



Microsoft

Azure Analytics

R Services

Microsoft Research

Alexandre Gattiker

(algattik@microsoft.com)

Data Platform Solution Architect

Platform Services

Security & Management

- Service Creation & Configuration
- User/Group Directory Store
- Identity Sign-Up and sign-in
- Multi-Factor Authentication
- Scheduled Service Management
- Task Scheduler
- Encryption Key Store
- Software/Solution Marketplace
- Pre-Build VM Images

Services Compute

- Stateless Compute
- Distributed Compute
- Scheduled Compute Jobs
- Virtual App Streaming

Integration

- Simple Queuing
- B2B Integration
- Hybrid Connections
- Pub/Sub Queuing

Media & CDN

- Live & OD Media Streaming
- Content Delivery Network (CDN)

Web and Mobile

- Web Apps Infrastructure
- API App Infrastructure
- Mobile Backends
- Business Process Automation
- API Management
- Push Notifications

Developer Services

- Development Tools
- Software Development Kits
- Software Lifecycle Management
- Application Instrumentation

Data

- Relational SQL Database
- Data Warehouse
- Document Database Service
- Distributed In-Memory Cache
- Search
- Simple Key/Value Store

Analytics & IoT

- Big Data Analytics
- Predictive Analytics
- Data Stream Analytics
- Big Data Storage
- Data Pipelines
- Device Data Collection
- Data Source Management
- IoT Device Management
- Mobile Analytics

Hybrid Operations

- Directory Health Monitoring
- Privileged Identity Management
- Domain Join & Policy Management
- Server Data Backup
- Operational Analytics
- Bulk Data Import And Export
- Disaster Recovery
- Hybrid/Intelligent Data Backup

Infrastructure Services

OS/Server Compute

- Virtual Servers
- Containers

Storage

- Disk based Object/File Storage
- Shared Storage
- SSD based Object/File Storage

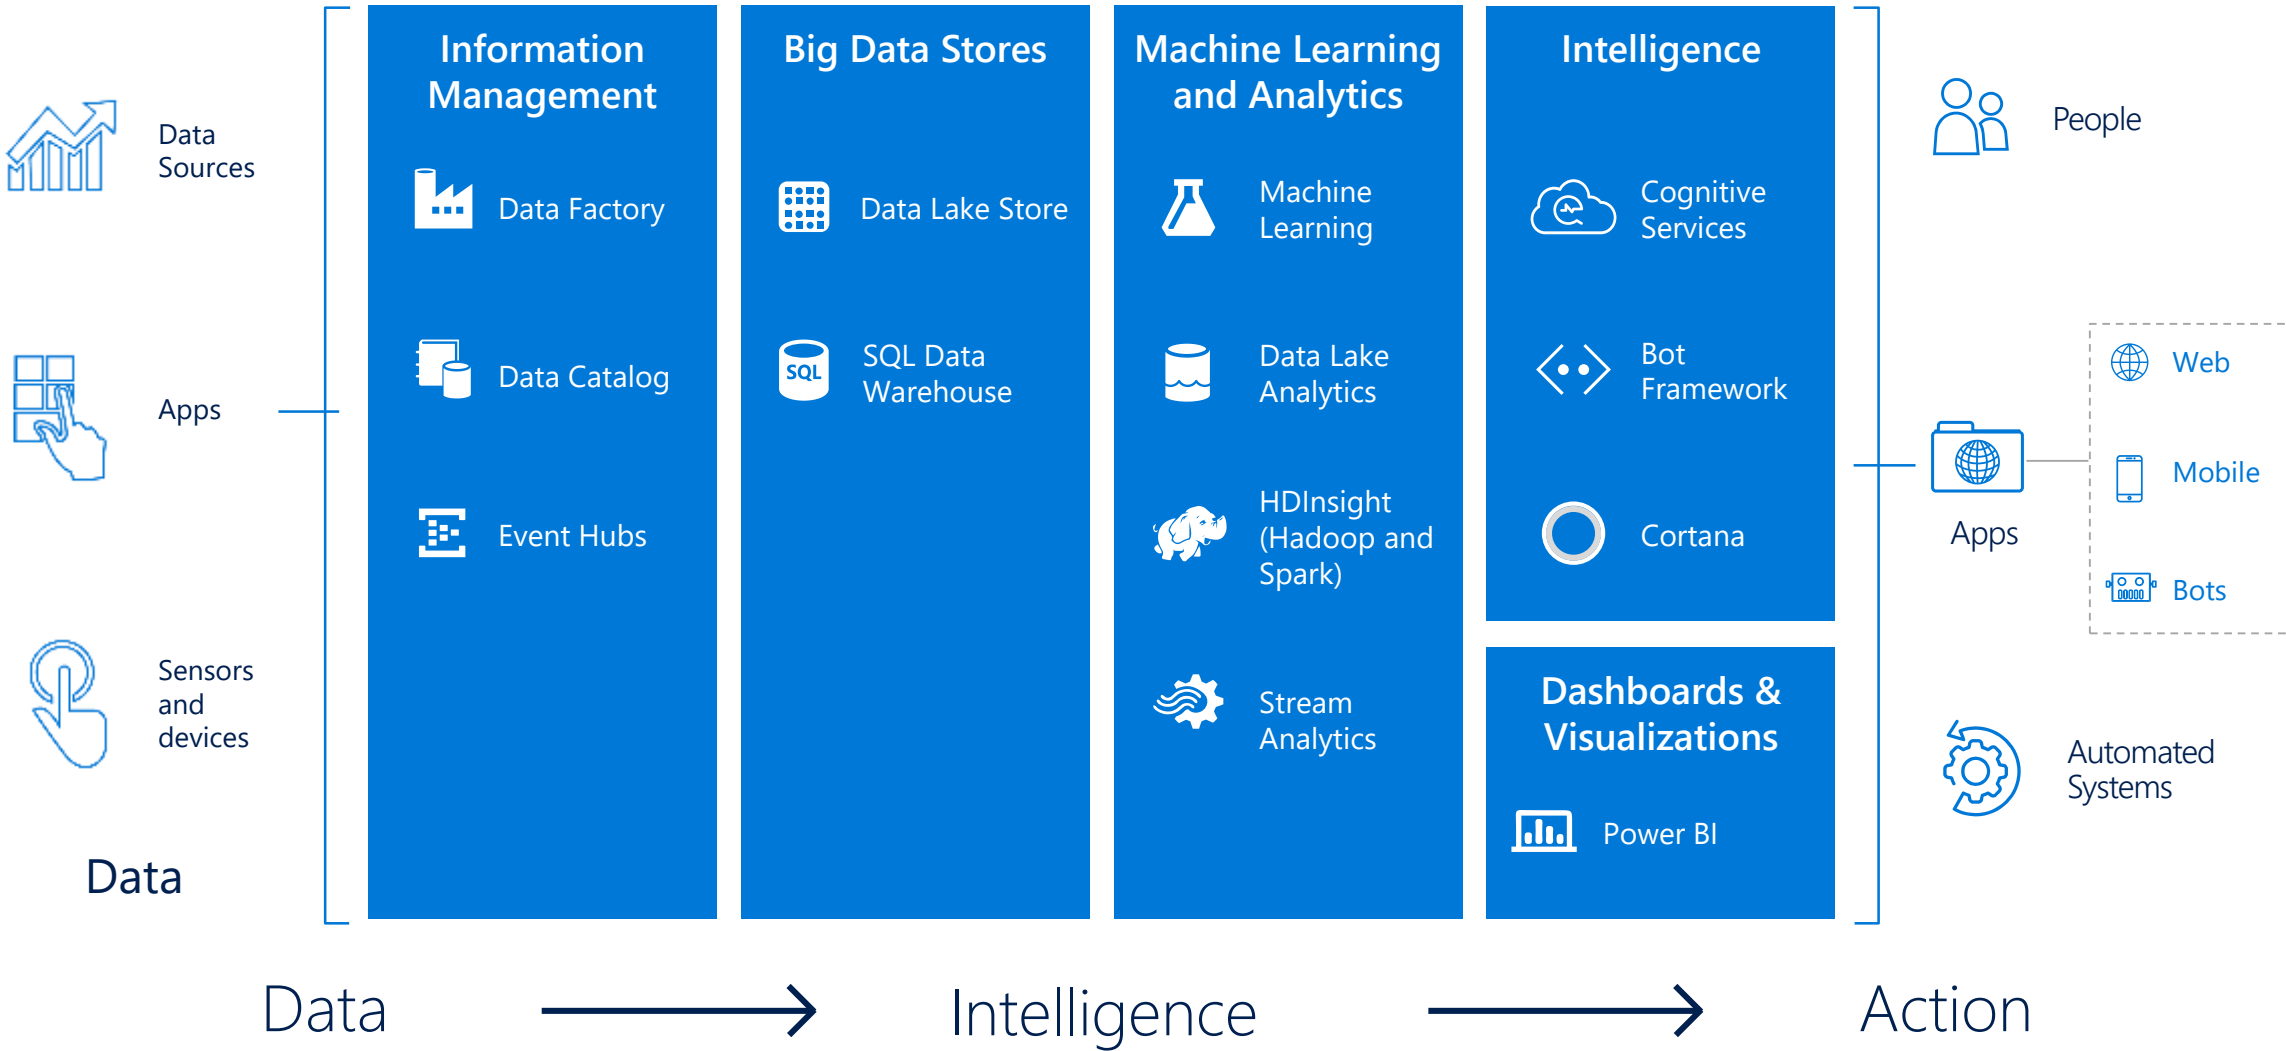
Networking

- Virtual Network
- VM Load Balancer
- DNS
- Direct Network Connections
- Traffic Distribution
- VPN Gateway
- HTTP Load Balancer

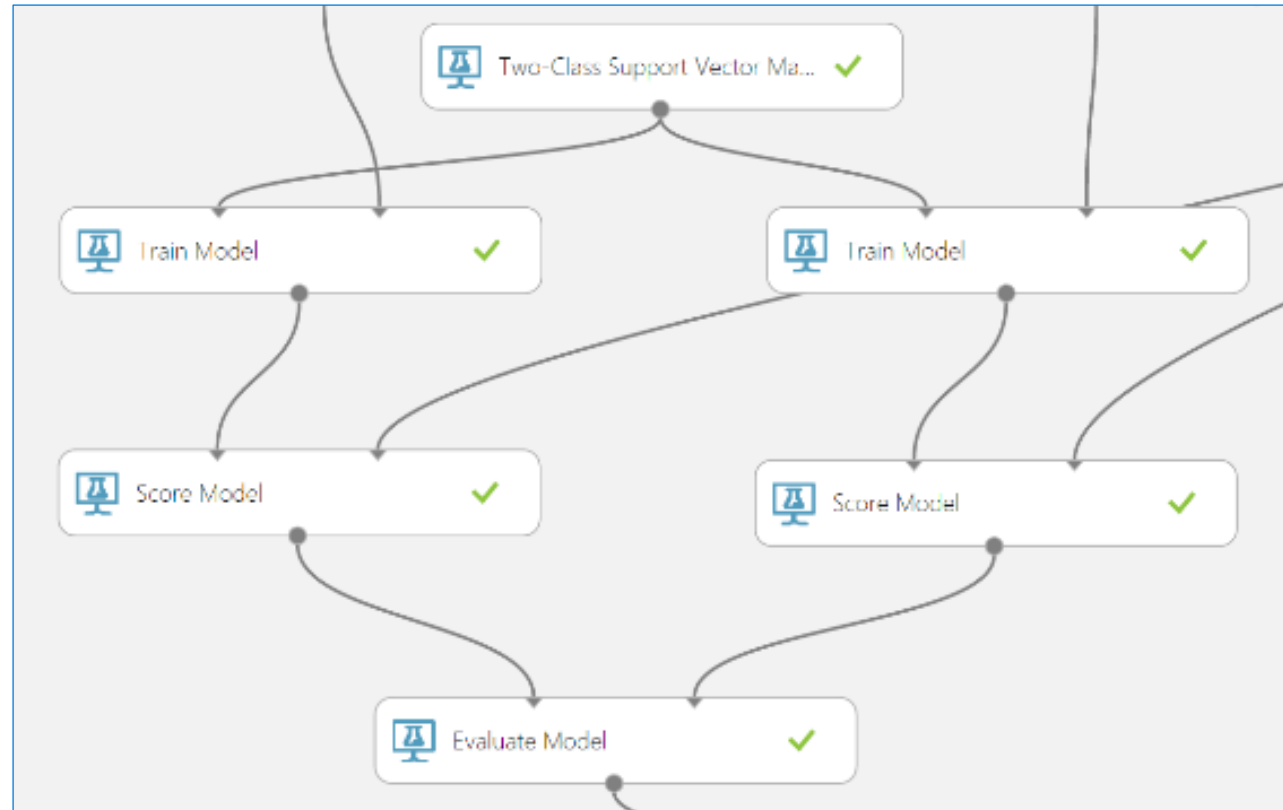
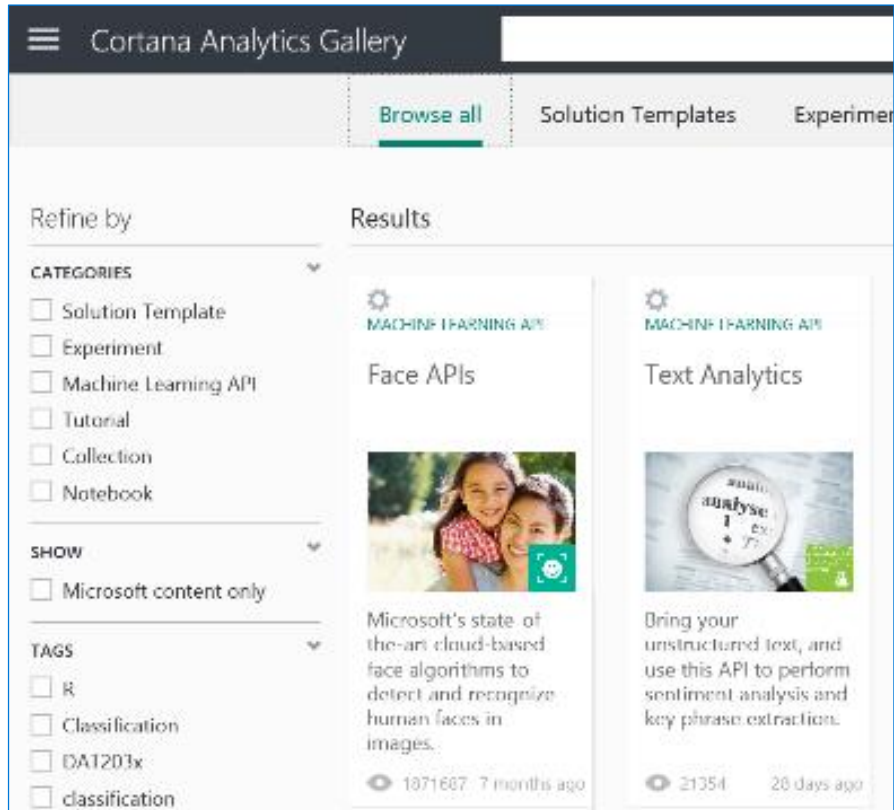
Datacenter Infrastructure (28 Regions, 22 Online)



Transform data into intelligent action



Easily build, deploy, and share predictive analytics solutions



- Simple, scalable, cutting edge. A fully managed cloud service that enables you to easily build, deploy, and share predictive analytics solutions.
- Deploy in minutes. Azure Machine Learning means business. You can deploy your model into production as a web service that can be called from any device, anywhere and that can use any data source.
- Publish, share, monetize. Share your solution with the world in the Gallery or on the Azure Marketplace.

Azure Stack

Microsoft Azure Stack

Power of Azure in your datacenter

Try it now >

Available now in Technical Preview




Why Azure Stack?

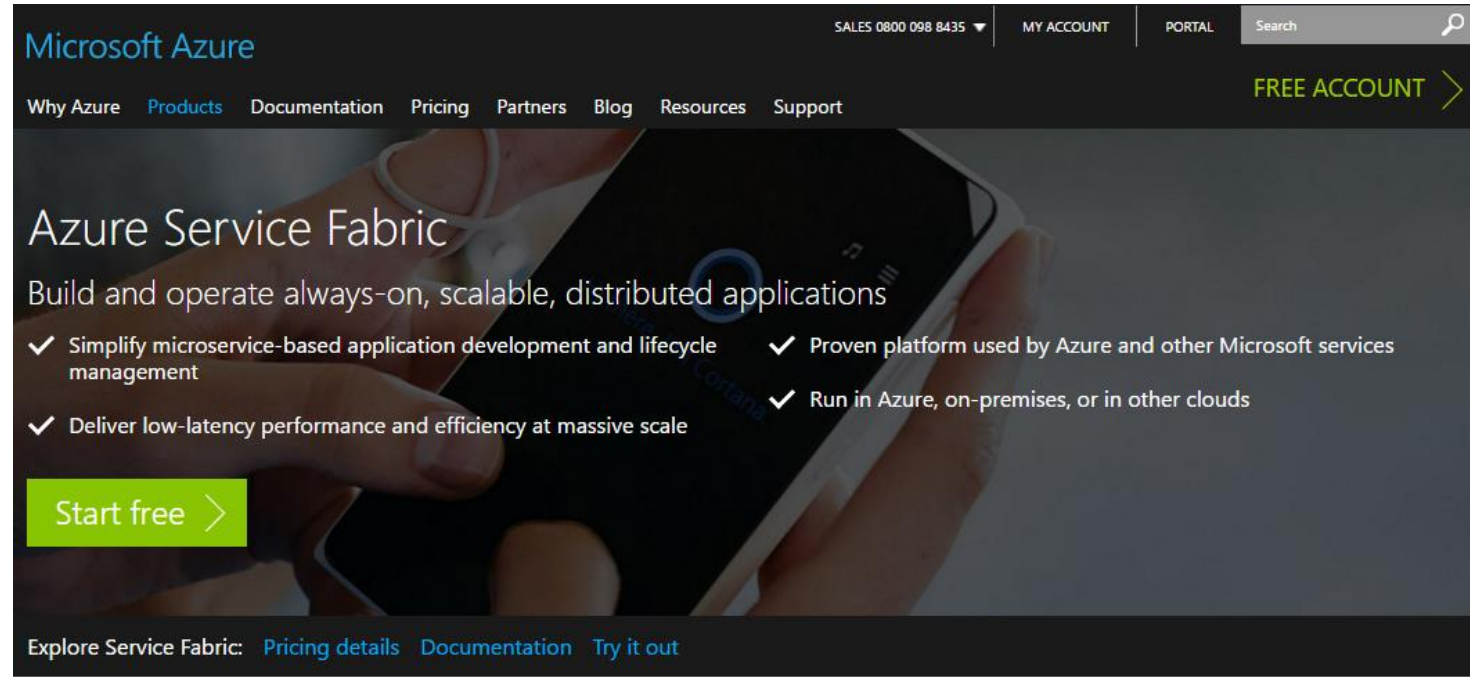
Microsoft Azure Stack is a new hybrid cloud platform product that enables your organization to deliver Azure services from your own datacenter to help you achieve more. Get the power of cloud services, yet maintain control of your datacenter for true hybrid cloud agility. You decide where to keep your data and applications—in your own datacenter or with a hosting service provider. Easily access public cloud resources to scale at busy times of the year, for dev-test, or whenever you need them. Only Microsoft builds and runs its own hyper-scale datacenters and delivers that proven innovation to your datacenter.

[Read the blog to learn more about our Azure Stack vision](#)

[Read a white paper about Azure Stack fundamentals](#)



Azure Service Fabric



The screenshot shows the Microsoft Azure website for Azure Service Fabric. At the top, there is a navigation bar with the Microsoft Azure logo, a phone number (SALES 0800 098 8435), links for MY ACCOUNT and PORTAL, a search bar, and a prominent green button for FREE ACCOUNT. Below the navigation bar, there is a main heading "Azure Service Fabric" and a sub-heading "Build and operate always-on, scalable, distributed applications". Three bullet points with checkmarks describe the benefits: "Simplify microservice-based application development and lifecycle management", "Deliver low-latency performance and efficiency at massive scale", "Proven platform used by Azure and other Microsoft services", and "Run in Azure, on-premises, or in other clouds". A green "Start free" button is positioned below the first two bullet points. At the bottom of the main content area, there are links for "Explore Service Fabric: Pricing details", "Documentation", and "Try it out".


Simplify microservice-based application development and lifecycle management

- **Fast time to market:** Service Fabric lets developers focus on building features that add business value to their application, without the overhead of designing and writing additional code to deal with issues of reliability, scalability, or latency in the underlying infrastructure.
- **Choose your architecture:** Build stateless or stateful microservices—an architectural approach where complex applications are composed of small, independently versioned services—to power the most complex, low-latency, data-intensive scenarios and scale them into the cloud.







Predictive maintenance IoT template

Aircraft map



Sensor history



Remaining Useful Life (RUL)
IN CYCLES

173
ENGINE #1

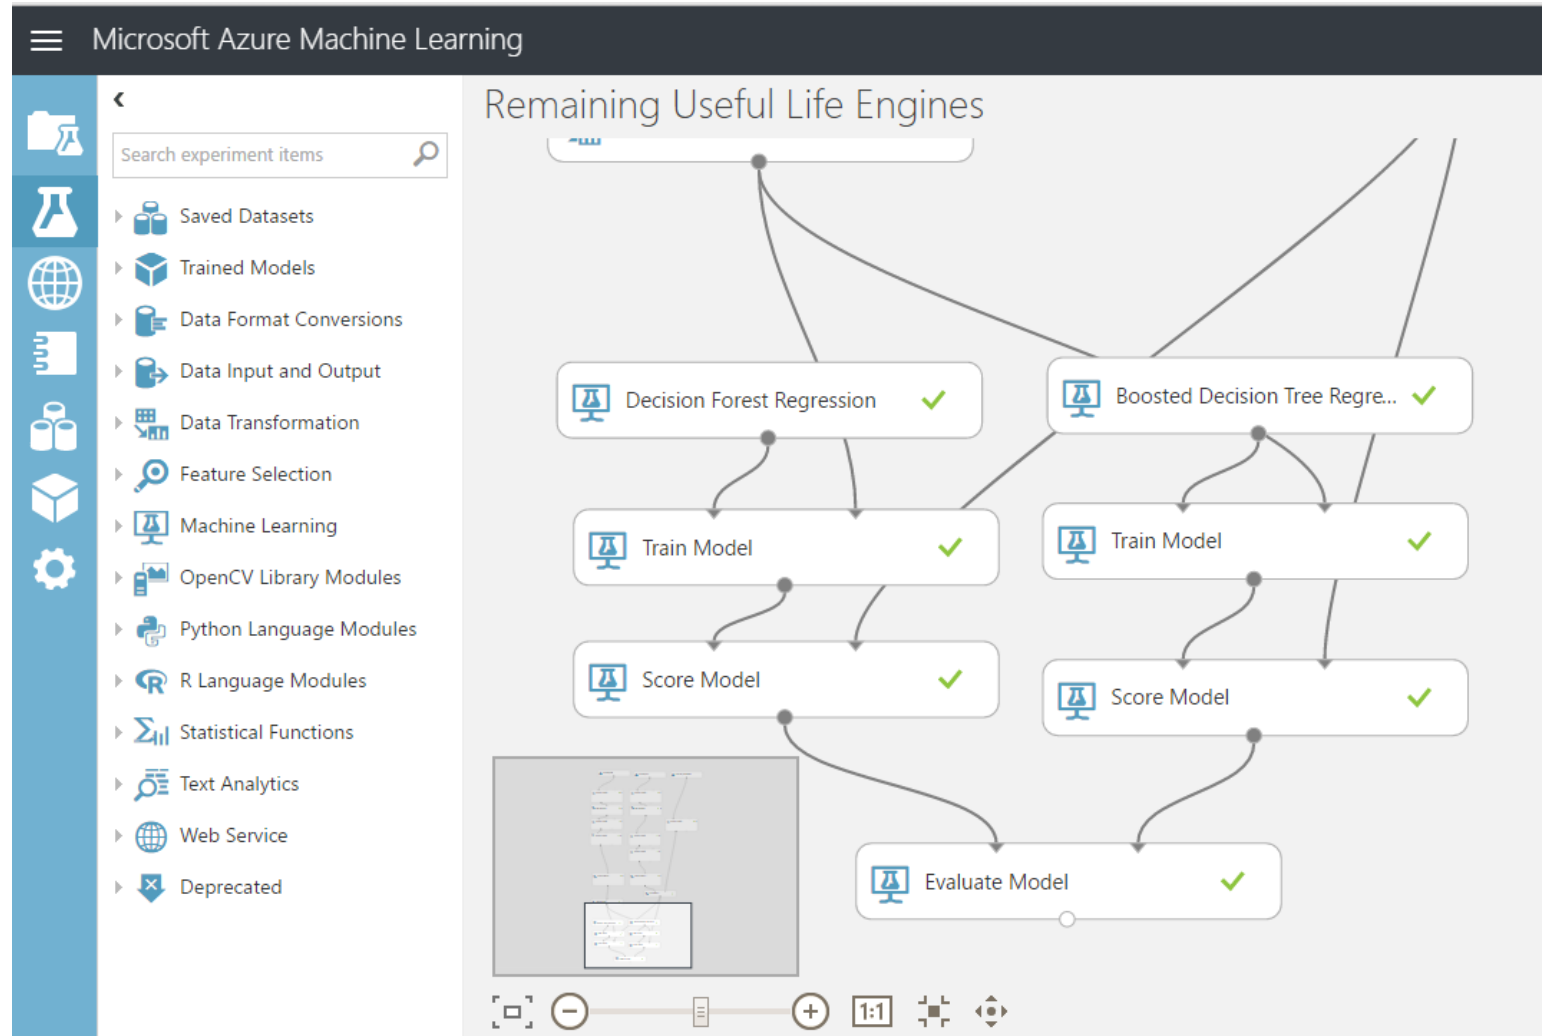
159 ⚠
ENGINE #2

Cycles
#

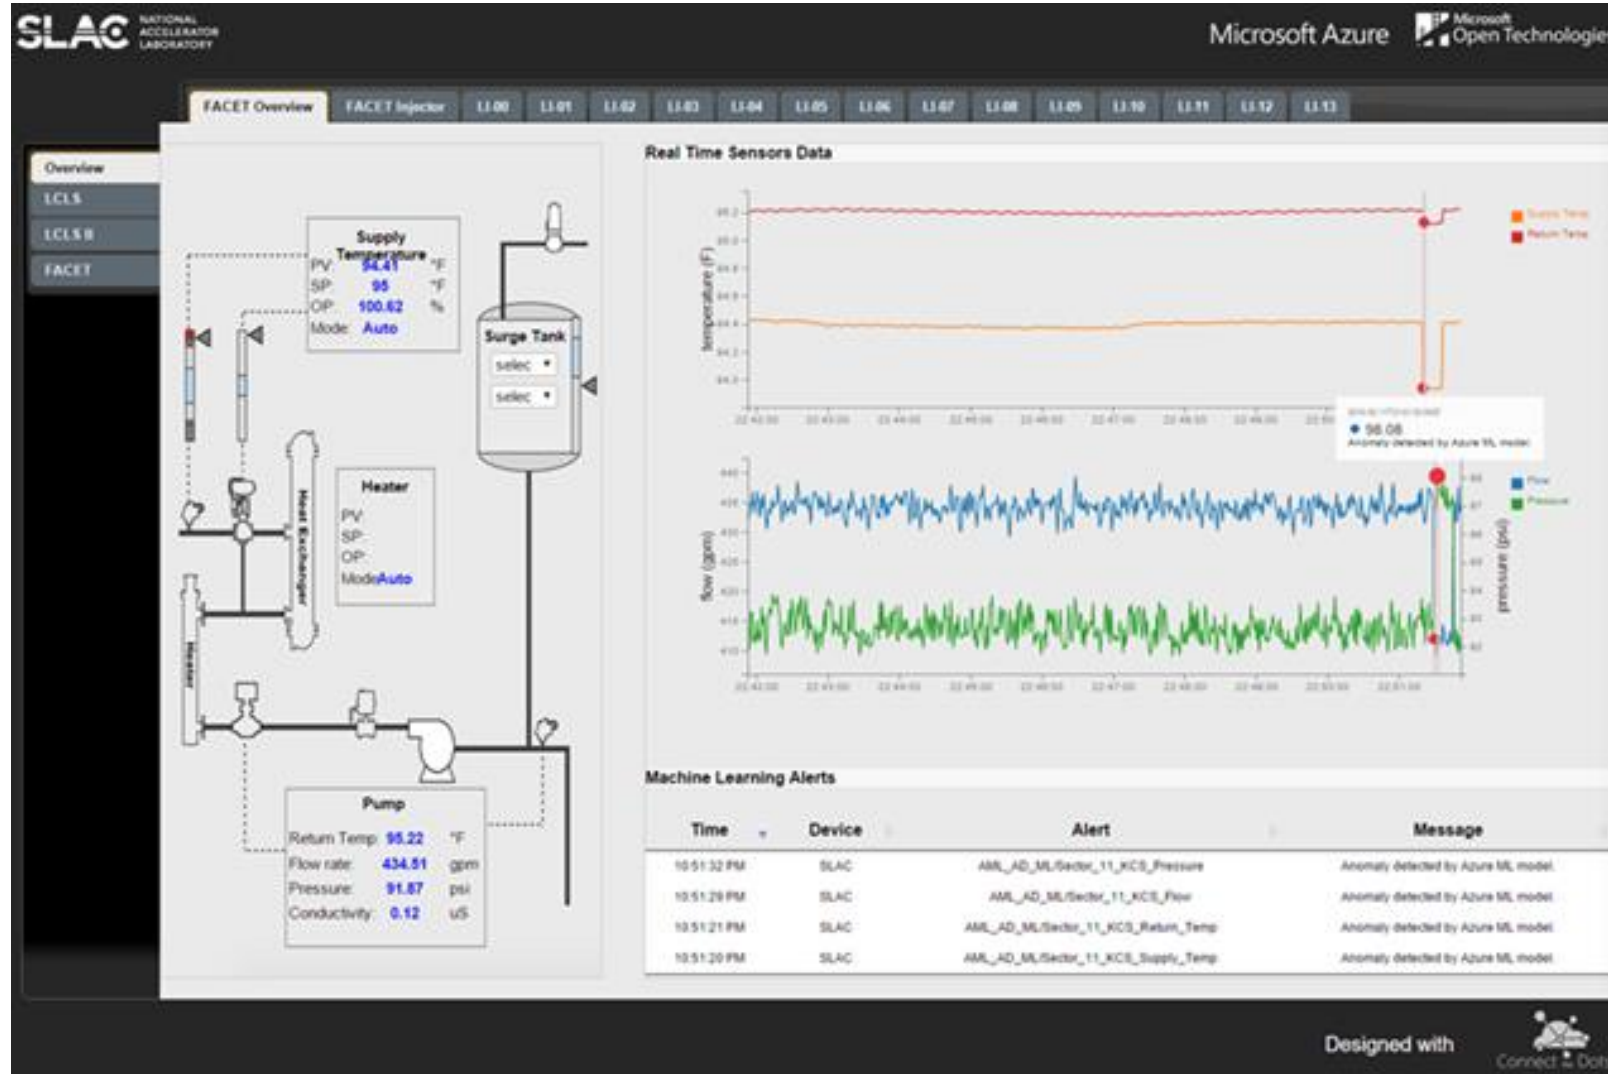
54
ENGINE #1

55
ENGINE #2

Predictive maintenance machine learning



SLAC: Azure ML & IoT



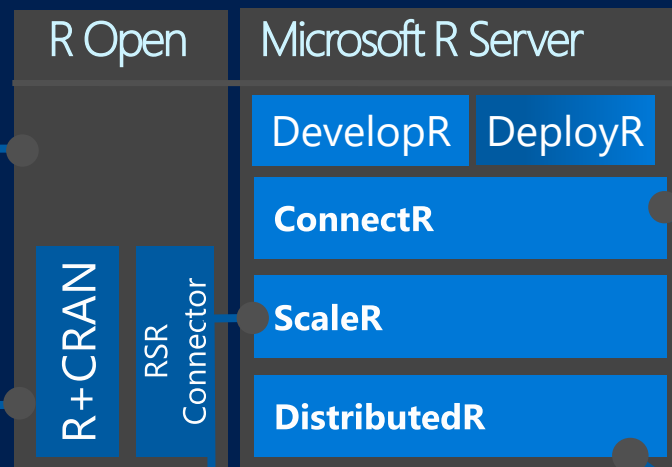
The Microsoft R Server Platform

R+CRAN

- Open source R interpreter
 - R 3.1.2
- Freely-available huge range of R algorithms
- Algorithms callable by RevoR
- Embeddable in R scripts
- 100% Compatible with existing R scripts, functions and packages

RevoR

- Performance enhanced R interpreter
- Based on open source R
- Adds high-performance math library to speed up linear algebra functions



ScaleR

- Ready-to-Use high-performance big data big analytics
- Fully-parallelized analytics
- Data prep & data distillation
- Descriptive statistics & statistical tests
- Range of predictive functions
- User tools for distributing customized R algorithms across nodes
- Wide data sets supported – thousands of variables

ConnectR

- High-speed & direct connectors

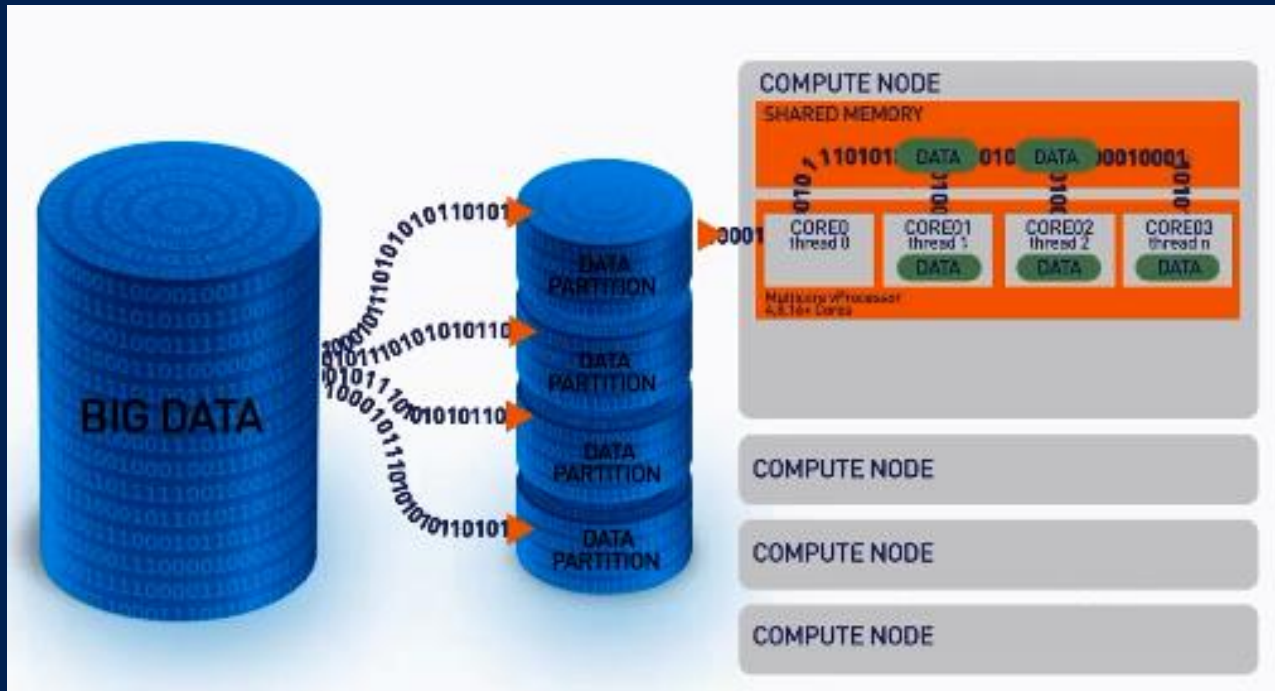
Available for:

- High-performance XDF
- SAS, SPSS, delimited & fixed format text data files
- Hadoop HDFS (text & XDF)
- Teradata Database & Aster
- EDWs and ADWs
- ODBC

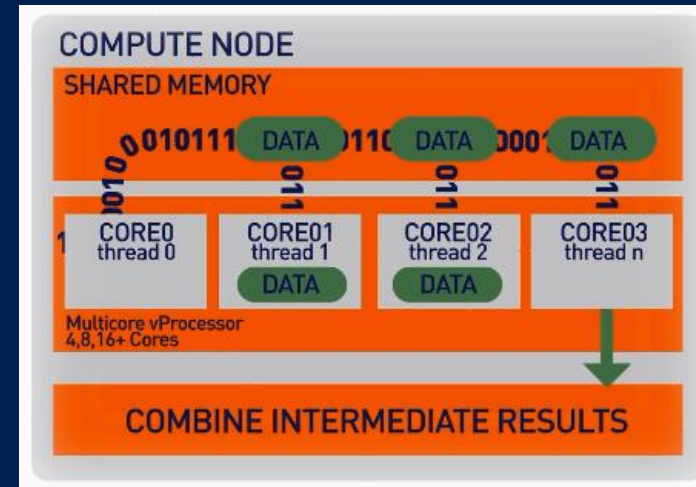
DistributedR

- Distributed computing framework
- Delivers cross-platform portability

ScaleR – Parallel + “Big Data”



Our ScaleR algorithms work inside multiple cores / nodes **in parallel** at high speed



Interim results are collected and combined analytically to produce the output on the entire data set

Stream data in to RAM in blocks. “Big Data” can be any data size. We handle Megabytes to Gigabytes to Terabytes...

XDF file format is optimised to work with the ScaleR library and significantly speeds up iterative algorithm processing.

R Server for Azure HDInsight



Familiarity of R algorithms


+

Scalability of Hadoop + Spark

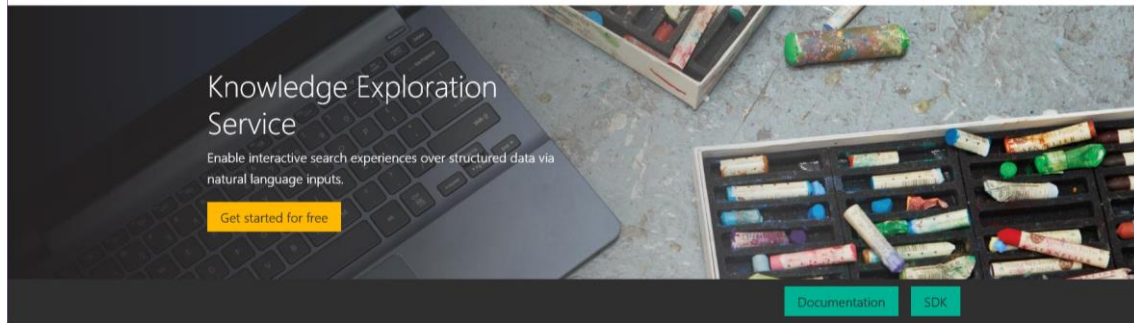
=

More Accurate Predictions

CNTK open-source deep-learning toolkit



Cognitive Services



Knowledge Exploration Service
Enable interactive search experiences over structured data via natural language inputs.

[Get started for free](#)

[Documentation](#) [SDK](#)

Natural Language Understanding

To interpret natural language queries as structured query expressions.



Query Auto-Completion

To reduce user effort and help with discovery of rich capabilities.



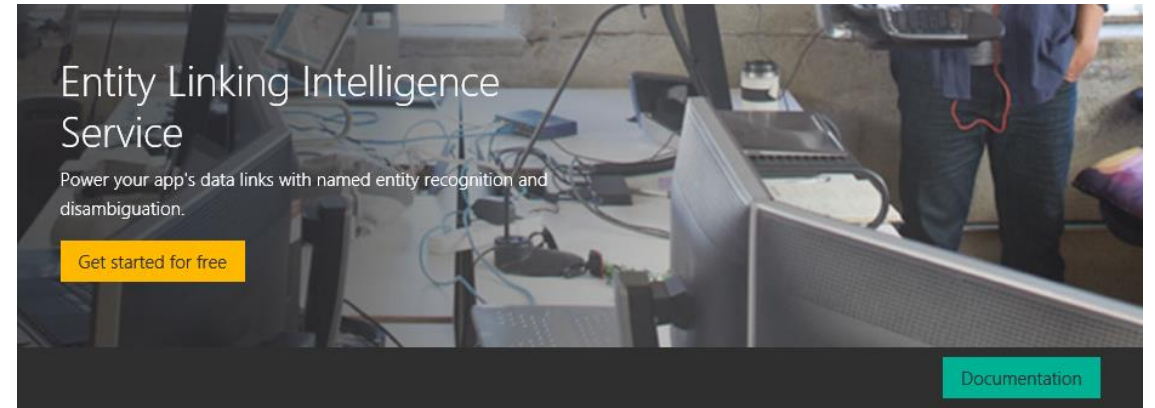
Structured Query Evaluation

To efficiently retrieve detailed information about matching objects.



Attribute Histograms

To enable rich visualizations and interactive faceted experiences.



Entity Linking Intelligence Service
Power your app's data links with named entity recognition and disambiguation.

[Get started for free](#)

[Documentation](#)

Entity Linking

Sometimes in different contexts, a word might be used as a named entity, a verb, or another word form within a given sentence. For example, in the case where "times" is a named entity, it still may refer to two separately distinguishable entities, such as "The New York Times" or "Times Square". Given a specific paragraph of text within a document, the Entity Linking Intelligence Service will recognize and identify each separate entity based on the context.

For months, the four scientific instruments at the heart of the James Webb Space Telescope have been sealed in what looks like a huge pressure cooker. It's a test chamber that simulates the grueling operating conditions they will face after Webb is launched into orbit in 2018. But in fact, "pressure cooker" is an apt metaphor for the whole project. The infrared Webb observatory is the biggest, most complex, and most expensive science mission that NASA has ever attempted. Like that of its predecessor, the Hubble Space Telescope, Webb's construction has been plagued by redesigns, schedule slips, and cost overruns that have strained relationships with contractors, international partners, and supporters in the U.S. Congress. Lately the project has largely stuck to its schedule and its \$8 billion budget. But plenty could still go wrong, and the stakes are high: Both the future of space-based astronomy and NASA's ability to build complex science missions depend on its success.

[Highlight Content](#)

[JSON](#)



For months, the four scientific instruments at the heart of the [James Webb Space Telescope](#) have been sealed in what looks like a huge pressure cooker. It's a test chamber that simulates the grueling operating conditions they will face after [Webb](#) is launched into orbit in 2018. But in fact, "pressure cooker" is an apt metaphor for the whole project. The infrared [Webb](#) observatory is the biggest, most complex, and most expensive science mission that [NASA](#) has ever attempted. Like that of its predecessor, the [Hubble Space Telescope](#), [Webb's](#) construction has been plagued by redesigns, schedule slips, and cost overruns that have strained relationships with contractors, international partners, and supporters in the [U.S. Congress](#). Lately the project has largely stuck to its schedule and its \$8 billion budget. But plenty could still go wrong, and the stakes are high: Both the future of space-based astronomy and [NASA's](#) ability to build complex science missions depend on its success.

Q&A

Microsoft Academic Graph

<http://aka.ms/academicgraph>

aka.ms/AcademicAPI

Author (> 40M)


Publication (> 100M)

Event (> 46K)

Venue (> 23K)



Citations (billions)




Institution (20K)

Field of Study (> 50K)

Knowledge API

Microsoft
Cognitive Services

My account 

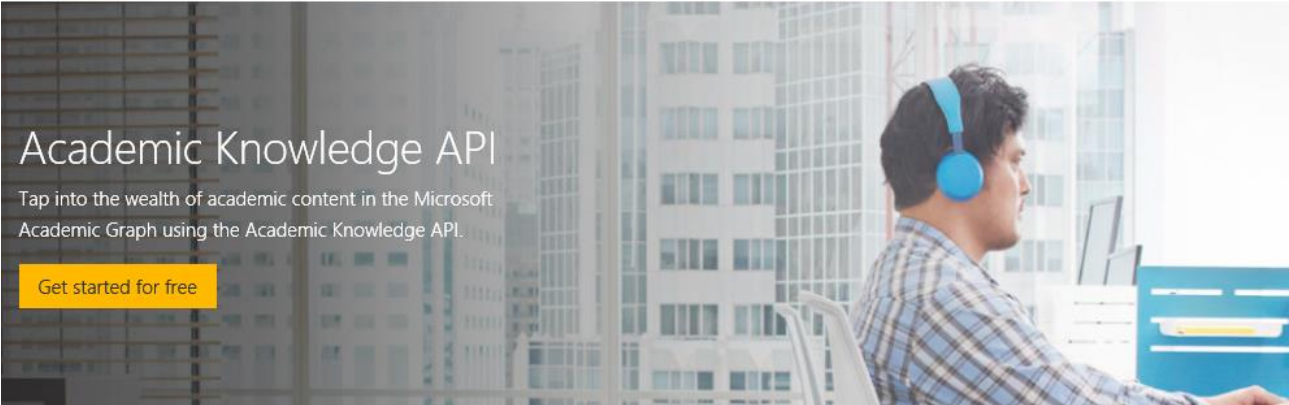
Home APIs Applications Developers Pricing

Academic Knowledge API

Tap into the wealth of academic content in the Microsoft Academic Graph using the Academic Knowledge API.

Get started for free

Documentation API Reference



Interpret

Interprets a natural language user query string. Returns annotated interpretations to enable rich search-box auto-completion experiences that anticipate what the user is typing.

Evaluate

Evaluates a query expression and returns Academic Knowledge entity results.



Author

aka.ms/AcademicAPI



© 2016 Microsoft Corporation. All rights reserved.

