

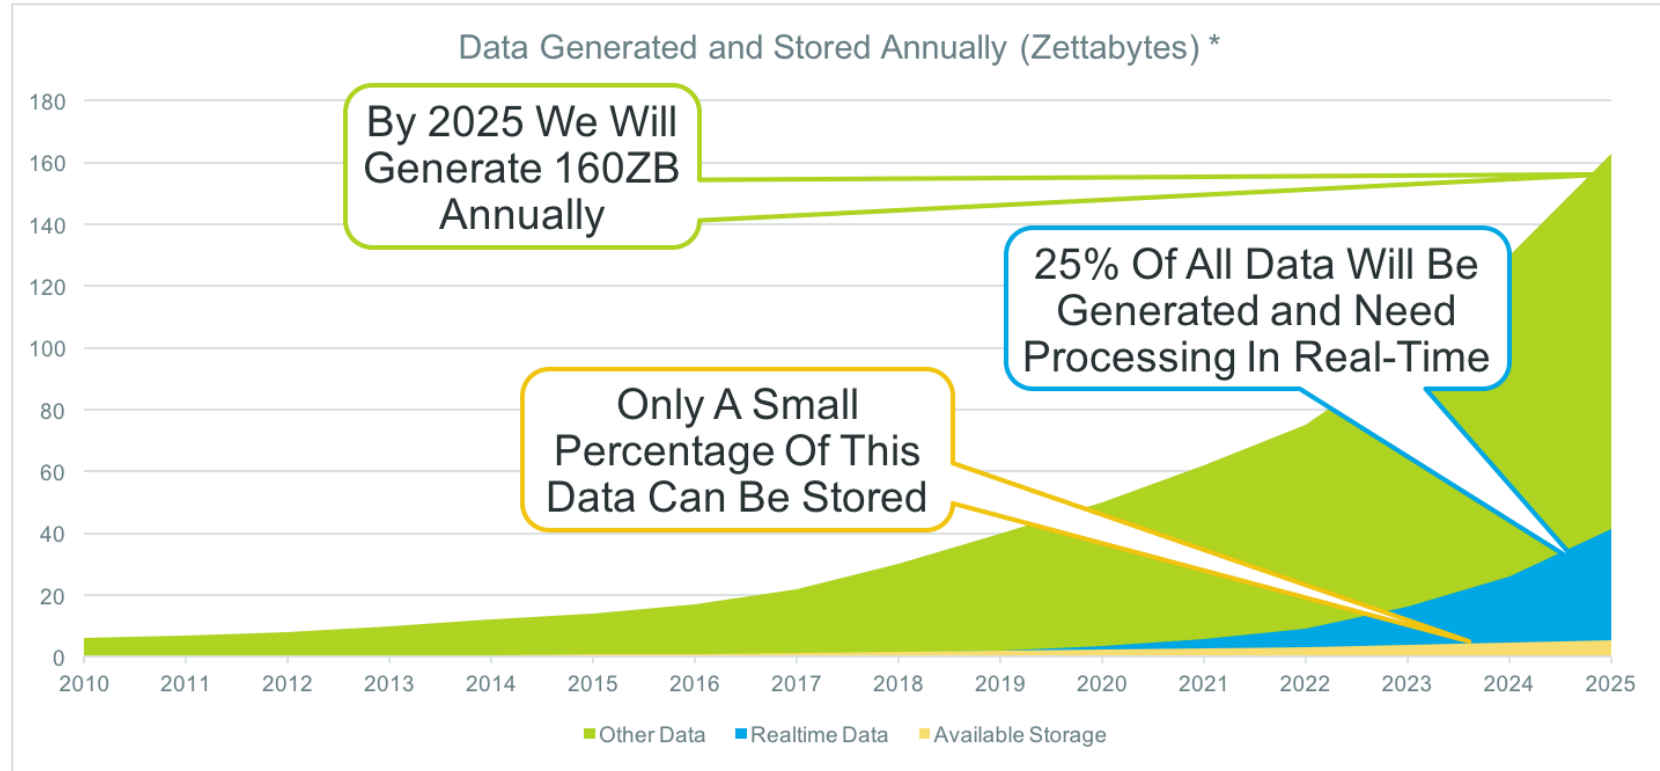
Real-Time Actionable Insights with IBM Streams

—
Roger Rea
IBM Streams Senior Offering Manager
rrea@us.ibm.com

Data and AI Forum 2019
Session 309



Real-time processing needs are growing



Continuous intelligence

- What if you could analyze data as it's created?
- What if you could visualize your business?
- What if you could better predict your customers' needs?
- What if you could gain insights from unstructured data like audio, text or video?
- What if could automate immediate actions?
- What if you always knew where your assets were, and where they would be?
- What if you could update machine learning models continuously?

And do it all in real time?

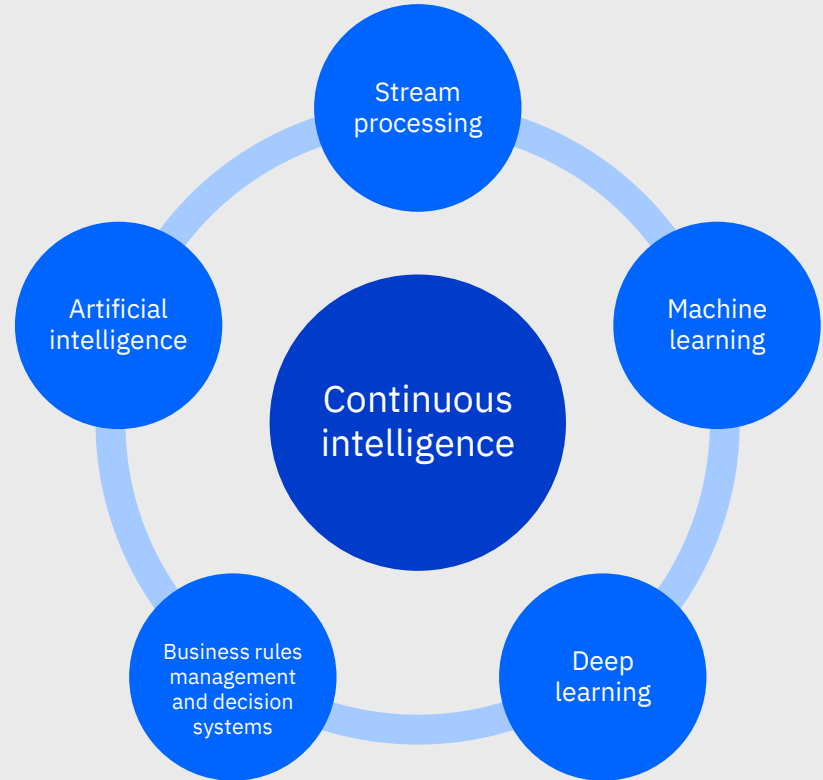


“Continuous intelligence is a design pattern in which real-time analytics are integrated into a business operation, processing current and historical data to prescribe actions in response to business moments and other events.”

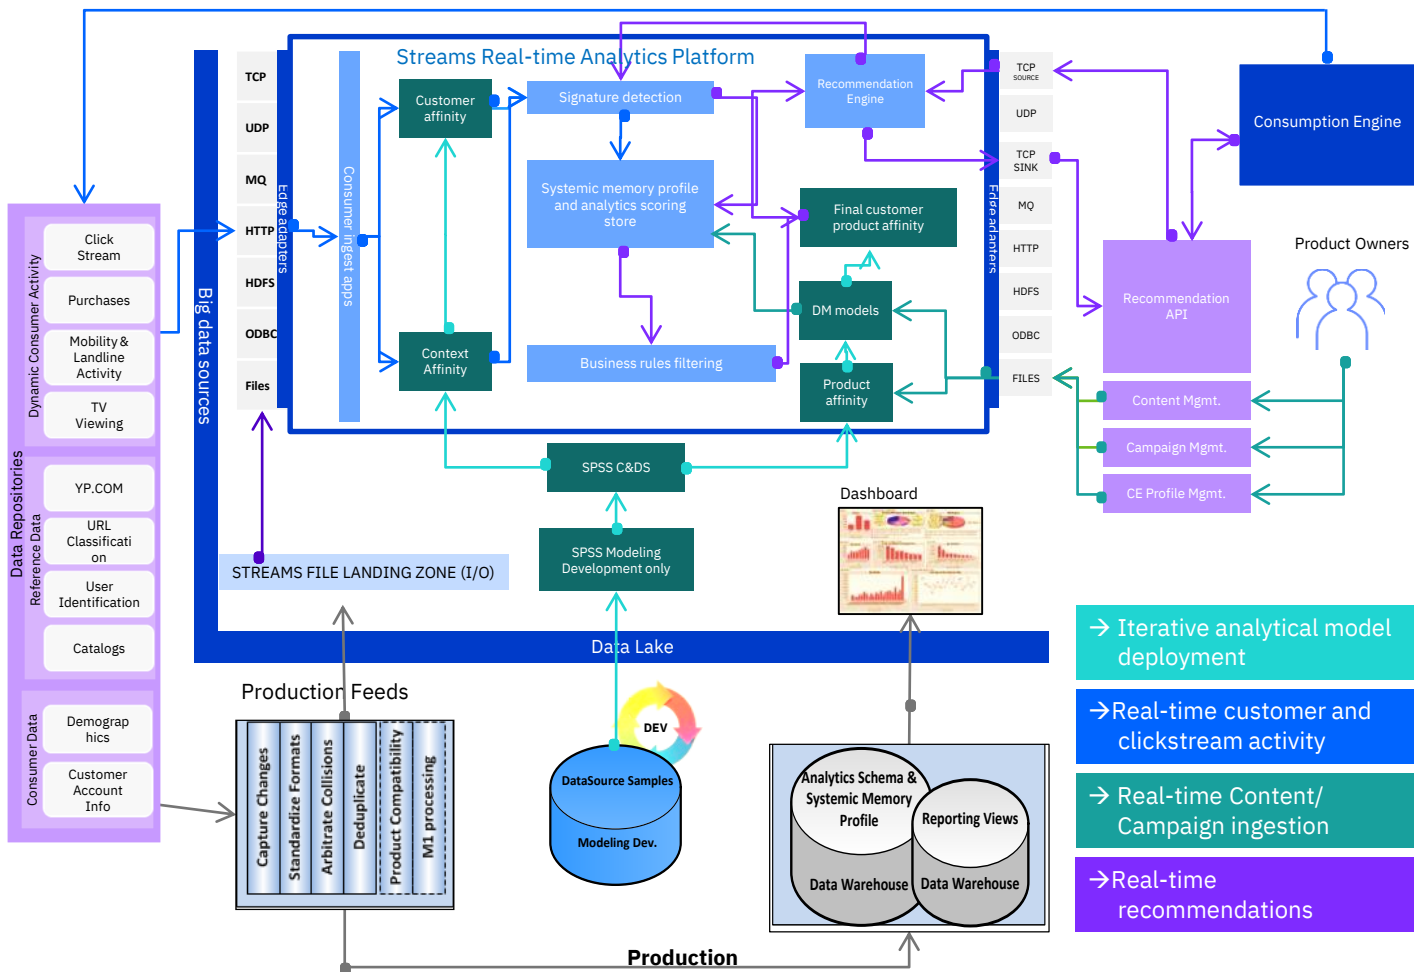
Gartner: Innovation Insight for Continuous Intelligence

Continuous intelligence

- Engage data from inside and outside of applications or business operations
- Enable fast-paced digital business decisions and process optimization
- Leverage AI, ML, data analytics, real-time analytics and streaming event data to deliver business optimized solutions
- Ensure you can take advantage of the “perfect storm” of rising supply and demand for real-time situation awareness and responsiveness



Continuous intelligence example running over 1700 models in production



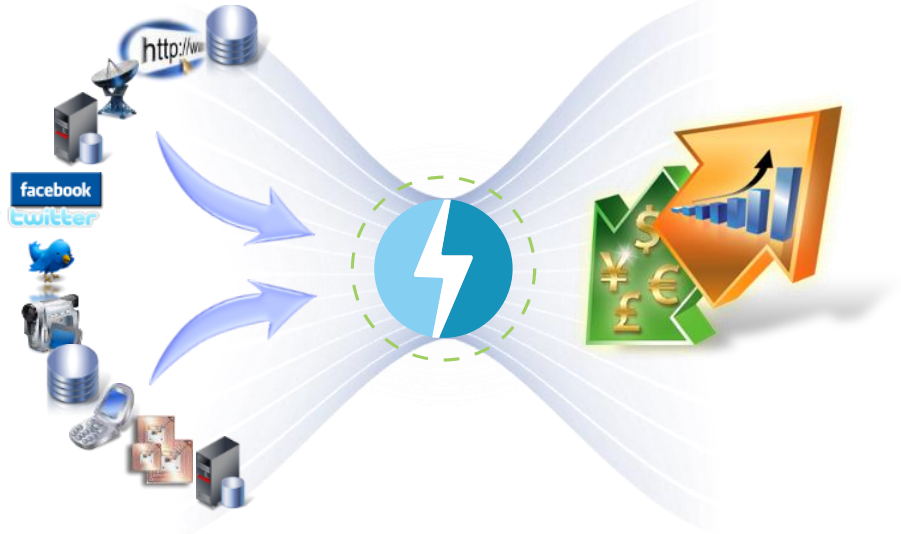
IBM Streams to act on all your data in real time

Market leader in streaming analytics

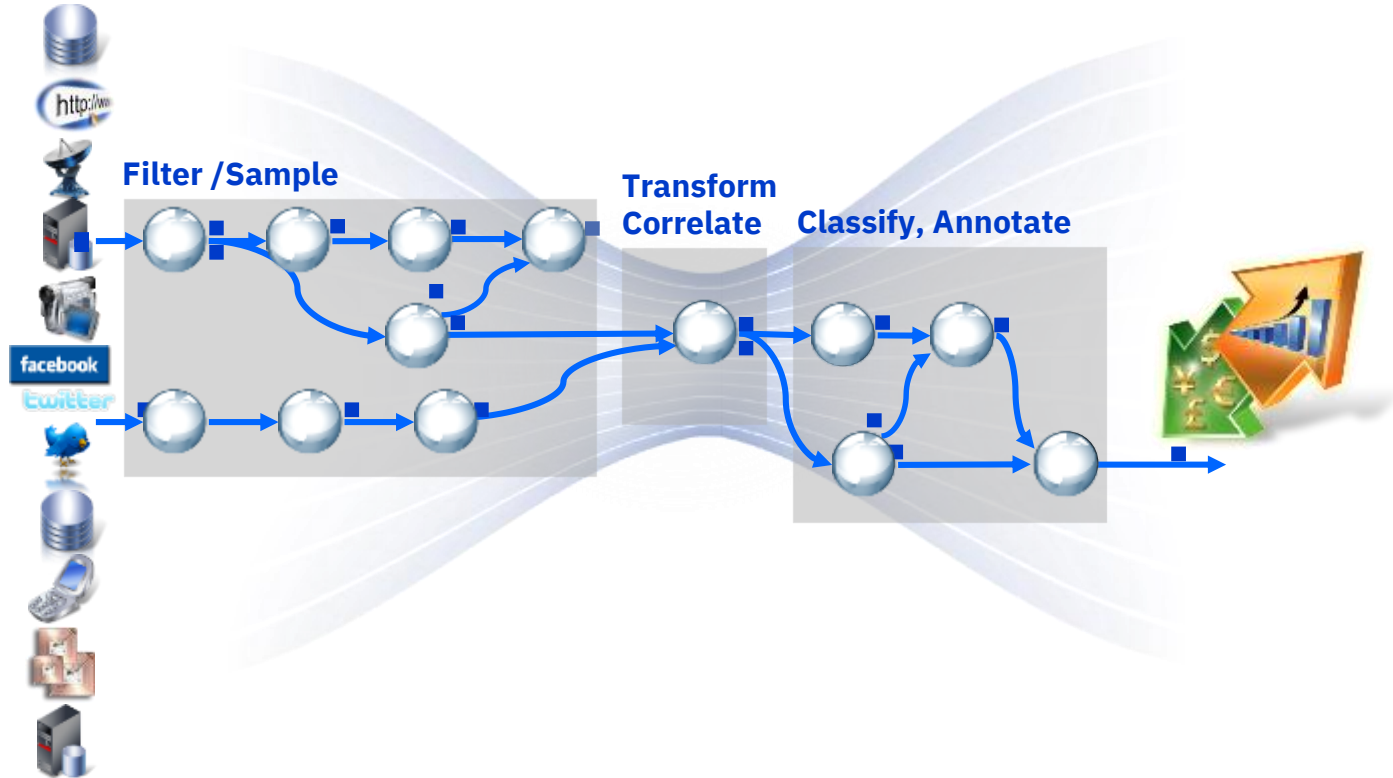
- Machine Learning
- Model Scoring
- Geospatial
- Video/Image
- Text, Speech to Text, Predictive, Descriptive

Enterprise Ready: included with Cloud Pak for Data

- Visual development
- Web console
- Management
- Enterprise connectors like JSON, JMS, MQ, MQTT



Streams to act on all your data in real time



IBM Streams offerings

Streams v4.3.1 July 2019

Streams runtime

Baremetal/VM

Available 1Q19

Streams for IBM Cloud Pak
for Data

Streams runtime

Containers 

Free [Lite Plan](#) – 50
hours a month

Streaming Analytics

Streams runtime

Containers 

SaaS

Common runtime: develop anywhere and deployment everywhere

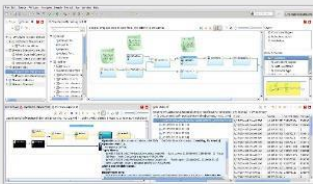
Moved from VM to
Containers in 2018



IBM Streams development options

Streams Developer Edition, Streams Quick Start Edition

- Dedicated Streams development IDE and tools
- Local machine

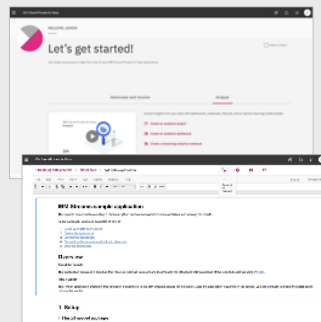


Quick Start free
for non-production

IBM Cloud Pak for Data

- Integrated data and AI project experience
- Containers*

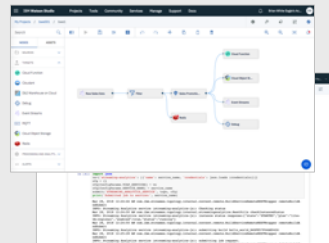
Python Notebooks



*Moved from VM to Containers in 2018

Watson Studio Streams Flows in IBM Cloud

- Integrated data and AI project experience
- SaaS



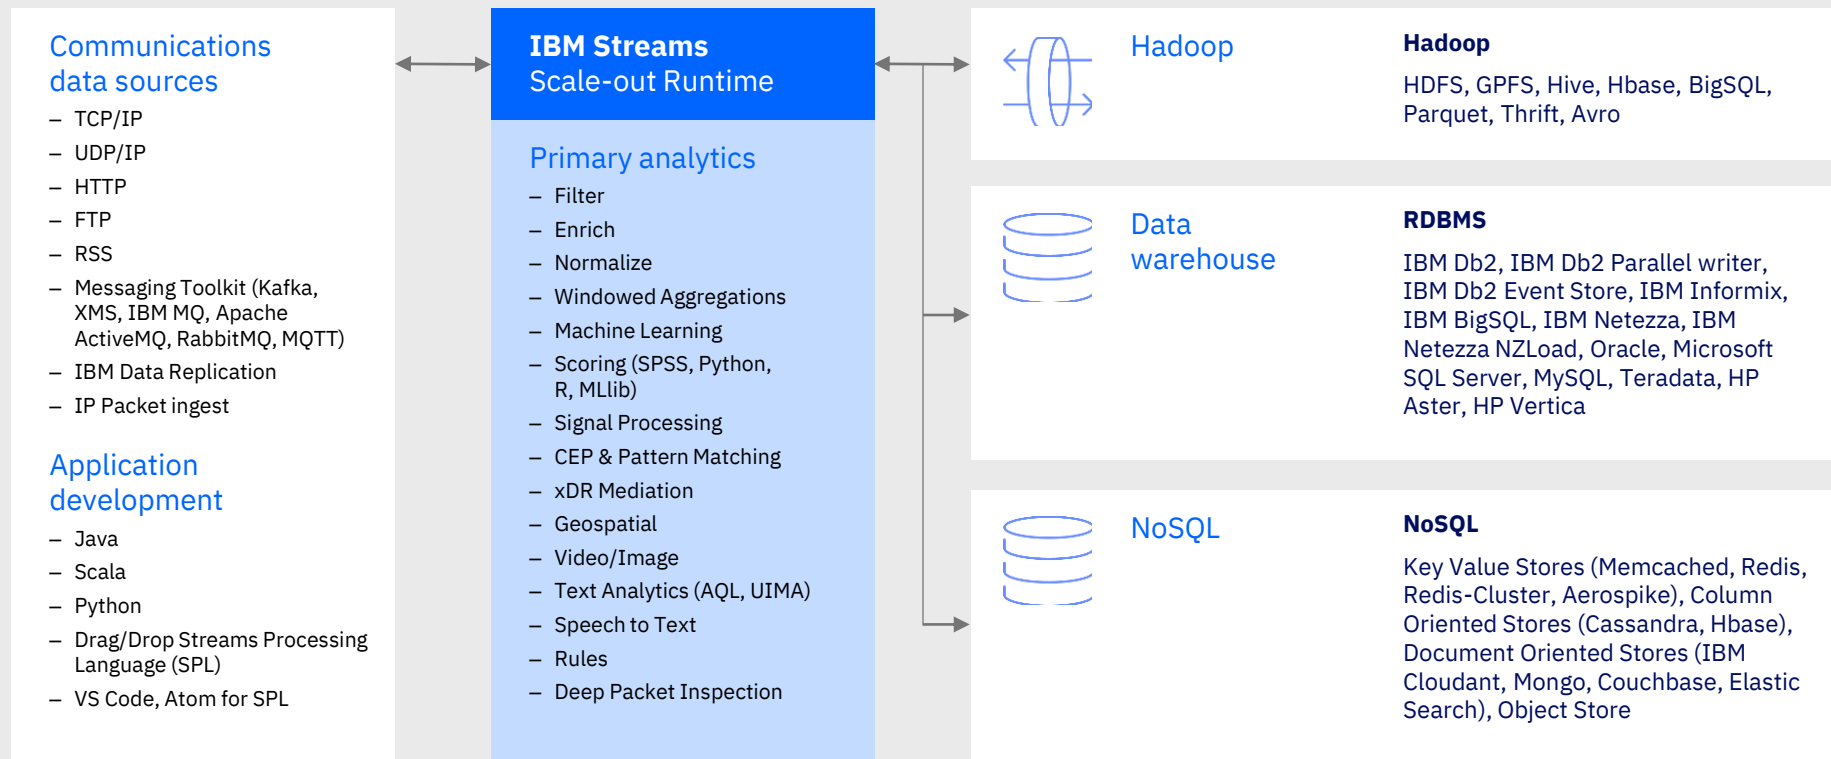
Available 4Q18

Plus IBM Runner for Apache Beam, Java, Scala, Python and beta plug-ins for VSCode and Atom

IBM Streams at a glance

Out of the box with over 200 operators with 1300 functions

And even more at
github.com/IBMStreams



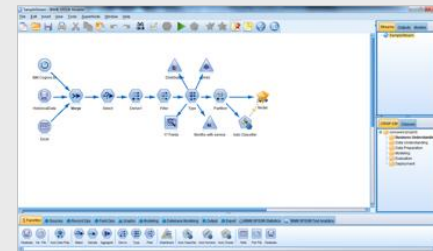
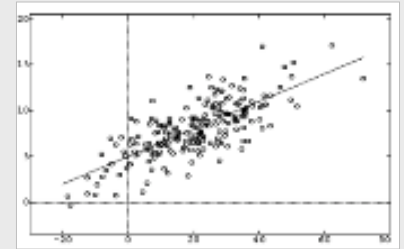
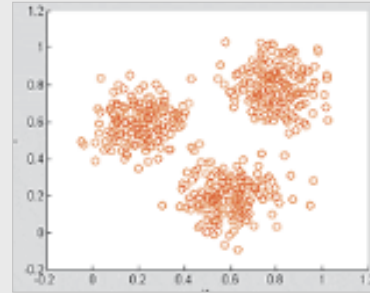
Machine learning and real time analytics with IBM streams

Many different approaches to Machine Learning:

- Mechanisms: Supervised and Unsupervised
- Algorithms: Decision Trees, Regressions, Classification, Clustering, etc.
- Inputs: Single and multi-variant, rich feature vectors
- Streams has 20 ML algorithms that learn as you go in real time
- Streams scores models created offline from popular tools including IBM SPSS, Watson Machine Learning, SparkMLLib, Python, R & PMML libraries
- Native ML and Model scoring can all be integrated within a Streams application. With this approach, real time scores can be generated on the incoming data

“The science of getting computers to act without being explicitly programmed”¹

Use machine learning models created with Watson Studio to score live data on Streaming Analytics using Watson Studio Streams Flows.



Streams: Use the language of choice

Streams Processing Language (SPL)

- Tailored to stream processing
- High-level, declarative composition language
- Graphical editor support

Create topologies in Java

- Indirect support for Scala

Python topologies and operators

- Integration with Jupyter notebook
- Integration with IBM Watson Studio
- Add Python functions in line with SPL code

Publish/subscribe data exchange

- Between applications written in any language

Streams Processing Language

```
stream<rstring item> Sale = Join(Bid; Ask)
{
    window Bid: sliding, time(30);
           Ask: sliding, count(50);
    param match : Bid.item == Ask.item
           && Bid.price >= Ask.price;
    output Sale: item = Bid.item;
}
```

Java Topology

```
/*
 * Declare a source stream (hw) with String as tuples that sends
 * two tuples "Hello" and "World!", and prints them to output.
 */
Topology = new Topology("HelloWorld");
TStream<String> hw = topology.strings("Hello", "World!");
hw.print();
StreamsContextFactory.getEmbedded().submit(topology).get();
```

Python Topology

```
12 class ECGPatientData:
13     def __init__(self, username, password, sample_r
14         self.username = username
15         self.password = password
16         self.sample_rate = sample_rate
17         self.target_sample_rate = 100
18         self.patient_id=patient_id
19
20     def run(self):
21         ## Create topology
22         topo = Topology("ECGPatientDataViz")
```

Advantage: Scale out with ease

Create logical application flow

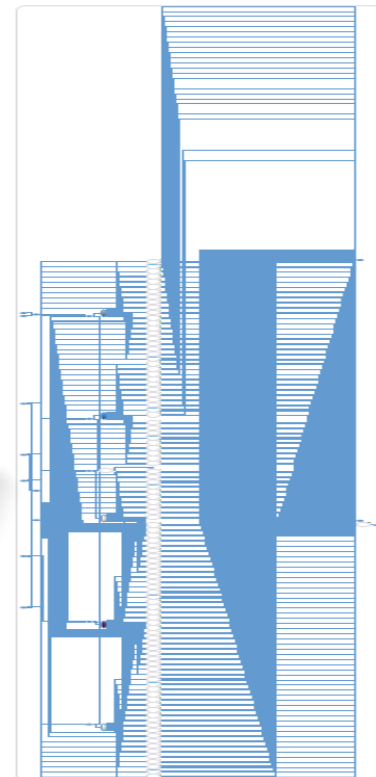
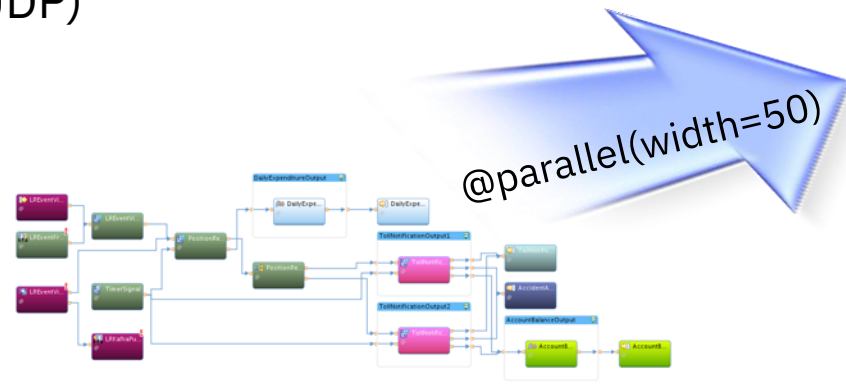
- Without concern for throughput limitations

As needed, add parallel paths

- Based on runtime performance profiling

User-Defined Parallelism (UDP)

- Simple change in code or graphical editor
 - @parallel annotation



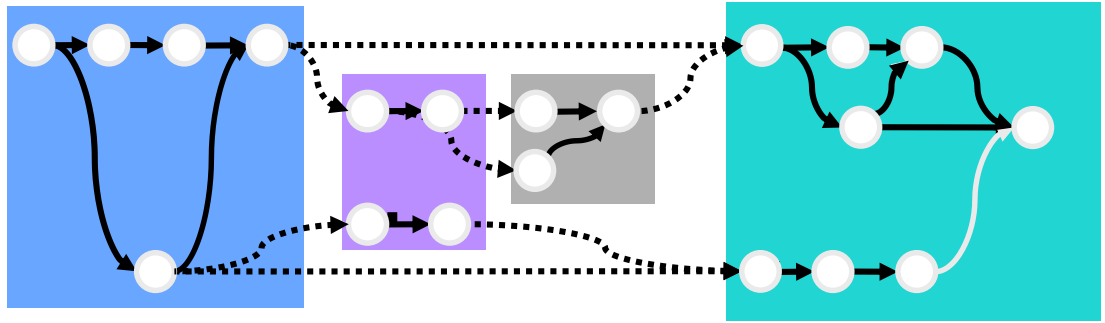
Static vs. dynamic composition

Static connections

- Specified at application development-time and do not change at run-time

Dynamic connections

- Partially specified at application development-time (Name or properties)
- Established at run-time, as new jobs come and go
 - Specifications can also be updated at run-time



Dynamic application composition

- Incremental deployment of applications
- Dynamic adaptation of applications

Application scenarios and real-world use cases

[IBM Streams](#) is being applied in many industries

- Market and Customer Intelligence
- Call Center Customer Care
- Manufacturing
- Personalized Customer Experience
- Network Analytics
- IoT, Connected Car and Telematics
- Cyber Security
- Health and Improved Patient Outcomes
- Operational Optimization



[Watch the Video](#)



[Watch the video](#)



[Watch the Video](#)



[Watch the Video](#)



[Watch the Video](#)
[Read the Case Study](#)



[Watch the video](#)



[Watch the video](#)



[Read the Case Study](#)



[Read the Case Study](#)



[Watch the video](#)



[Watch the Video](#)



[Watch the Video](#)



[Watch the Video](#)



[Watch the Video](#)



[Watch the Video](#)



[Watch the Video](#)



Science & Technology
Facilities Council

[Watch the Video](#)



[See the Presentation](#)



[Read the Case Study](#)



[Read the Case Study](#)

Medtronic Sugar.IQ

IBM Technology

Watson Platform for Health

- Mobile app management
- Data management
- Integration
- Analytics with IBM Streams

Medtronic Sugar.IQ

Past–Present–Future

How have I done?

- Important **glucose management*** information

How am I doing?

- **Near real-time personalized insights*** for better decision making

What should I be doing?

- **Predictive alerts**** to help avoid incidents to stay ahead

The Results: Sugar.IQ with Watson

90%

AUC accuracy level when predicting the risk of hypoglycemia two to four hours in advance. ***

Sugar.IQ 1.0 ‘Learning Launch’ showed*:

655

Hypo-related insights

699

Hyper-related insights

36 minutes

Average more in range per day

* Requires a Medtronic CGM device

** Planned feature

*** Data on File

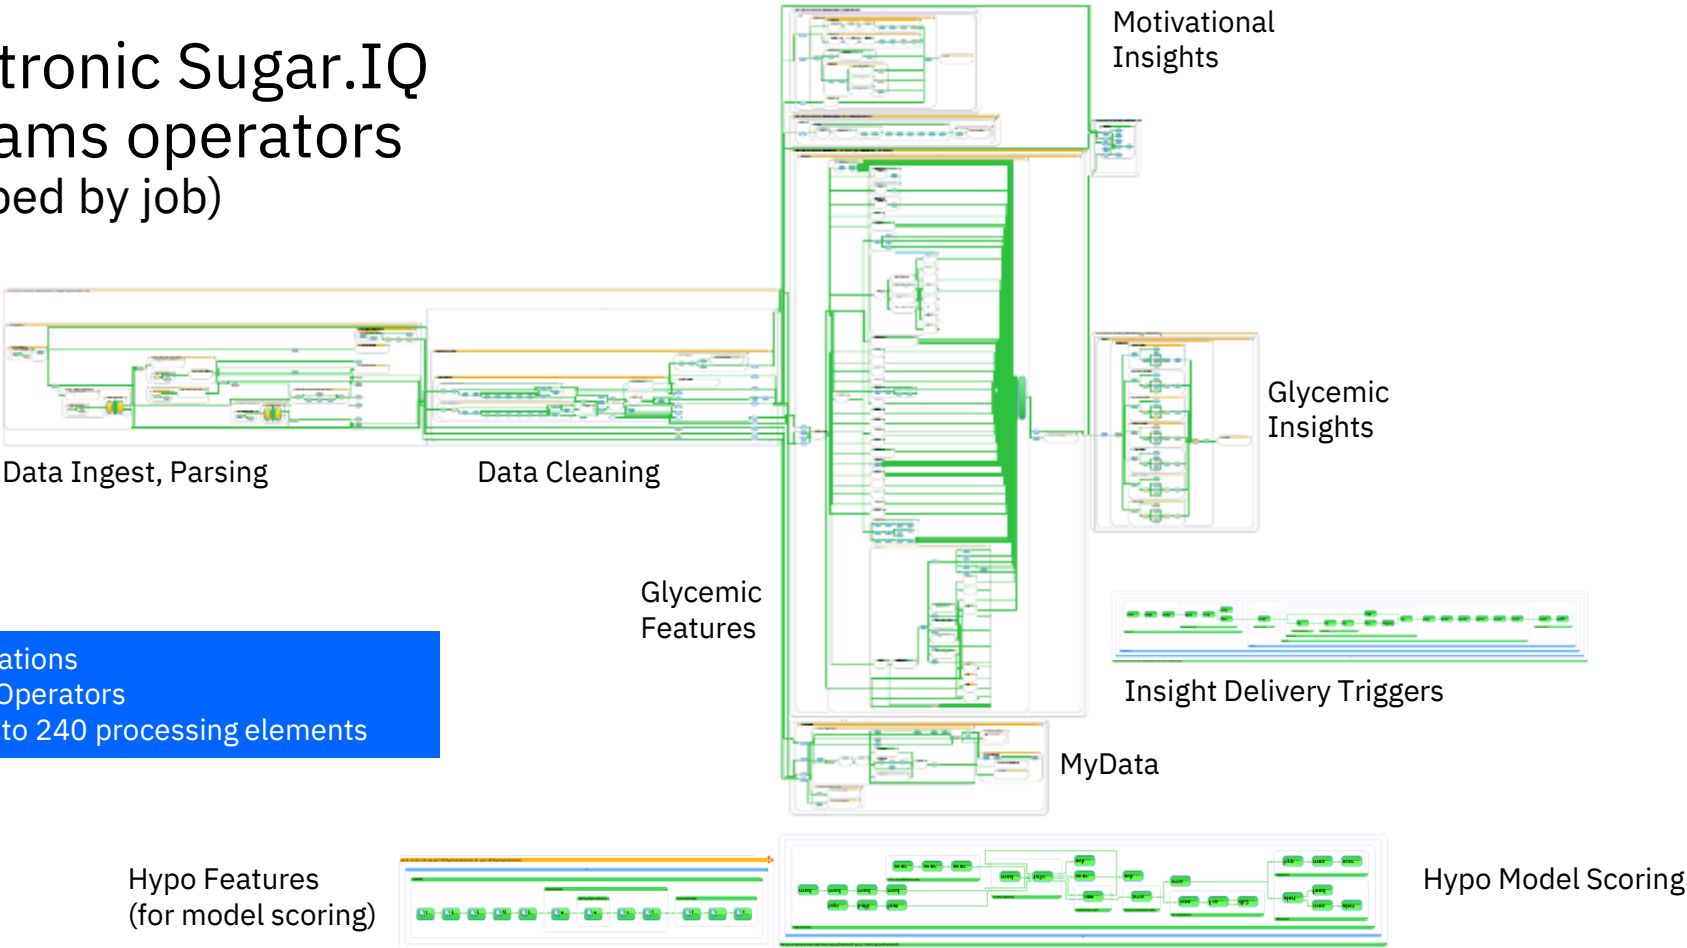
+ Data from all of the patients from Sugar.IQ ‘Learning Launch’ with Ver 1.0, Apr-Aug 2017. 256 total users

Medtronic Sugar.IQ

Streams operators

(grouped by job)

8 Applications
14,070 Operators
Fused into 240 processing elements



For continuous intelligence, **IBM Streams** is the clear choice

Continued leadership in high volume,
low latency streaming

- Ingest and analyze **massive volumes of streaming data**

Extend and embrace open source

- **Over 100 included** in Streams (Spark, Eclipse, Yarn)
- **Apache Edgent** donated to open source
- **github.com/IBMStreams** with over 50 projects

Simplified development

- **Java, Scala and Python** native development
- Rules development with deployment to Streams
- **Drag and Drop** visual development
- Simplified **Web based development** in Watson Studio

Applications

- Telco and Finance Solutions from IBM Analytics Solutions
- Partner applications in Healthcare, Telco
- Accelerators for Customer Care and Clickstream analytics

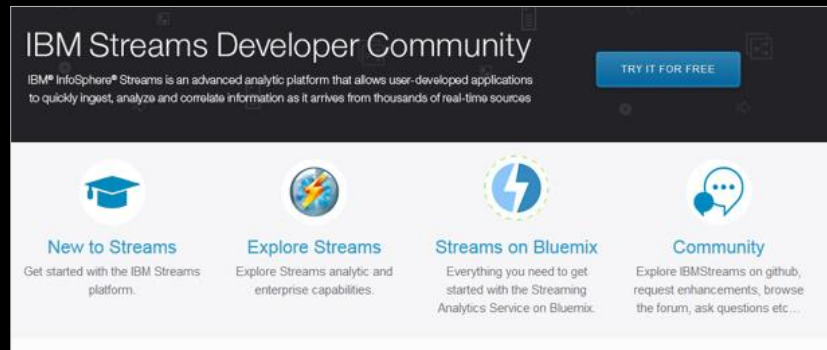
Multicloud deployment options with IBM Cloud Pak for Data

What is the next step?

Try Streams with [Cloud Pak for Data](#) and score models built in Watson Studio against real time data for Continuous Insights.

Try [Streaming Analytics on IBM Cloud](#) to capture data and enable intelligent applications so you can spot opportunities and risks sooner than the competition.

Join the [IBM Streams developer community](#) which is a direct channel to IBM Streams developers and a place to discuss, learn and share ideas.



The screenshot shows the IBM Streams Developer Community homepage. At the top, the title "IBM Streams Developer Community" is displayed, followed by a description: "IBM® InfoSphere® Streams is an advanced analytic platform that allows user-developed applications to quickly ingest, analyze and correlate information as it arrives from thousands of real-time sources". A "TRY IT FOR FREE" button is in the top right. Below this, there are four main sections, each with an icon and a title: "New to Streams" (graduation cap icon), "Explore Streams" (lightning bolt icon), "Streams on Bluemix" (lightning bolt icon), and "Community" (speech bubble icon). Each section has a brief description of what it offers.

IBM Streams Developer Community

IBM® InfoSphere® Streams is an advanced analytic platform that allows user-developed applications to quickly ingest, analyze and correlate information as it arrives from thousands of real-time sources

TRY IT FOR FREE

- New to Streams**
Get started with the IBM Streams platform.
- Explore Streams**
Explore Streams analytic and enterprise capabilities.
- Streams on Bluemix**
Everything you need to get started with the Streaming Analytics Service on Bluemix.
- Community**
Explore IBMStreams on github, request enhancements, browse the forum, ask questions etc...



The screenshot shows a promotional banner for Streaming Analytics. It features a dark background with a glowing, curved light trail on the right side. The text "Streaming Analytics" is prominently displayed at the top. Below it, a subtitle reads: "Leverage continuously available data from all sources to discover opportunities faster." At the bottom left, there is a blue button labeled "Try Streaming Analytics". To its right, there is a link with an external icon and the text "Read Forrester report".

Streaming Analytics

Leverage continuously available data from all sources to discover opportunities faster.

Try Streaming Analytics

Read Forrester report



Copyright © 2019 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights – use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environment. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a particular purpose.**

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com, Aspera®, Bluemix, Blueworks Live, CICS, Clearcase, Cognos®, DOORS®, Emptoris®, Enterprise Document Management System™, FASP®, FileNet®, Global Business Services®, Global Technology Services®, IBM ExperienceOne™, IBM SmartCloud®, IBM Social Business®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®, Netcool®, OMEGAMON, OpenPower, PureAnalytics™, PureApplication®, pureCluster™, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, Smarter Commerce®, SoDA, SPSS, Sterling Commerce®, StoredIQ, Tealeaf®, Tivoli® Trusteer®, Unica®, urban@code®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.

Source:

1. <https://online.stanford.edu/course/machine-learning-1>

