize well life

Field data
Asset data
Demographics
Production data



Digital oil field/ oil production

Production optimization
Integrate exploration
and seismic data
Minimize lease
operating expenses
Decline curve analysis

**Faster innovation
for revenue
growth**

Grid operations, asset inventory optimization

Sensor stream data
UAVs images
Inventory data
Production data



Industrial IoT

Pipeline monitoring
Preventive maintenance
Smart grids and microgrids
Grid operations, field service
Asset performance
as a service

**Improved outcomes
and increased
revenue**

Supply-chain optimization

Transaction data
Demographics
Purchasing history
Trends



Supply-chain optimization

Trade monitoring,
optimization
Retail mobile applications
Vendor management -
construction, transportation,
truck and delivery
optimization

**Optimizing supply-
chain with machine
learning**

Risk optimization

Sensor stream data
Transport
Retail data
Grid production data
Refinery tuning parameters



Safety and security

Real-time anomaly detection
Predictive analytics
Industrial safety
Environment health and safety

**Predictive analytics
transforms safety
and security**

Recommendations engine

Clickstream data
Products
Services
Market data
Competitive data
Demographics

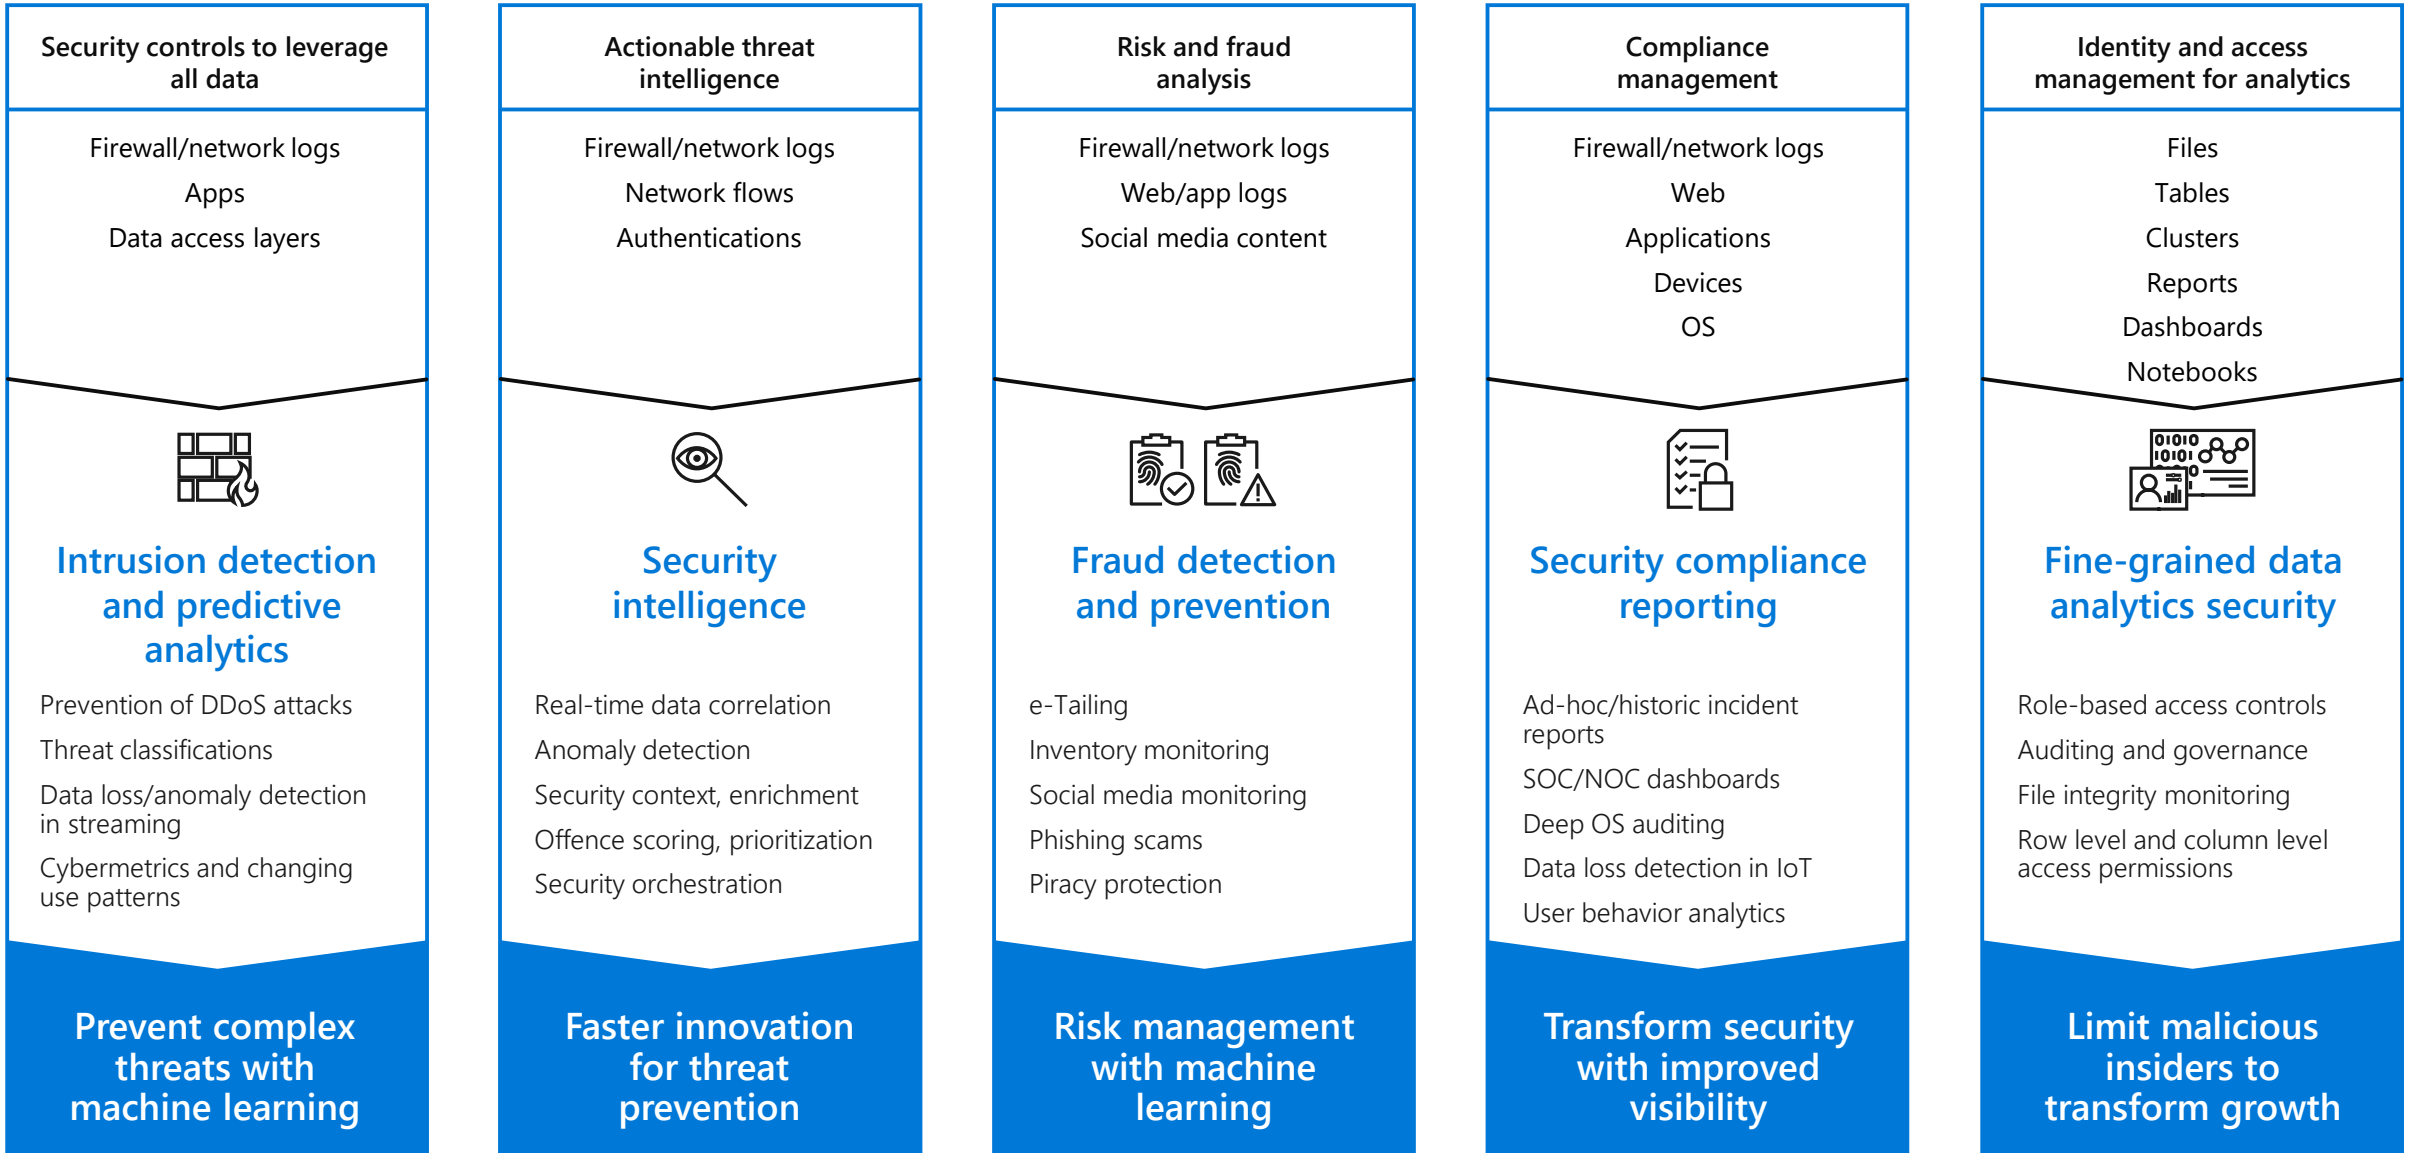


Sales and marketing analytics

Fast marketing and
multi-channel engagement
Develop new products and
monitor acceptance of rates
Predictive energy trading
Deep customer insights

**Improved customer
engagement with
machine learning**


Security use cases



Our unique differentiators

Microsoft AI breakthroughs

Vision

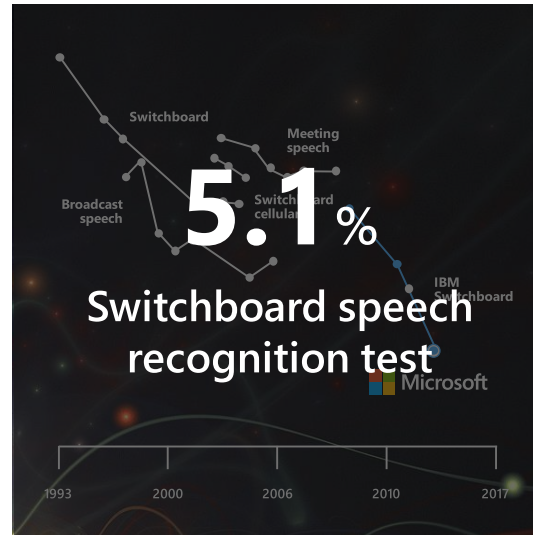


2016

First to achieve

Object recognition
Human parity

Speech



2017

First to achieve

Speech recognition
Human parity

69.9%


MT research system

March 2018

First to achieve

Machine translation
Human parity

Language




January 2018

First to achieve


Machine reading comprehension
Human parity

Ongoing Momentum



Vision



Speech



Language



Our differentiated value proposition

Accelerate time to value
with agile tools and services



Pretrained AI
services



Powerful
tools



Comprehensive
platform

Innovate with AI everywhere –
in the cloud, at edge and on-premises



Cloud



Edge



On-premises

Use any language, any development
tool and any framework




Benefit from industry-leading security, privacy,
compliance, transparency, and AI ethics standards

>90% of Fortune 500 companies
use Microsoft Cloud


It's all on Microsoft Azure

Microsoft Azure




Productive

Accelerate time to market




Hybrid

Optimize your infrastructure



Intelligent

Innovate at scale



Trusted

Develop with confidence

Companies using Azure for ML and AI



UBER



PROGRESSIVE

JABIL



accenture



TaylorWessing



HEALTHVILLAGE.FI

pivothead



PLEXURE



dutchcrafters



HChavarri

PicCOLLAGE

devon

R·I·T



JUNIPER

EQUADEX



365mc




Rich partner network

Data integration



System integration



BI and analytics



Thank you

Appendix

ONNX is the new open ecosystem for AI models


> [Announcing ONNX Runtime open source](#)

Create

Frameworks



Services



ONNX Model

Deploy

Azure

Azure Machine Learning services

Ubuntu VM

Windows Server 2019 VM

Devices

Windows devices

Other devices (iOS, etc.)

Productive Services

To empower data science and development teams




Azure Machine Learning

Python-based machine learning service

Develop models faster with automated machine learning

Use any Python environment and ML frameworks

Manage models across the cloud and the edge.



Azure Databricks

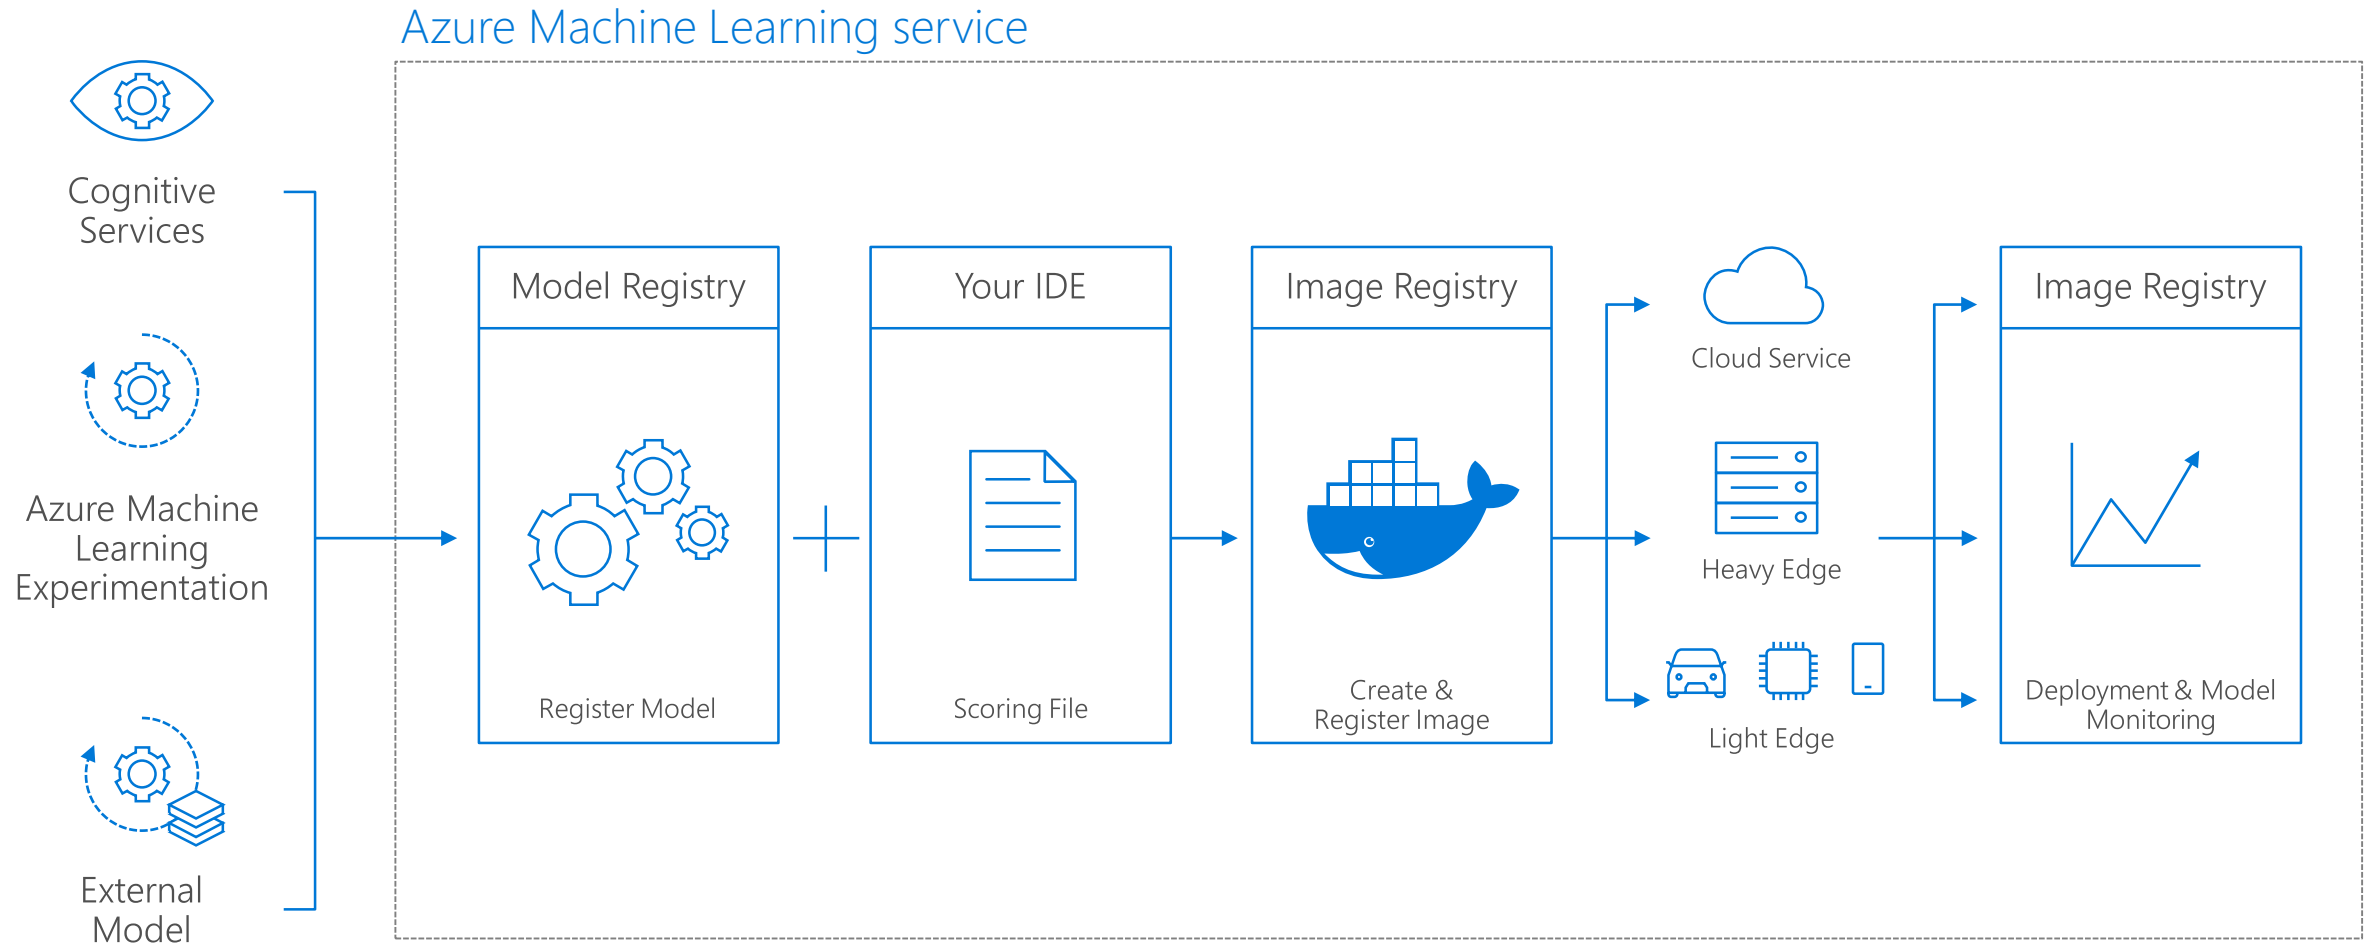
Apache Spark-based big-data service

Prepare data clean data at massive scale

Enable collaboration between data scientists and data engineers

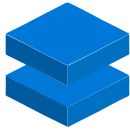
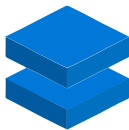




Access machine learning optimized clusters

Deploy Azure ML models at scale




Productive Services

What to use when?

Customer use case	Data Prep	Build & Train	Manage and Deploy
Big Data/Apache Spark	 Azure Databricks (Apache Spark Dataframes, Datasets, Delta, Pandas, NumPy etc.)	 +  Azure Databricks + Azure ML service (Spark MLlib and OSS frameworks + Automated ML, Model Registry)	 Azure ML service (containerize, deploy, inference and monitor)
Data Science	Pandas, NumPy etc.	 Azure ML service (OSS frameworks, Hyperdrive, Pipelines, Automated ML, Model Registry)	 Azure ML service (containerize, deploy, inference and monitor)

DevOps loop for data science

Prepare





Prepare
Data

Experiment



Build model
(your favorite IDE)



Train &
Test Model



Register and
Manage Model

Deploy



Build
Image



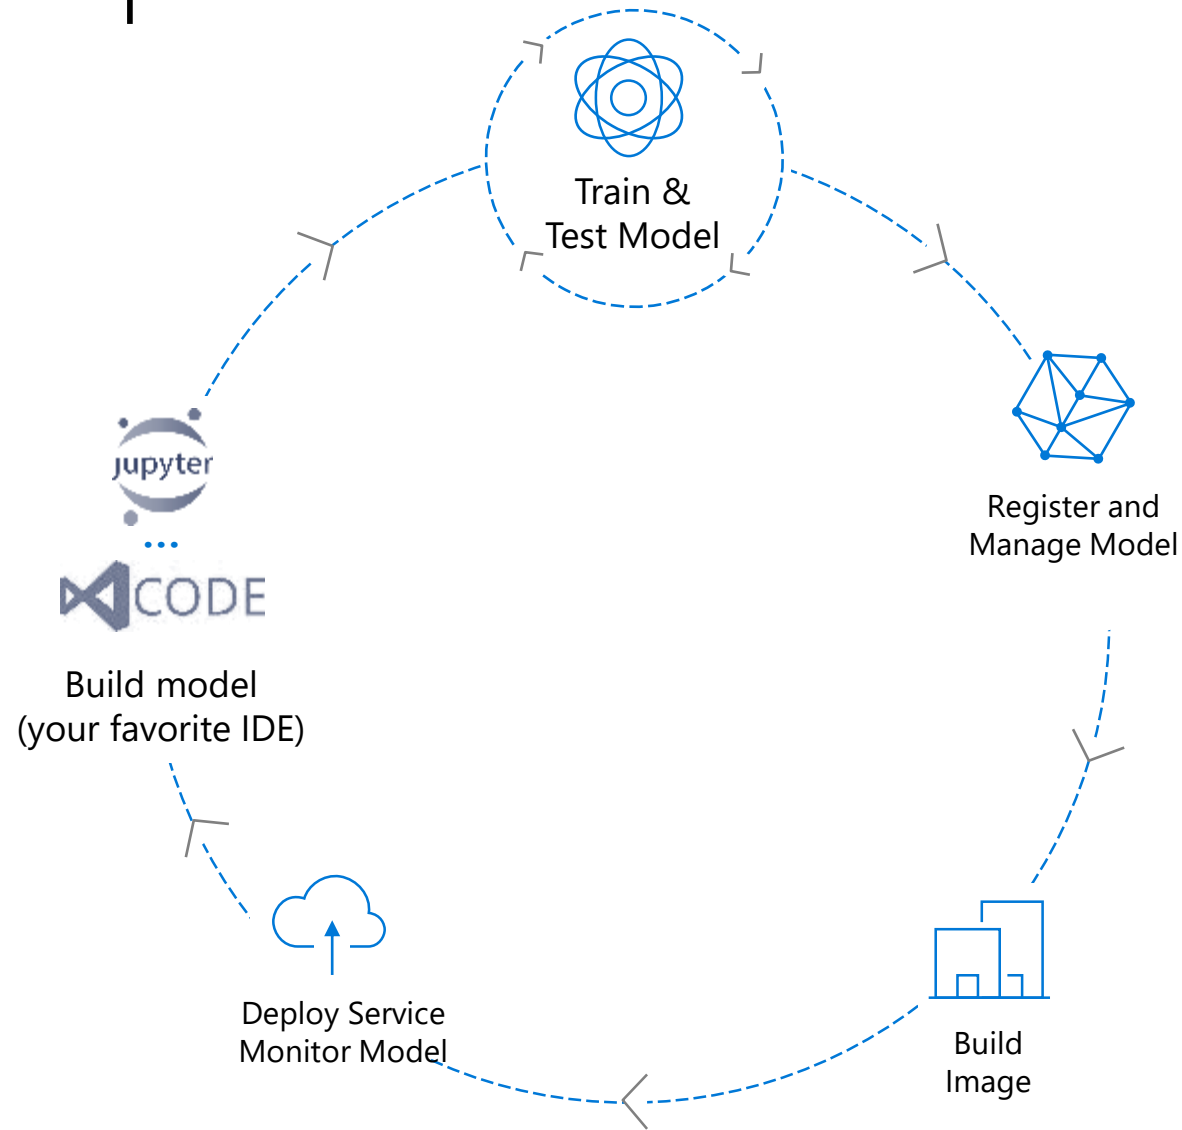
Deploy Service
Monitor Model

DevOps loop for data science

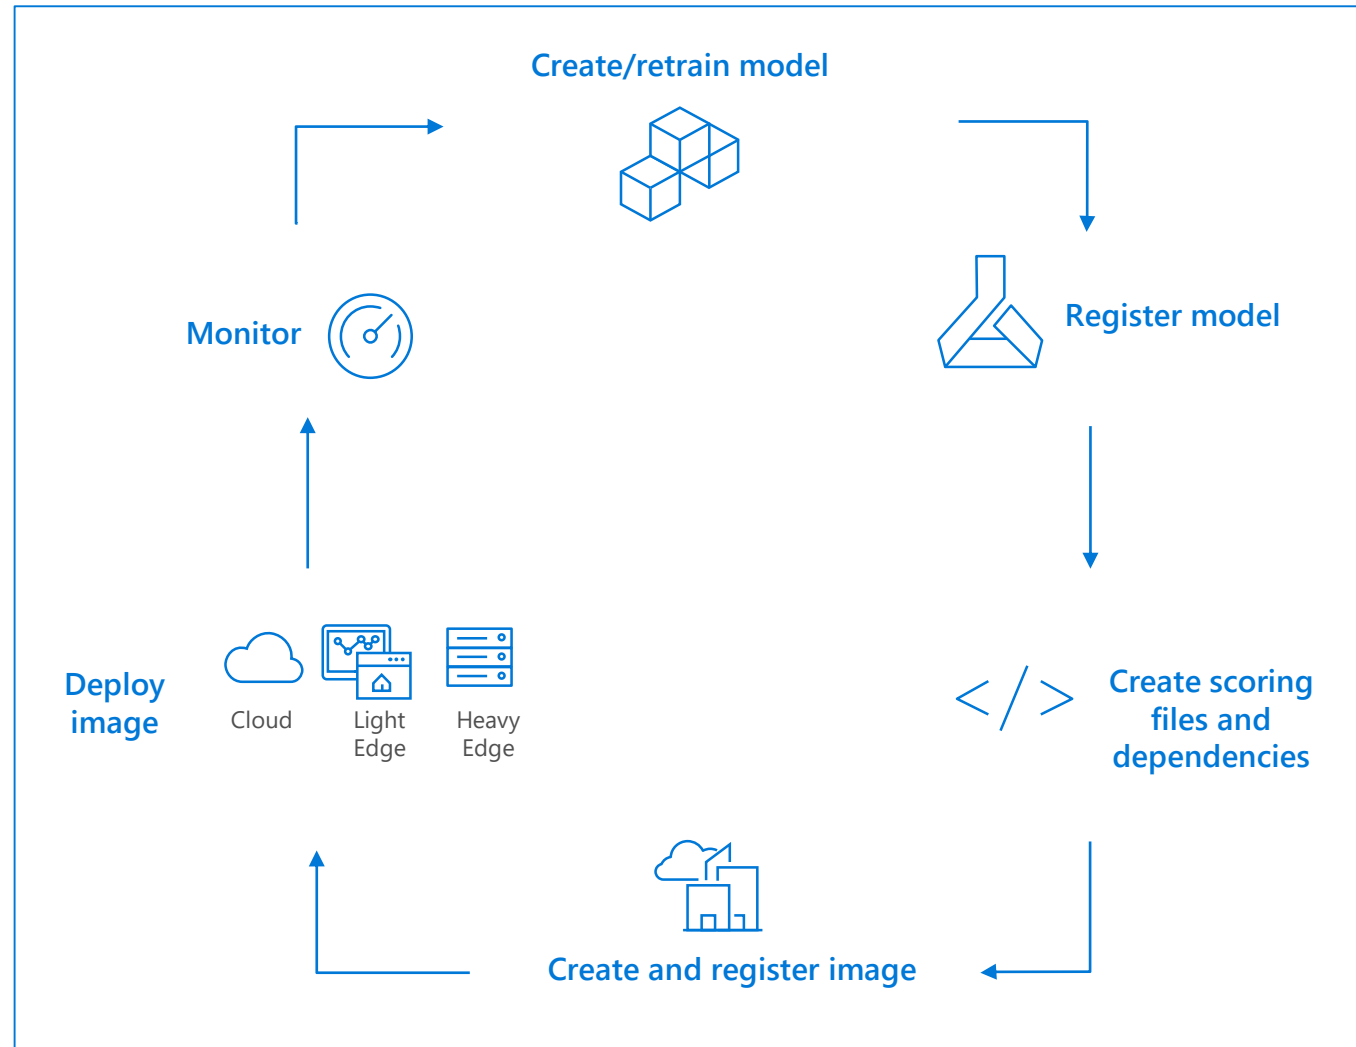
Prepare




Prepare
Data



Model management in Azure Machine Learning




Experimentation




Leverage service-side capture of run metrics, output logs and models



Manage training jobs locally, scaled-up or scaled-out



Use leaderboards, side by side run comparison and model selection



Conduct a hyperparameter search on traditional ML or DNN