Microsoft Azure Machine Learning

Cortana Analytics Suite

Transform data into intelligent action

Cortana Analytics Suite

Transform data into intelligent action

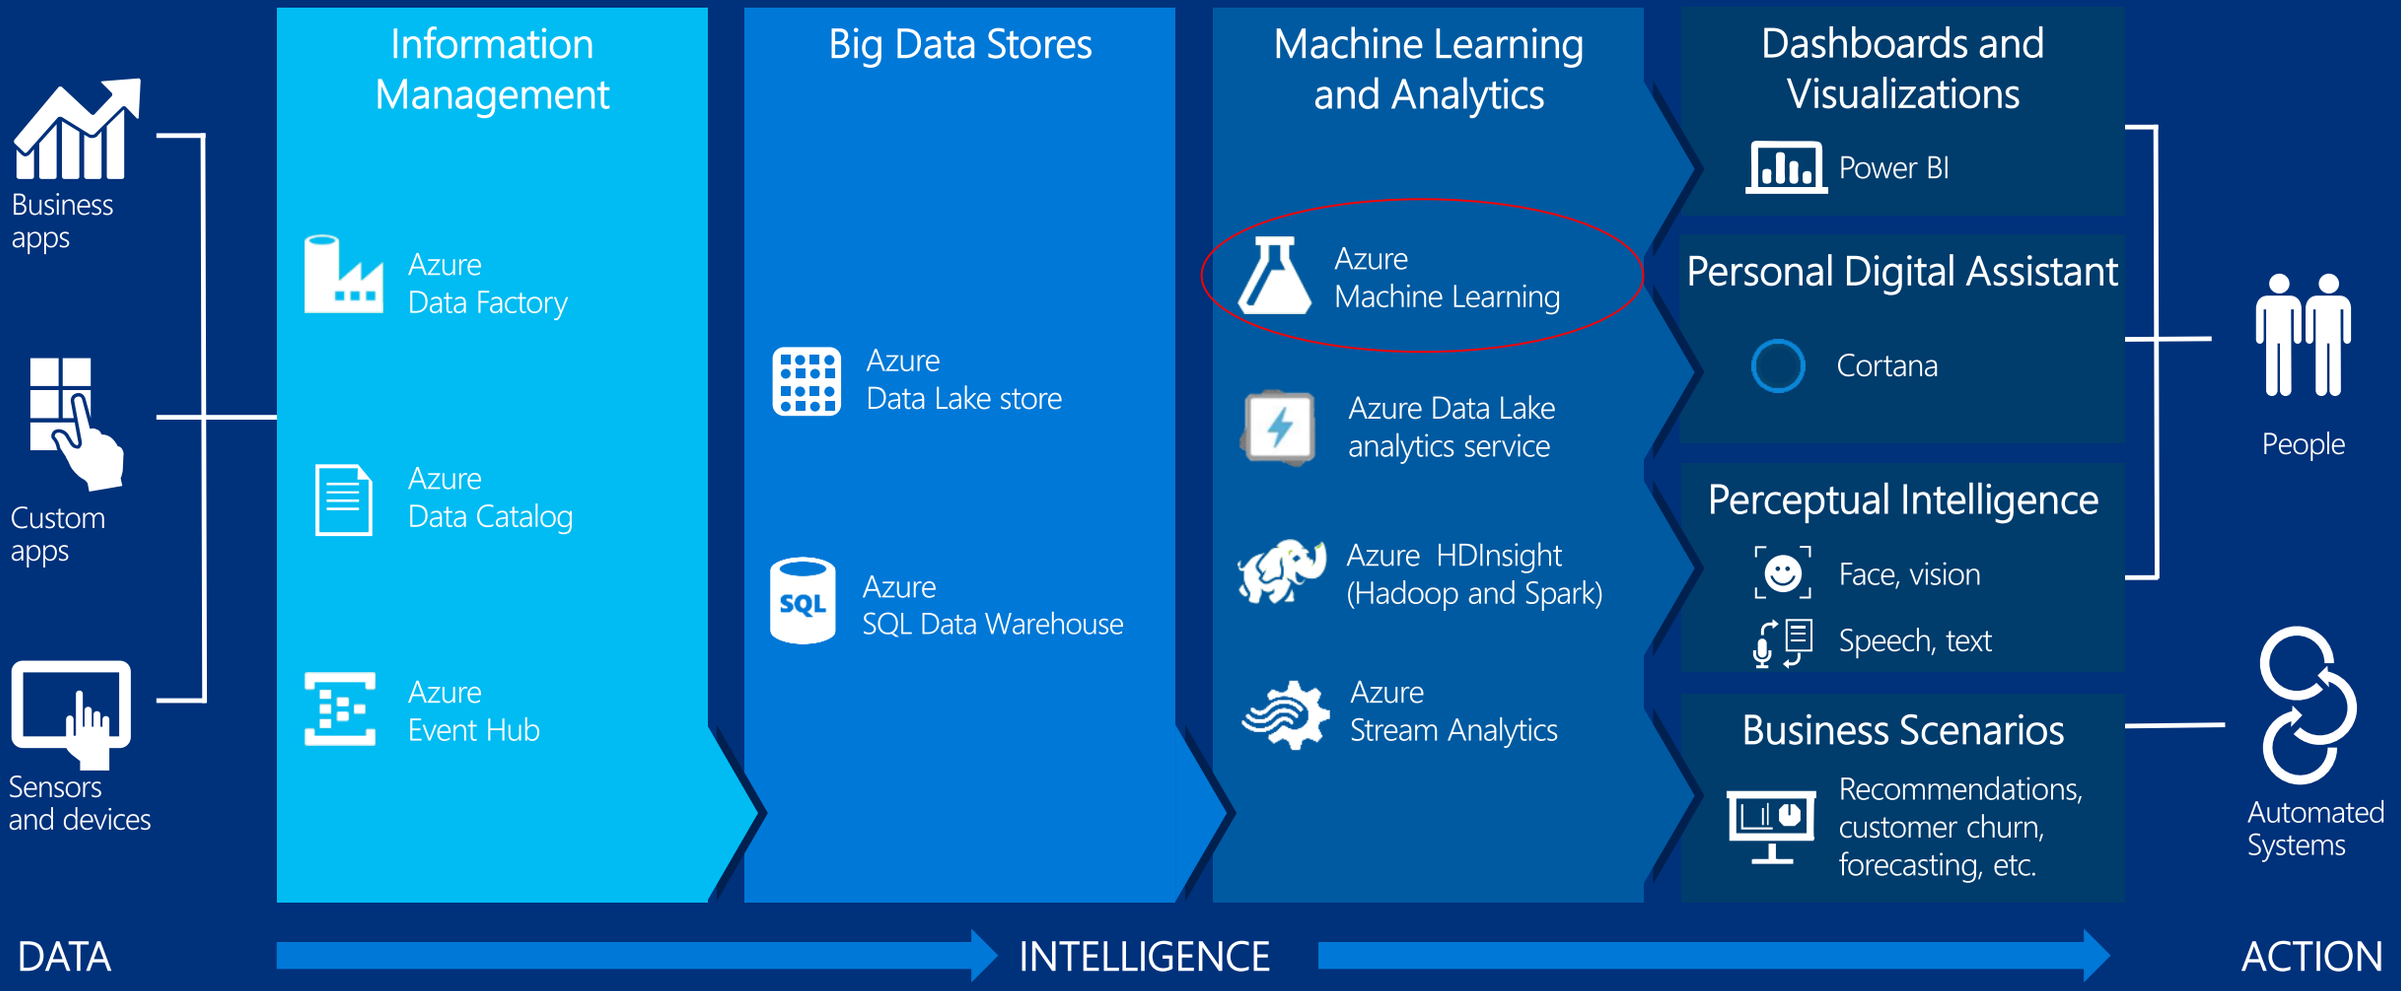

Microsoft & Machine Learning

Answering questions with experience

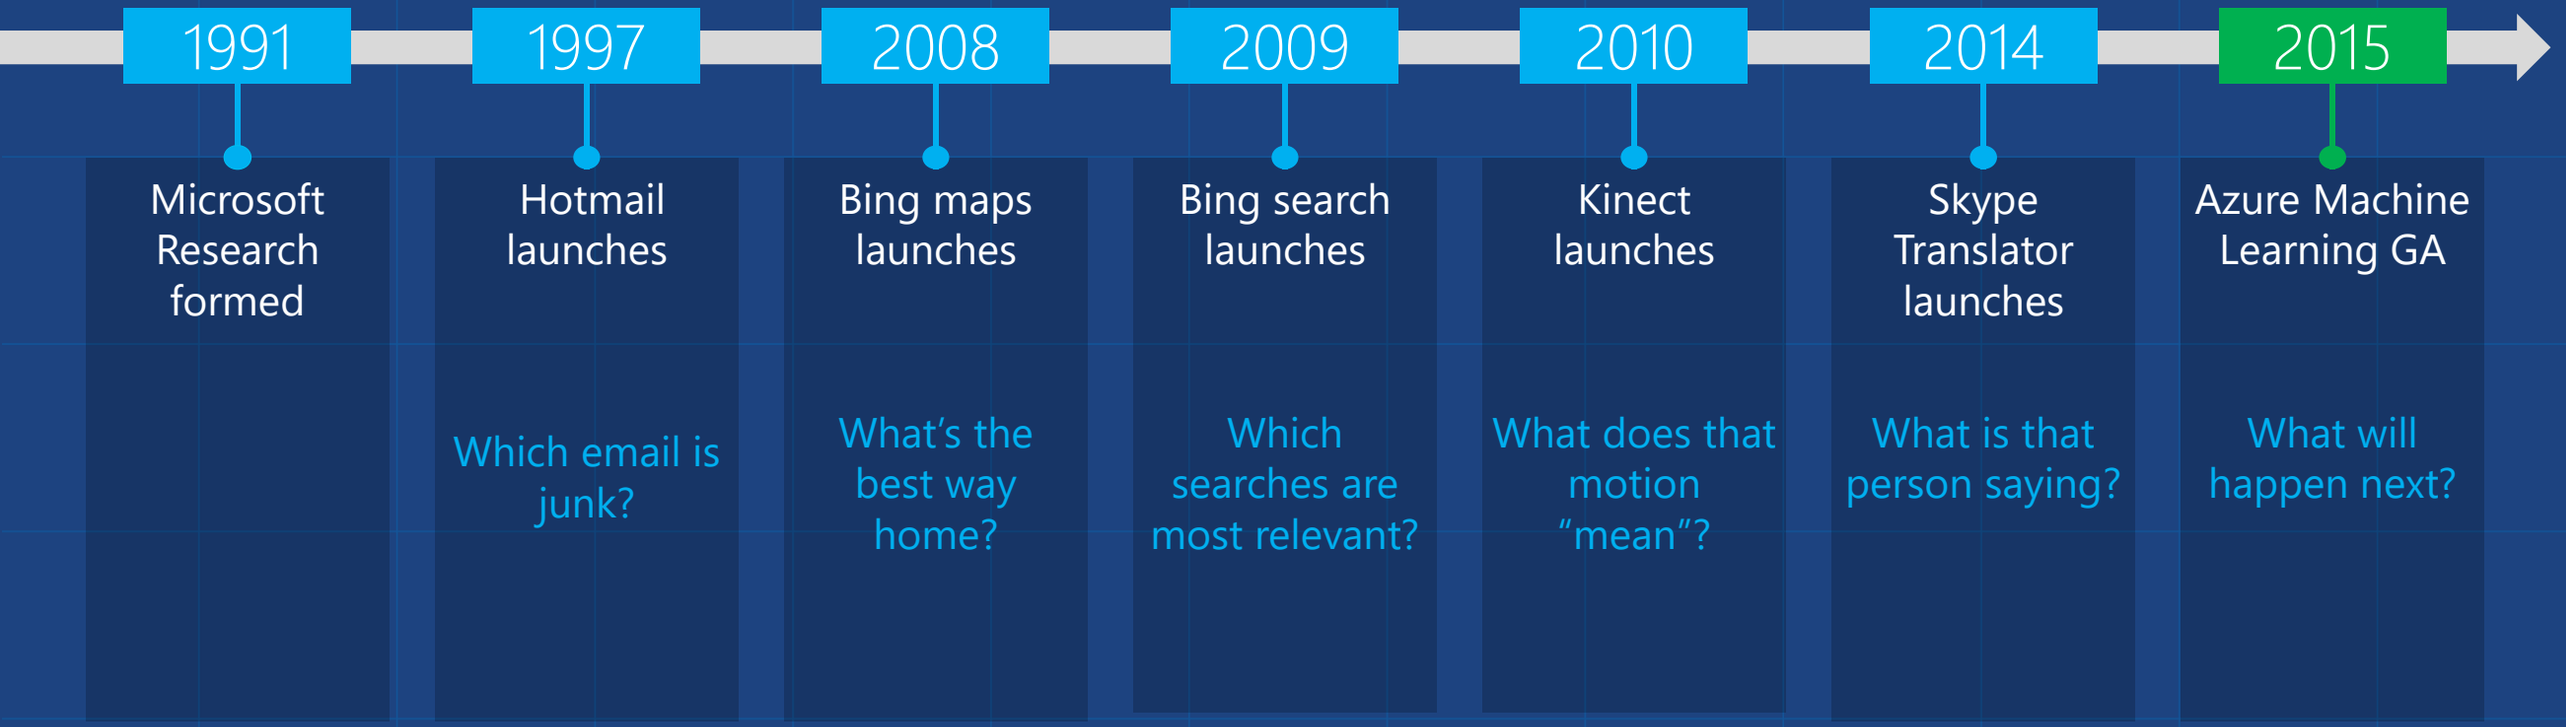

Machine learning is pervasive throughout Microsoft products.

Differentiation

Model Your Way

[Data Scientist]

All Skill Levels

Business-tested Algorithms

R & Python

Deploy in Minutes

[Data Scientist, IT & Developers]

One Click Deployment

Manage via Cloud Portal

Model accessed as a Web Service

Expand your Reach

[Ecosystem & Developers]

Product Gallery

Azure Marketplace

The Data Science “Inside”

Azure Machine Learning Service

Data -> Predictive model -> Operational web API in minutes

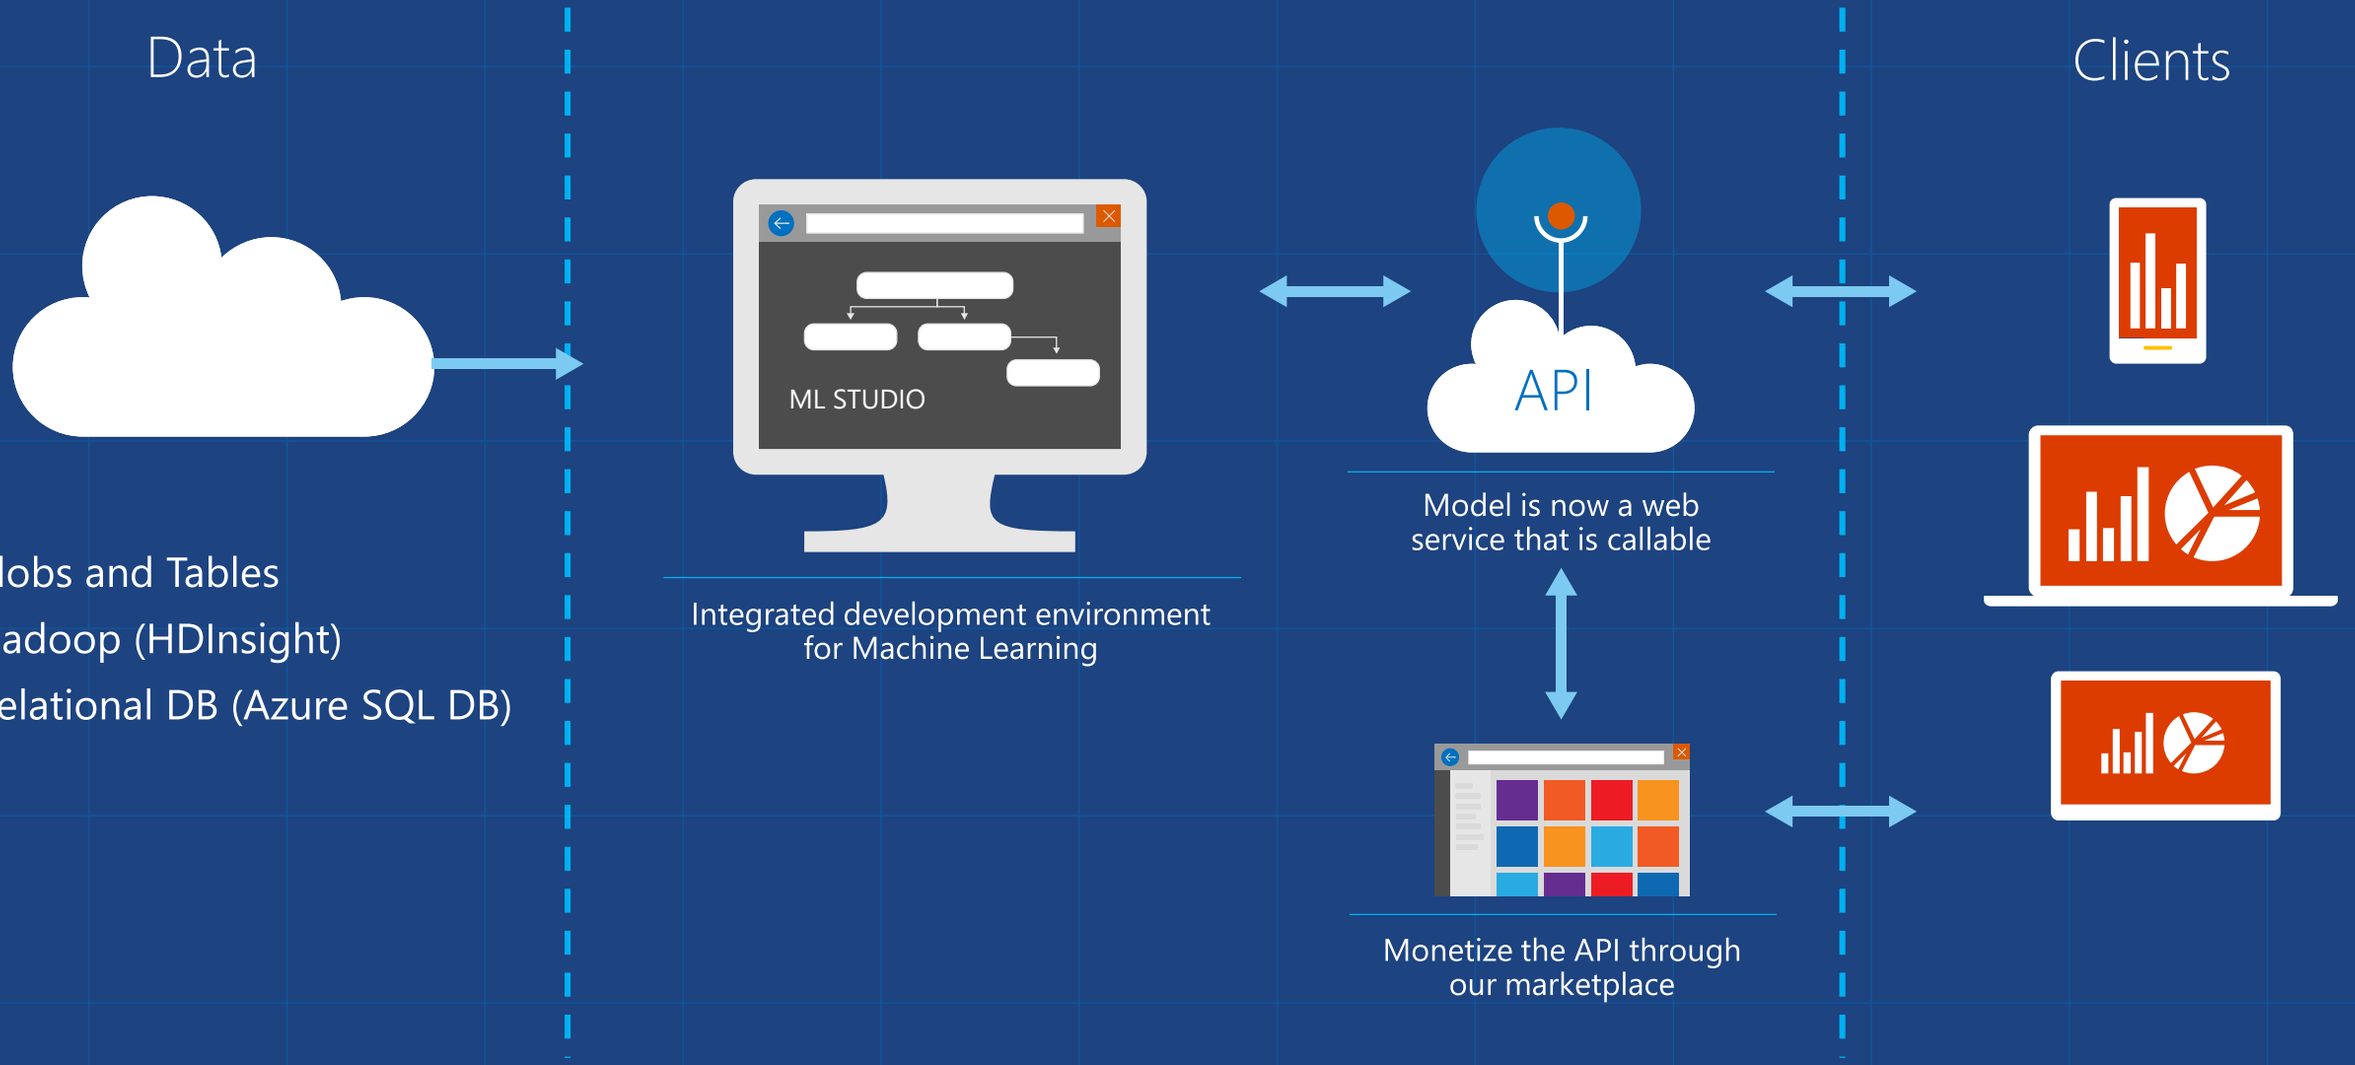

Azure Machine Learning Service

Data -> Predictive model -> Operational web API in minutes

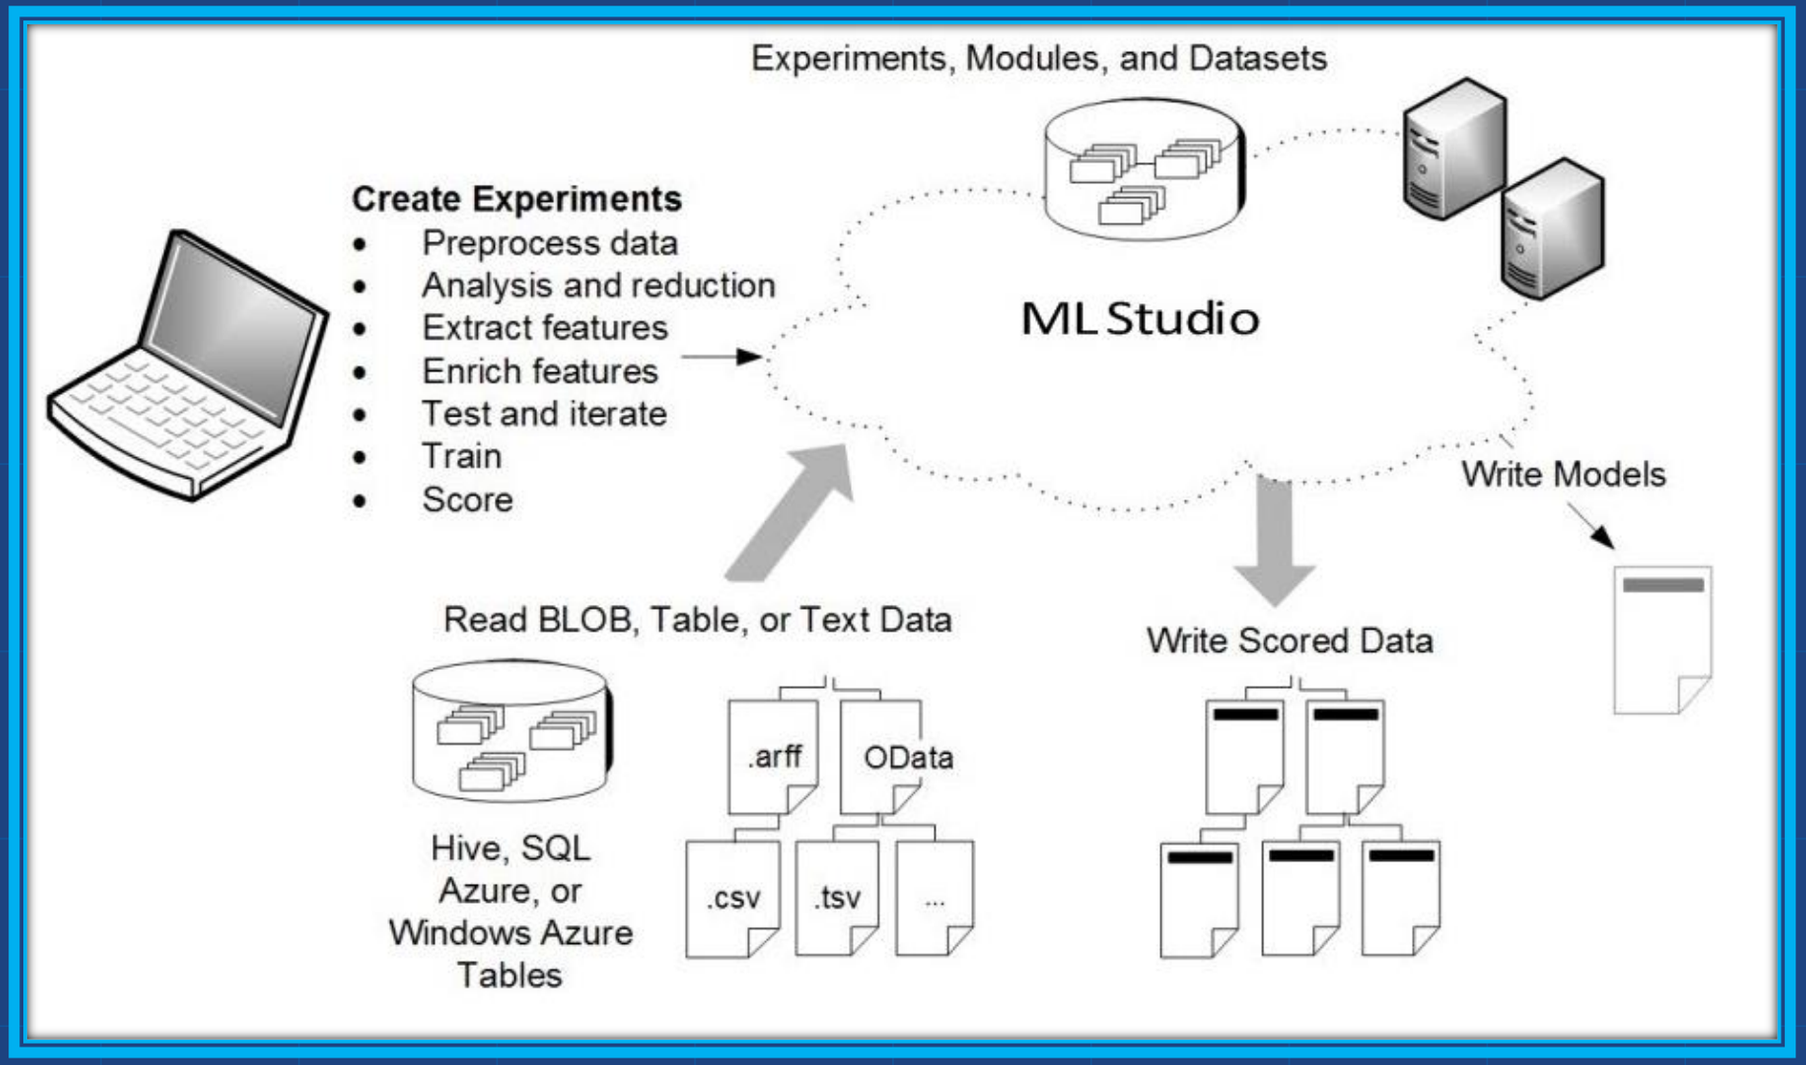

- Collaborative
- Drag-and-drop tool

- Build
- Test
- Deploy

Predictive Analytics solutions on your data.

AML publishes models as web services that can easily be consumed by custom apps or BI tools such as Excel

What can Azure ML do for you…?

Telemetry data analysis

Buyer propensity models

Social network analysis

Predictive maintenance

Web app optimization

Churn analysis

Natural resource exploration

Weather forecasting

Healthcare outcomes

Fraud detection

Life sciences research

Targeted advertising

Network intrusion detection

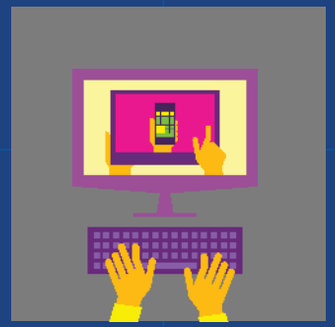

Smart meter monitoring

Machine Learning in ML Studio

Anomaly Detection
- One-class Support Vector Machine
- Principal Component Analysis-based Anomaly Detection
- Time Series Anomaly Detection*

Classification
- **Two-class Classification**
  - Averaged Perceptron
  - Bayes Point Machine
  - Boosted Decision Tree
  - Decision Forest
  - Decision Jungle
  - Logistic Regression
  - Neural Network
  - Support Vector Machine
- **Multi-class Classification**
  - Decision Forest
  - Decision Jungle
  - Logistic Regression
  - Neural Network
  - One-vs-all

Clustering
- K-means Clustering

Recommendation
- Matchbox Recommender

Regression
- Bayesian Linear Regression
- Boosted Decision Tree
- Decision Forest
- Fast Forest Quantile Regression
- Linear Regression
- Neural Network Regression
- Ordinal Regression
- Poisson Regression

Statistical Functions
- Descriptive Statistics
- Hypothesis Testing T-Test
- Linear Correlation
- Probability Function Evaluation

Text Analytics
- Feature Hashing
- Named Entity Recognition
- Vowpal Wabbit

Computer Vision
- OpenCV Library

https://studio.azureml.net

Guest Access Workspace: Free trial access without logging in.
Free Workspace: Free persisted access, no Azure subscription needed.
Standard Workspace: Full access with SLA under an Azure subscription.

Data/Model Visualization
- Scatterplots
- Bar Charts
- Box plots
- Histogram
- R and Python Plotting Libraries
- REPL with Jupyter Notebook
- ROC, Precision/Recall, Lift
- Confusion Matrix
- Decision Tree*

Training
- Cross Validation
- Retraining
- Parameter Sweep

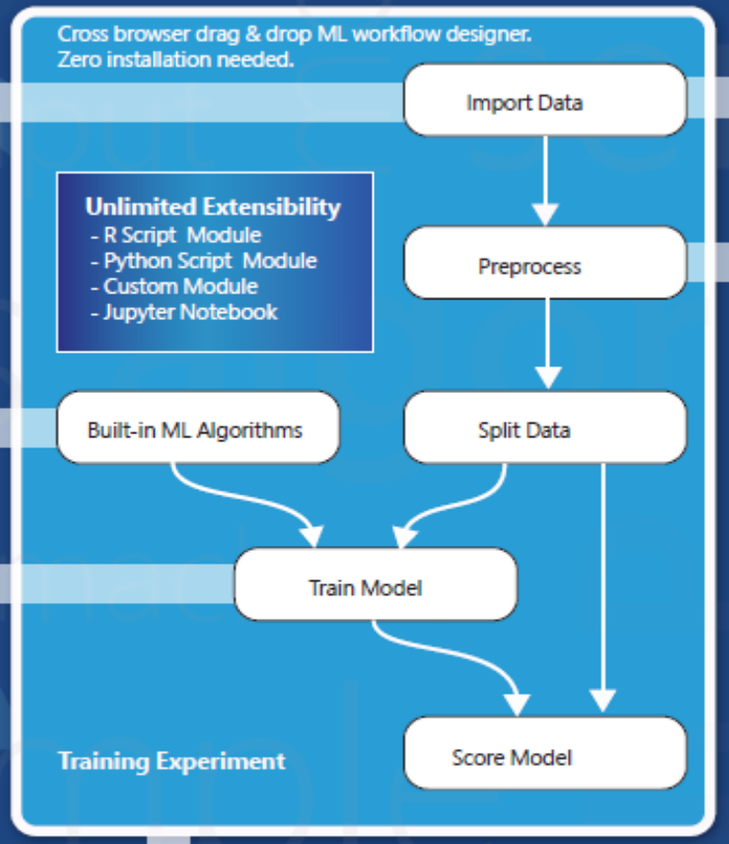

Unlimited Extensibility
- R Script Module
- Python Script Module
- Custom Module
- Jupyter Notebook

Data Source
- Azure Blob Storage
- Azure SQL DB
- Azure SQL DW*
- Azure Table
- Desktop Direct Upload
- Hadoop Hive Query
- Manual Data Entry
- OData Feed
- On-prem SQL Server*
- Web URL (HTTP)

Data Format
- ARFF
- CSV
- SVMLight
- TSV
- Excel
- ZIP

Data Preparation
- Clean Missing Data
- Clip Outliers
- Edit Metadata
- Feature Selection
- Filter
- Learning with Counts
- Normalize Data
- Partition and Sample
- Principal Component Analysis
- Quantize Data
- SQLite Transformation
- Synthetic Minority Oversampling Technique

Enterprise Grade Cloud Service
- SLA: 99.95% Guaranteed Up-time
- Azure AD Authentication
- Compute at Large Scale
- Multi-geo Availability
- Regulatory Compliance*

One-click Operationalization

Predictive Experiment

Make Prediction with Elastic APIs
- Request-Response Service (RRS)
- Batch Execution Service (BES)
- Retraining API

Community
- Gallery (http://gallery.azureml.net)
- Samples & Templates
- Workspace Sharing and Collaboration
- Live Chat & MSDN Forum Support

* Feature Coming Soon

Appendix

Model Your Way: Open source/our source

Script with R, SQLite or Python

CPython 2.7 support from inside AML Studio

numpy/scipy/panda/scikit-learn/etc.

Anaconda distro pre-installed

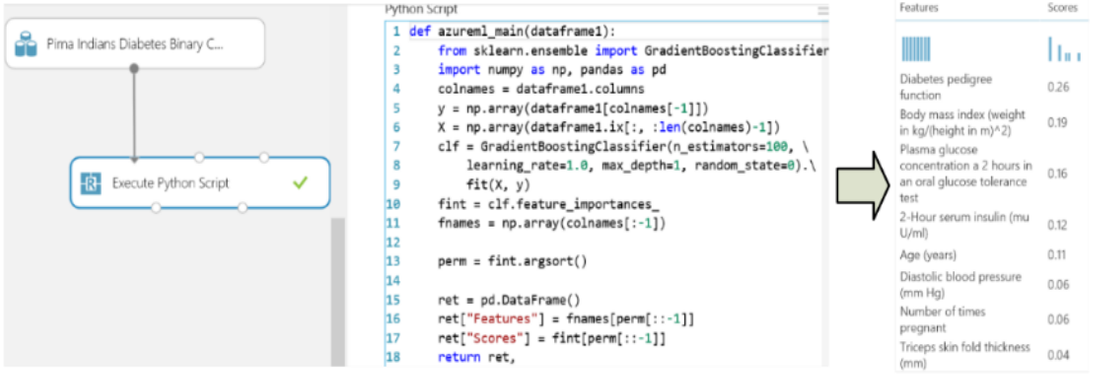

Python client library

Analyze data using Python and its libraries

Use IPython, PTVS, Eclipse to edit/debug

Big learning with counts

TB scale datasets

Modular: tune/monitor/replace in isolation

Monitorable and debuggable

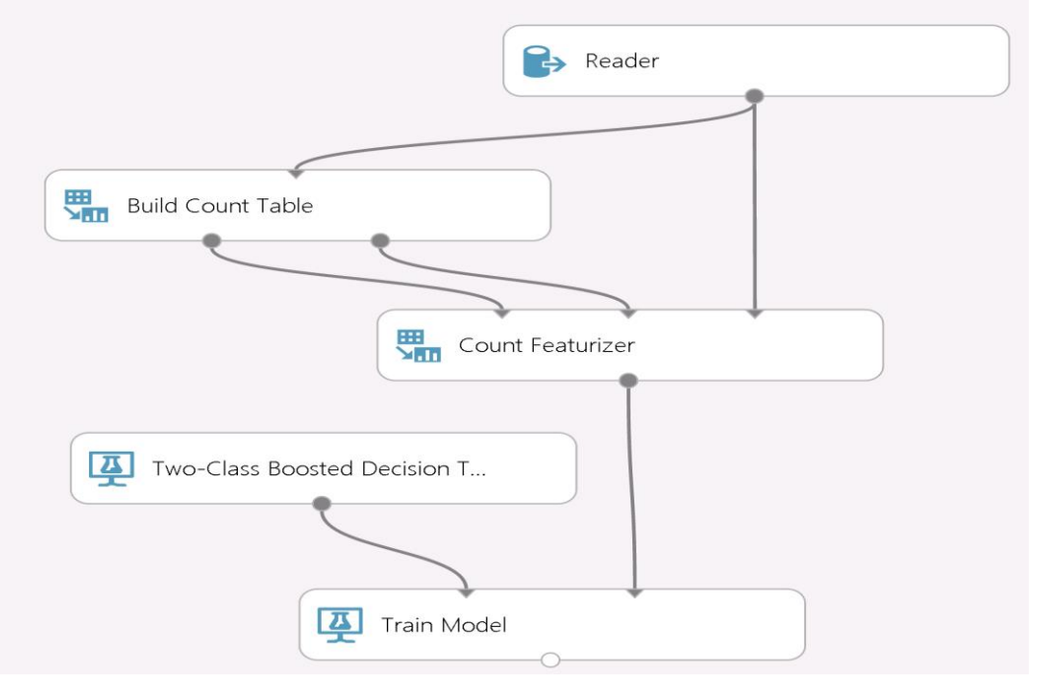

Deploy in Minutes

One click to production

Publish as a Web Service or to Gallery

Continuous updates to streamline process

Stay tuned to our blog for more

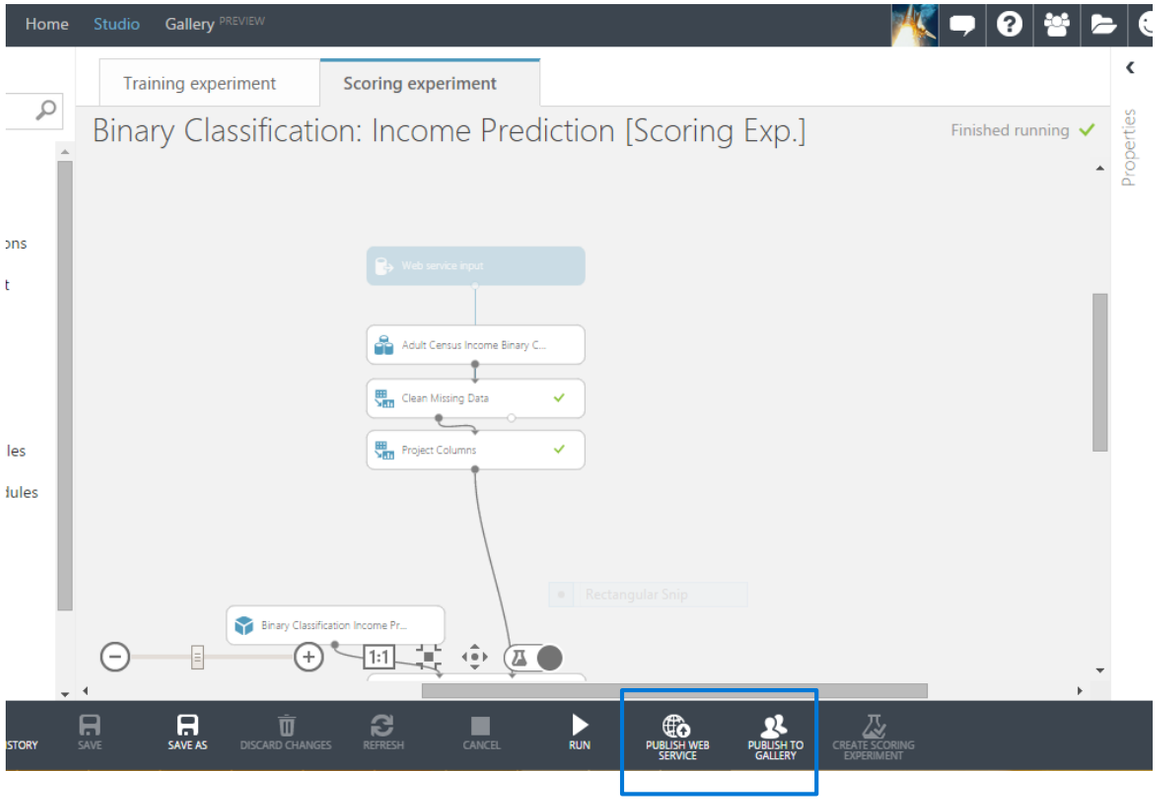

Expand your Reach

New in-product Gallery

Discover what others have built

Learn by dropping these into your workspace

Share your work with others

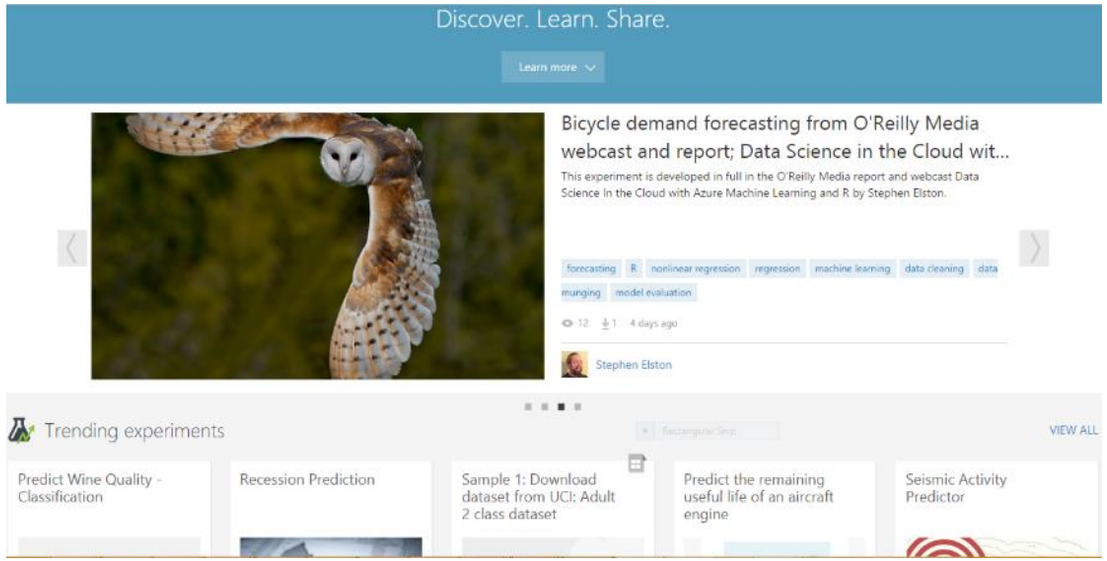

Carnegie Mellon University

“We see Azure ushering in an era of self-service predictive analytics for the masses. We can only imagine the possibilities”

Bertrand Lasternas
Carnegie Mellon

Smart buildings

CMU wanted to use sensor data for more than reactive repair and diagnosis

Smart Buildings

The CMU Center for Building Performance and Diagnostics studies the operational efficiency of its buildings worldwide. Before partnering with OSISoft and Microsoft, CMU was challenged to **use sensor data for more than reactive repair and diagnosis.**

Competitive advantage

20% savings on energy costs; several hundred thousand dollars campus-wide

Researchers focus on what's next; technicians fix problems before they start

Business transformation

Scale: same headcount with significantly more reach and impact

Carnegie Mellon University

"We see Azure ushering in an era of self-service predictive analytics for the masses. We can only imagine the possibilities."

Bertrand Lasternas
Carnegie Mellon