TDWI WEBINAR

Transforming Your Business With Data and AI

James Kobielus
Senior Director Of Research
for Data Management

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

SPONSOR

databricks

STARDOG

JAMES KOBIELUS

Senior Research
Director,
Data Management
TDWI

Digital business transformation rides on a new generation of applications—automated, intelligent, adaptive, and self-optimizing--that are powered by data and artificial intelligence.

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

Data is central to digital transformation

- Data that makes processes more agile, automated, predictive, and contextual
- Data for running processes in the cloud
- Data for orchestrating processes 24x7
- Data for enforcing strong governance and control of complex processes
- Data for accelerating process insights
- Data for cross-domain sharing of process intelligence

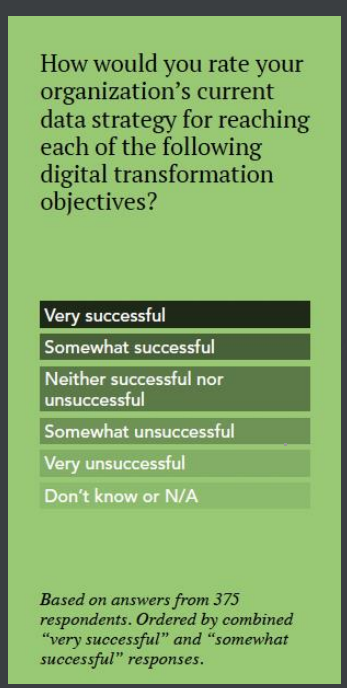

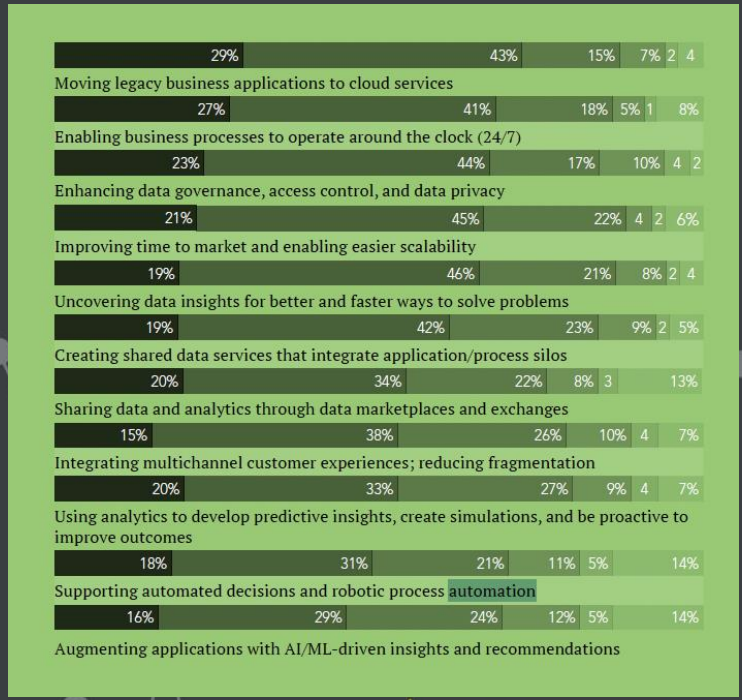

Source: TDWI Best Practices Report "Maximizing Business Value with Data Platforms, Data Integration, and Data Management," Q3 2022

TRANSFORMING DATA WITH INTELLIGENCE™

Data, analytics, & AI drive continuous transformation

AI, ML, and other data-driven process and decision logic

Process management and decisioning platforms

SaaS, cloud, mobile, edge, embedded, IoT, and other application platforms

Process automation

Autonomous operation

Sensor-driven actuation

Augmented human decisioning

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

AI-driven transformation depends on adaptive automation

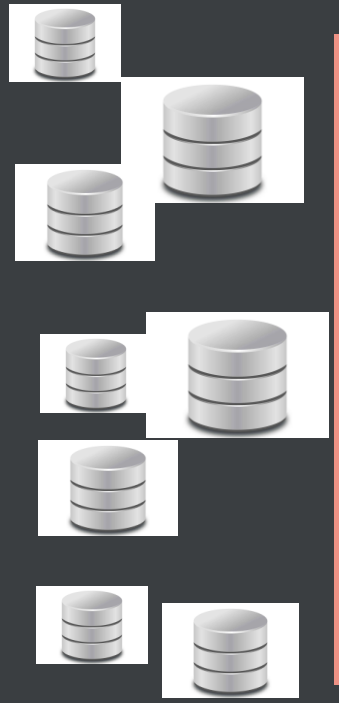

data, models, analytics, rules, policies, metadata, and other business logic

Process and decision automation engines

Automated actions


Sales

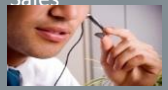
Marketing


Service


Logistics


Manufacturing

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

Enterprises are putting AI models to work in their organizations.

How many models (e.g.,predictive analytics/machine learning) does your organization currently have in production?

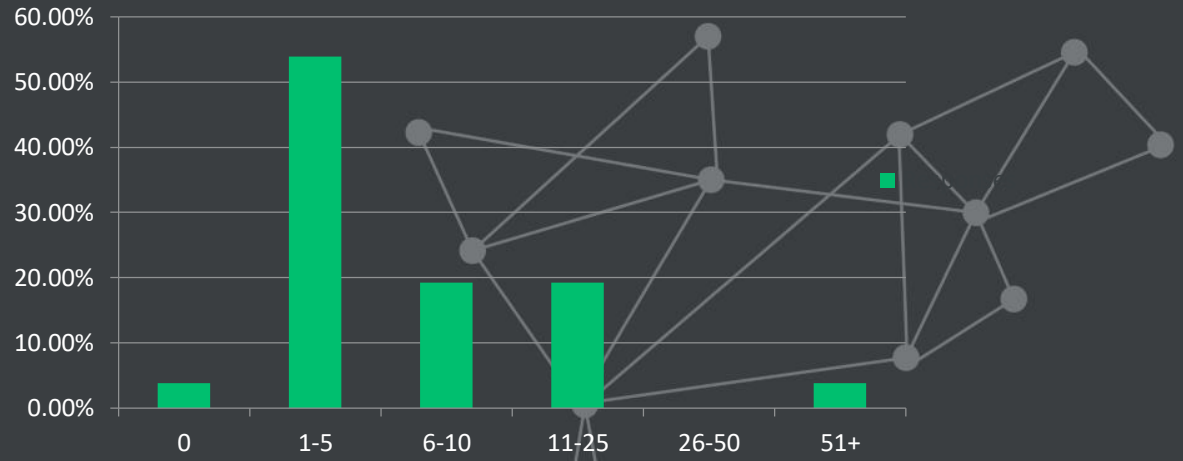

Source: TDWI Data & Analytics Survey 2022

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

AI/ML will continue to grow in footprint in enterprise application environments.

How has demand for more analytics such as machine learning or text analytics changed in your organization over the past year?

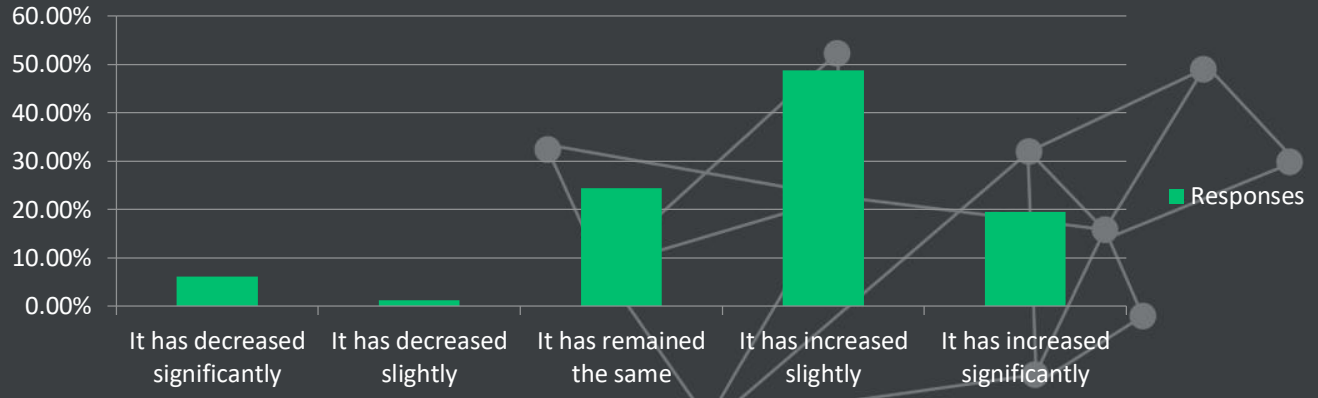

Source: TDWI Data & Analytics Survey 2022

TRANSFORMING DATA WITH INTELLIGENCE™

Enterprises are relying on professional data scientists to develop AI/ML apps, often in collaboration with business analysts and subject matter experts.

Who is building predictive analytics/machine learning models in your organization?

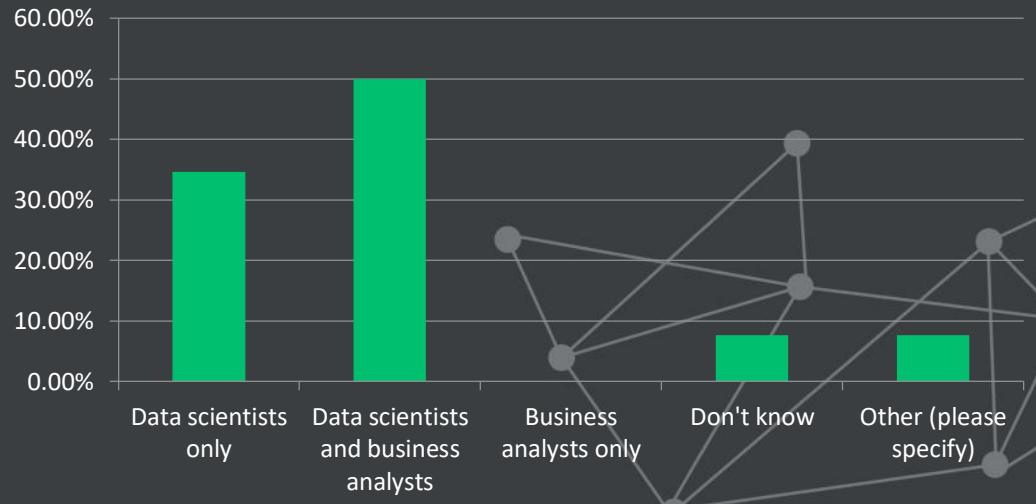

Source: TDWI Data & Analytics Survey 2022

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

Operationalizing the deployment and management of AI is a high enterprise priority.

Does your company have an MLOps team that deals with putting predictive analytics/machine learning models into production and monitoring them?

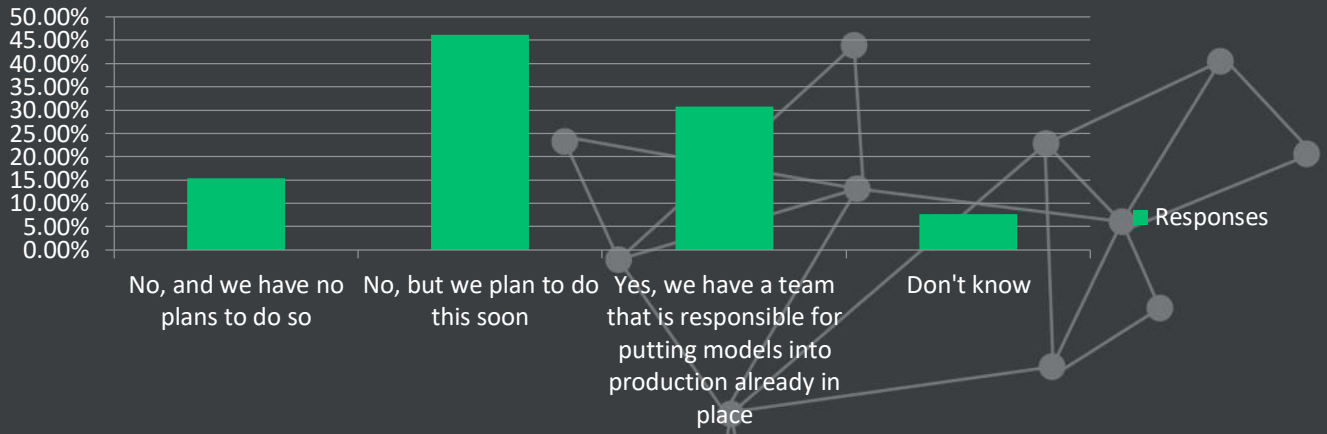

Source: TDWI Data & Analytics Survey 2022

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

BALA AMAVASAI

Global Technical Director
Databricks

6,000+
global employees

10,000+
Customers

$3B
in investment

Inventor and pioneer of the **data lakehouse**

databricks

The Lakehouse Company

Gartner-recognized Leader
Cloud Database Management Systems
&
Data Science and Machine Learning Platforms

Creator of

DELTA LAKE

APACHE Spark™

mlflow™

Lakehouse - one platform for Data and AI

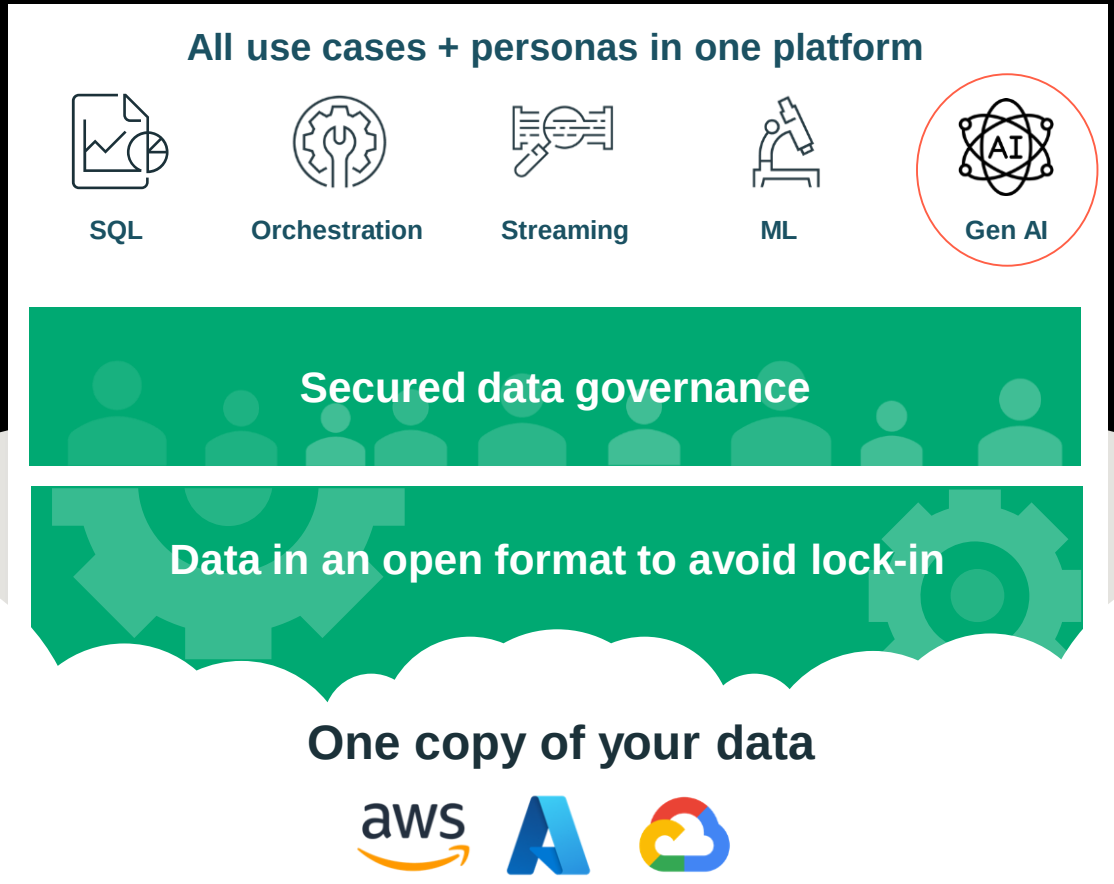

Databricks helps the biggest manufacturing brands build data-driven businesses

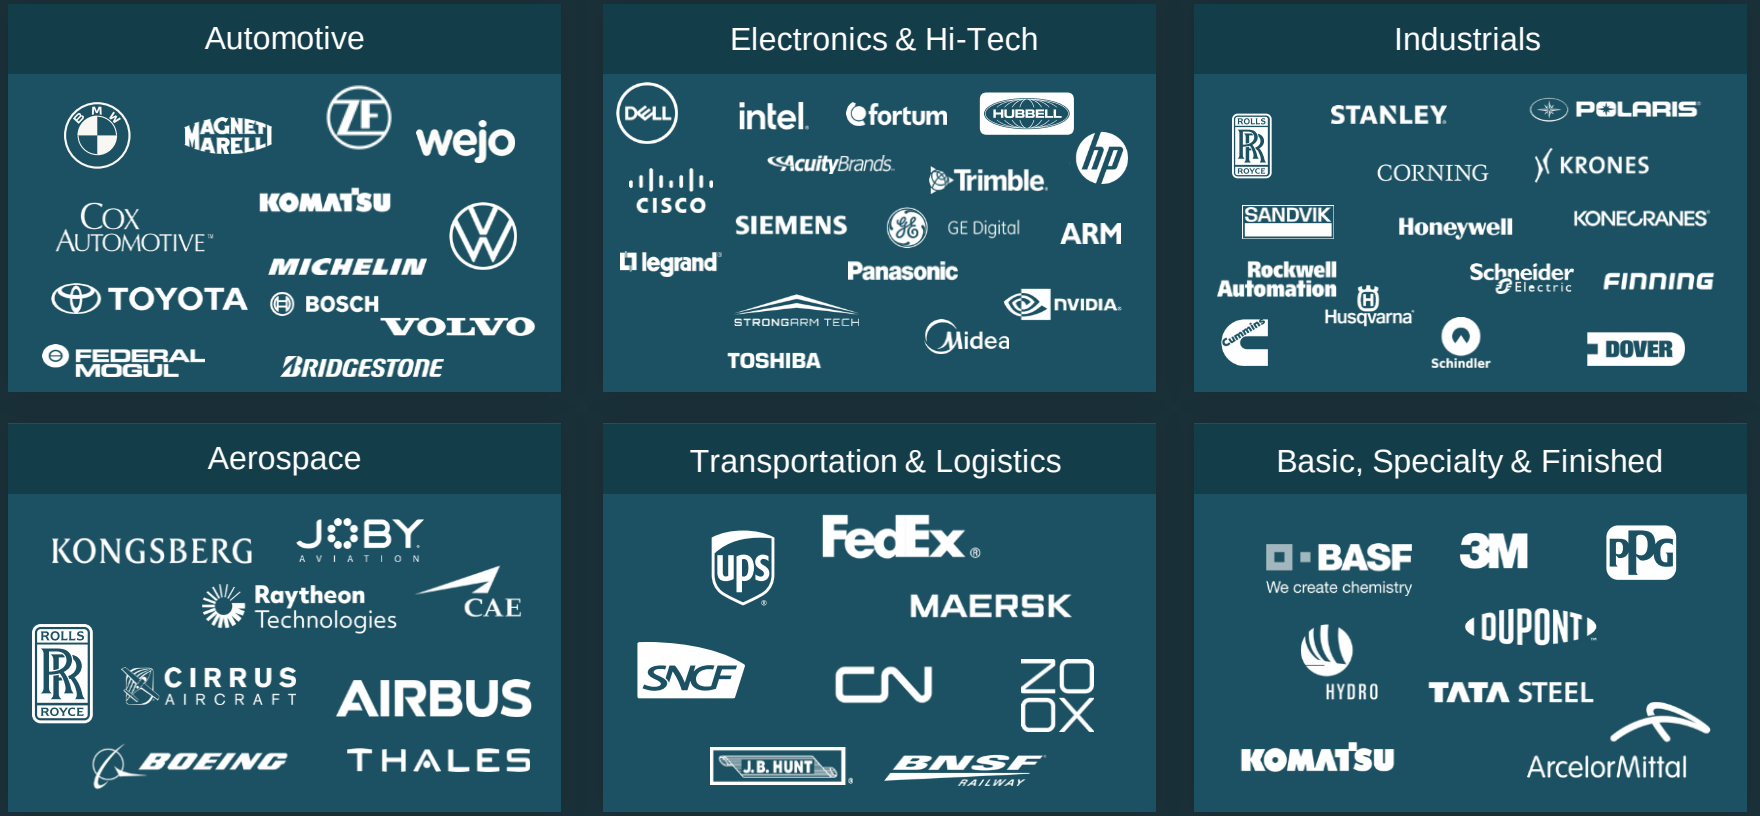

Accelerating Manufacturing towards Industry 4.0 and beyond

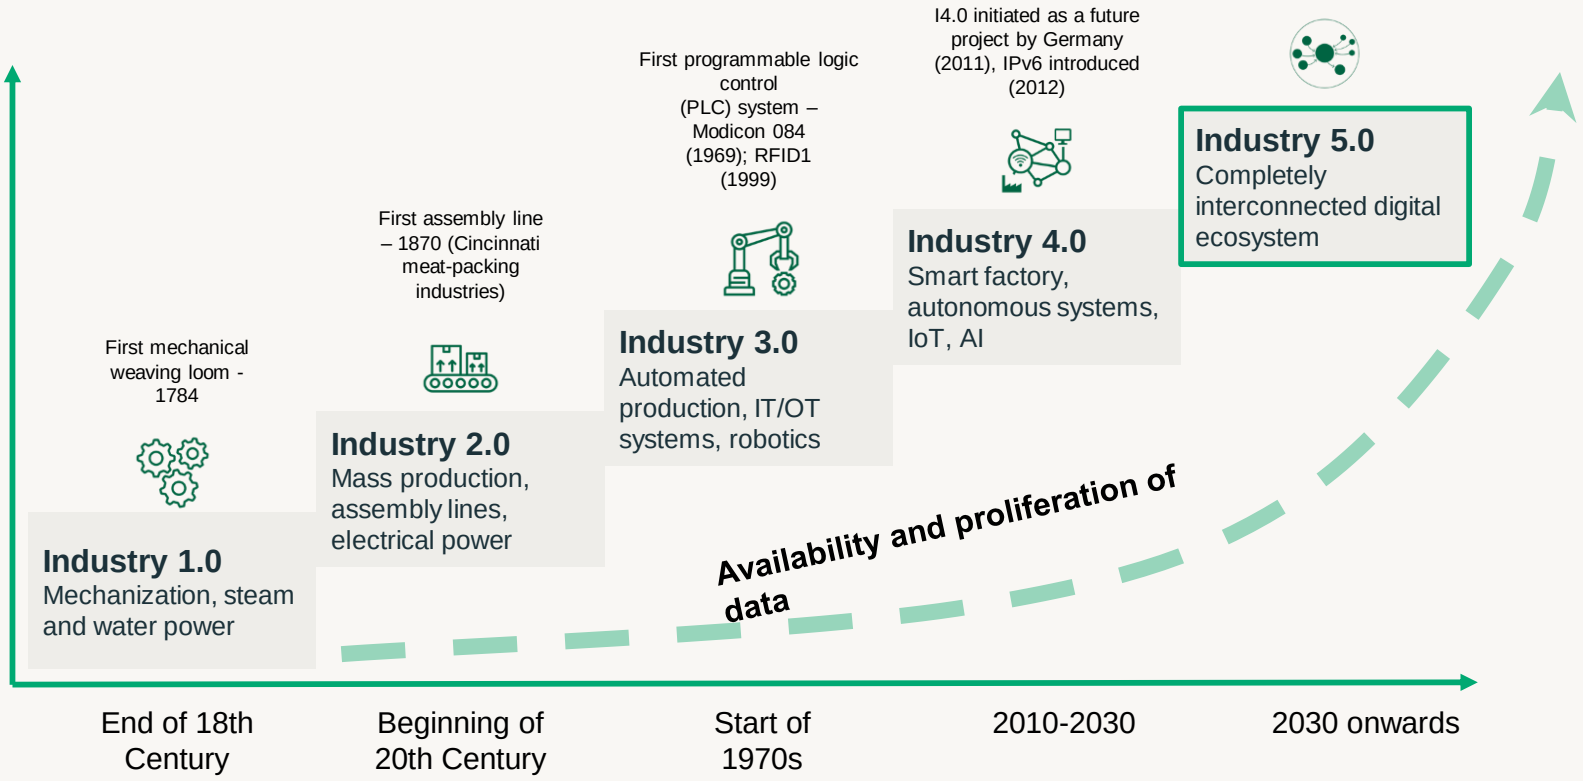

Databricks enables analytics processing of large-scale data using Apache Spark™, Delta, and MLflow.

- Best-in-class support from original creators of Spark
- 10-100x faster than Open Source Spark when using Databricks as ETL engine
- Lowest TCO through auto-scaling and auto-configuration capabilities
- Delta Lake provides the ability to specify your schema and enforce it, making it ready for analytics at petabyte scale

The Stardog Enterprise Knowledge Graph platform unifies data based on its meaning, powering a semantic data layer to unleash valuable insights and accelerate your investment in a Databricks lakehouse.

- Richest, most flexible semantic layer
- Most advanced graph data virtualization
- Easiest to develop and reuse data models
- Most intuitive search and discovery of data insights
- Designed for enterprise data and analytics work

databricks x STARDOG

Democratizing Data and AI across the Manufacturing stack

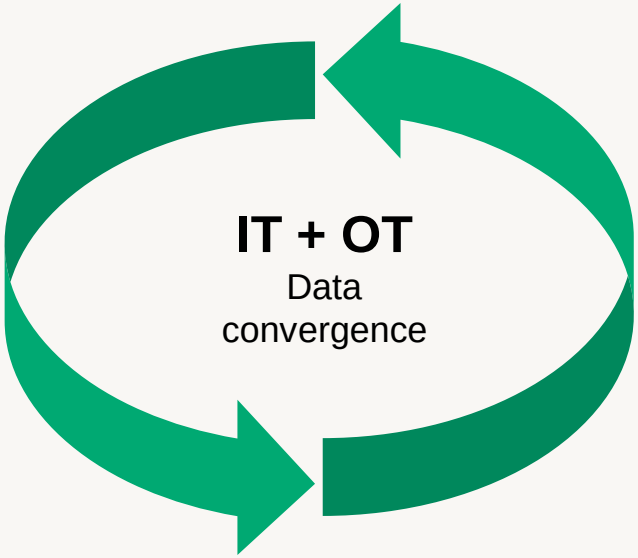

Data from ERP, MES, SCADA, PLCs & Sensors

Data democratization unlocks business value

Monitoring and Analytics

Tableau PowerBI Redash Qlik

Data Products

Key Use Cases

Asset & Operations Optimization
- Predictive Maintenance with Logistics
- Production Line Optimization
- Overall Equipment Effectiveness
- Digital Twin

Production Control
- Yield Optimization
- Automated Quality Control
- Automated Visual Inspection
- Process Control

Supply Chain & Inventory Management
- Inventory Planning & Optimization
- Material Tracking
- Part Forecasting
- Supply Chain Visibility

Demand Planning
- Demand Forecasting
- Demand Sensing
- Inventory Optimization

Finance
- Reporting: P&L, Profitability, Balance Sheet
- Forecasting: Revenue, Cashflow

AL BAKER

VP, Enterprise Solutions
Stardog

Stardog

Company Overview

Unite Data, Unleash Insight.

Stardog's **Enterprise Knowledge Graph platform** connects data based on business meaning into a flexible, reusable semantic data layer to get better insight faster.

We help **customers across many industries** empower data citizens to make knowledge-informed decisions.

> "Our primary objective is to provide data at a higher quality and relieve the heavy lifting up front so our data scientists can actually work with the data."
>
> — Head of IT Research, Top Global Pharma

Select Stardog Customers:

BOSCH

Schneider Electric

BAYER

NASA

EXFO

Boehringer Ingelheim

CIS Center for Internet Security®
Creating Confidence in the Connected World.

STARDOG

What is an Enterprise Knowledge Graph?

A flexible, semantic data layer for answering complex queries across data silos.

- Connects any data source / location and virtualizes access
- Enriches real-world context into data
- Infers new relationships, patterns and insights in data
- Semantically searches data by meaning

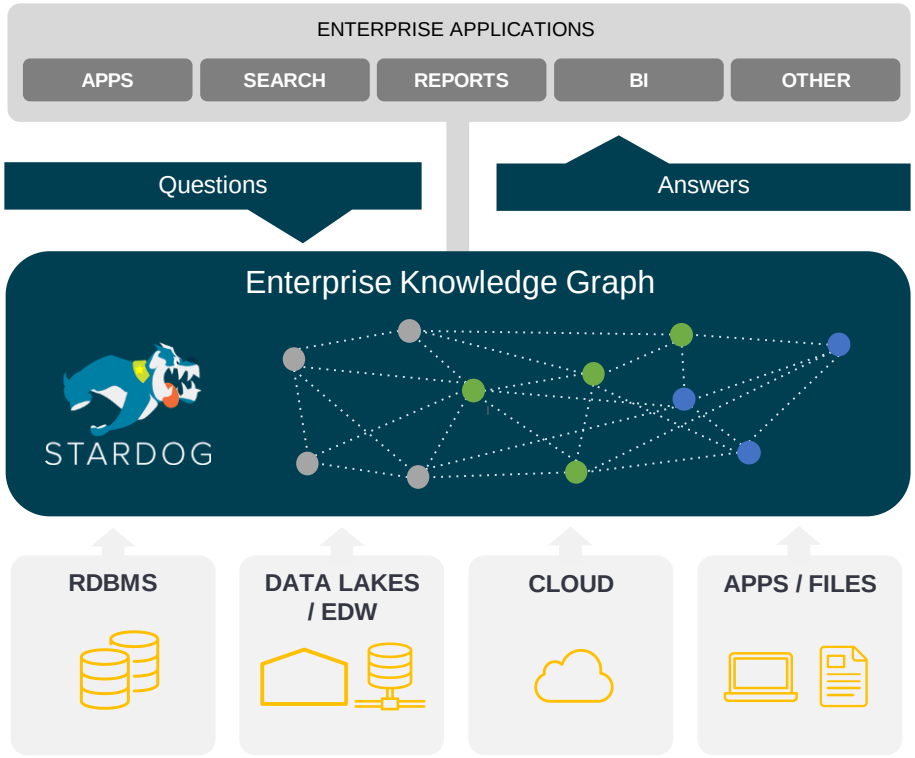

How Industrial Organizations Use Stardog

Supply Chain & Logistics

Product 360

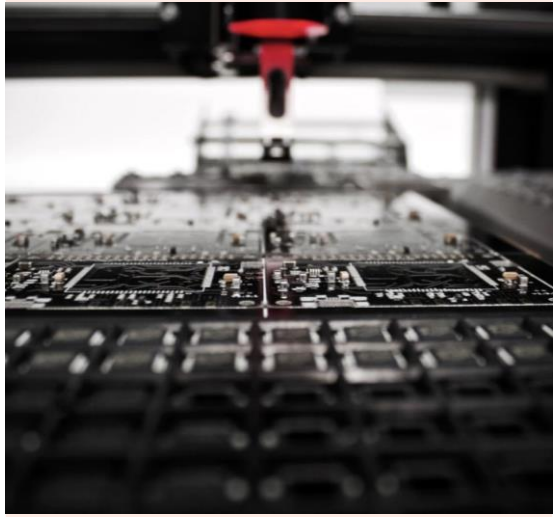

Modernize MBSE

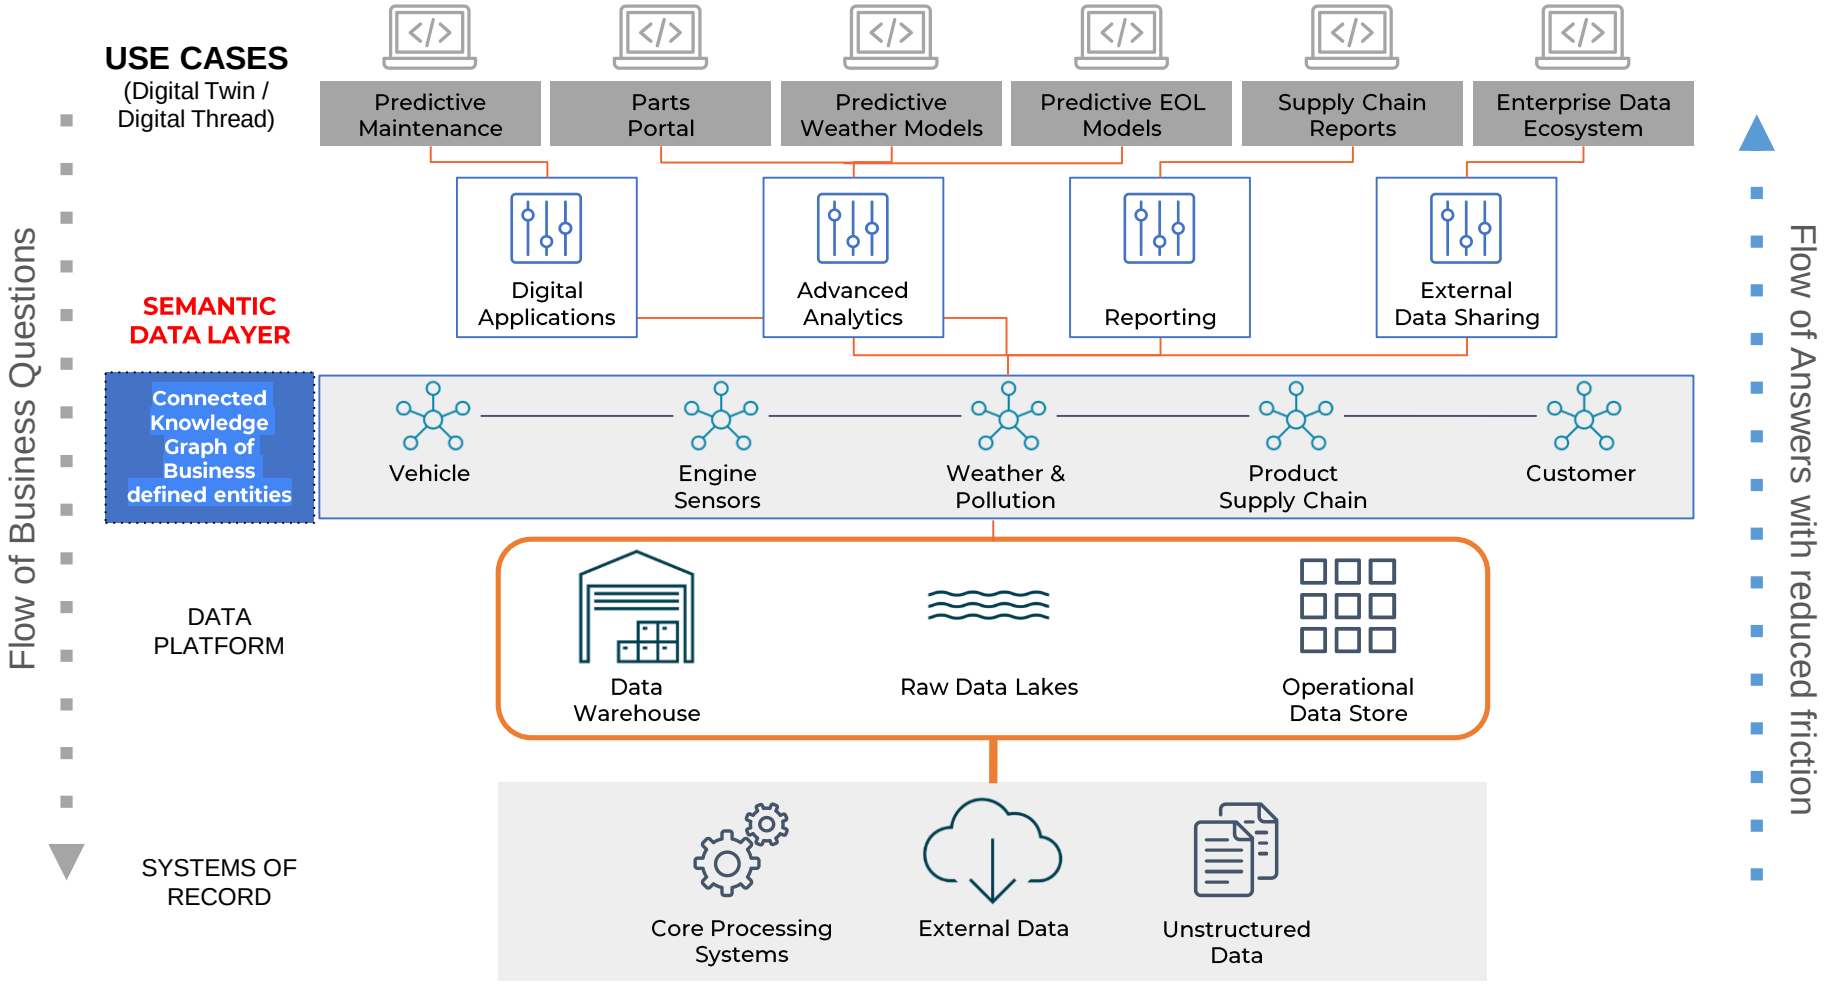
USE CASES
(Digital Twin / Digital Thread)
Predictive Maintenance
Parts Portal
Predictive Weather Models
Predictive EOL Models
Supply Chain Reports
Enterprise Data Ecosystem
SEMANTIC DATA LAYER
Digital Applications
Advanced Analytics
Reporting
External Data Sharing
Connected Knowledge Graph of Business defined entities
Vehicle
Engine Sensors
Weather & Pollution
Product Supply Chain
Customer
DATA PLATFORM
Data Warehouse
Raw Data Lakes
Operational Data Store
SYSTEMS OF RECORD
Core Processing Systems
External Data
Unstructured Data
Flow of Business Questions
Flow of Answers with reduced friction

Example: Root Cause Analysis

Markdown

Root cause analysis from flight to factory

Analysts at the Acme Aircraft Company have noticed an uptick in warning alerts triggered by engine sensors, but it's not immediately clear why only some engines of the same model are affected or why the effect is stronger on some flights than others.

In this demo, we'll use a knowledge graph to discover what common factors are associated with the warnings and track the cause all the way back to a specific badly manufactured lot of a component used in the engines.

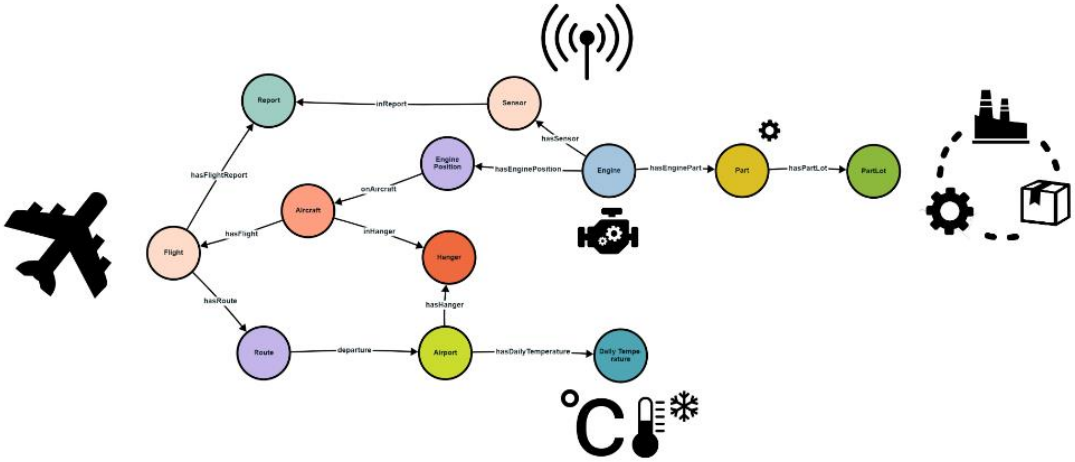

Example: Databricks and Stardog Notebook

Follow a pilot's hunch: does temperature influence the rate of sensor alarms?

To find out if the pilot's hunch was correct, we'll follow a path through the knowledge graph to connect sensor data with weather reports at the departure airport.

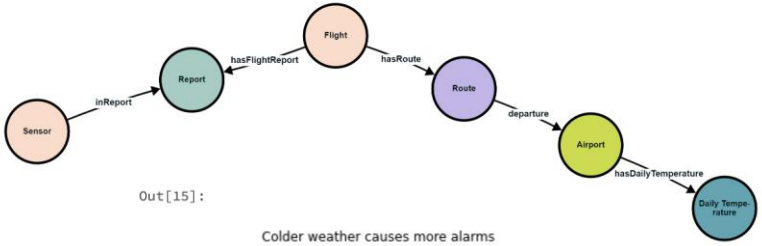

Out[15]:

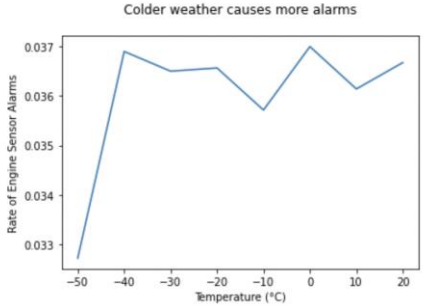


```
q = """
select ?bad_part_serial_no ?engine_serial_no ?position ?aircraft_tail_no {
    ?part acme:hasPartLot ?partLot ;
          acme:partSN ?bad_part_serial_no .
    values ?partLot {acme:lot2b-22}
    ?engine acme:hasEnginePart ?part ;
          acme:engineSN ?engine_serial_no ;
          acme:hasEnginePosition ?pos .
    ?pos acme:onAircraft ?aircraft ;
          acme:position ?position .
    ?aircraft acme:tailNumber ?aircraft_tail_no.
}
ORDER BY ?aircraft_tail_no
"""
bad_parts_df = sg.query(q)
bad_parts_df.head()
```

Out[14]:

	bad_part_serial_no	engine_serial_no	position	aircraft_tail_no
0	P2B-ph5	E-a8f	Tail	B-af
1	P2B-no9	E-h8o	Left	B-ag
2	P2B-tn8	E-k5v	Tail	B-bc
3	P2B-ix9	E-l8v	Right	B-cw
4	P2B-dg8	E-p4a	Left	B-dq

Data Analysis:

Data Segmentation into Knowledge for accelerating ML

Inferring Classifications for Different Types of Flights

The engines that Acme manufactures are expected to perform differently under different flight conditions

In this demo, we'll use a knowledge graph to classify types of flights and engine performance according to significant features

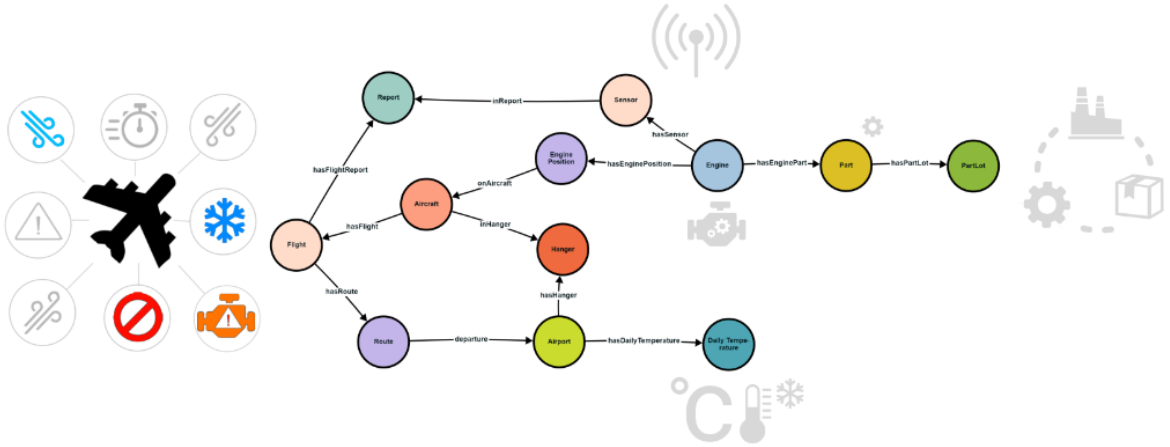

Our data scientists need the ability to classify flights or flight segments into analytical categories, for example:

- cold / warm weather
- headwind / tailwind / crosswind
- high / low altitude
- long / short duration
- long / short distance

Our goal: Formalize our technical and business logic as inference rules to classify flights

- Add inference rules to our ontology that apply flight classifications and helpful shortcuts
- Create a chain of inferences for a cascading effect that simplifies what would otherwise be complex queries

STARDOG

Using Datalog with Stardog and Databricks

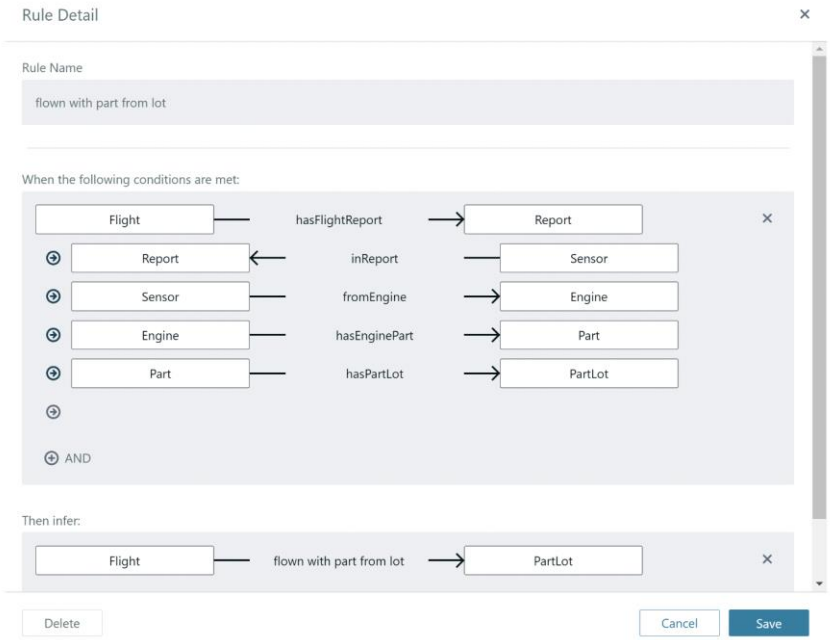

The simple query below has to do a lot!

It will...

- follow a hierarchy linking flight -> aircraft -> engines -> engine parts -> manufacturing lots
- coordinate that to another hierarchy of flight -> report -> engine -> sensor -> alarms
- calculate a ratio of observed alarms per minute of flight for each manufactured lot of each component part in the aircraft engines
- compare that ratio to a predetermined parameter
- classify the manufacturing lot according to whether the observed rate of alarms exceeded the parameter
- classify the flight by whether the aircraft contains a part from a bad lot
- cross-reference flight departure dates against each airport's daily weather data to classify especially cold weather flights
- and finally, classify the flight as `Grounded` if it both contains a bad part and is scheduled for takeoff on a very cold day

Now let's run the query.

Cmd 18

```
1  q = """select ?flight { ?flight a <http://acme.stardog.com/Grounded_Flight> } limit 10"""
2  df = sg.query(q, reasoning = True)
3  df.head(5)
```

Data Analysis:

Use Knowledge Graph logic over Time Series tables in Databricks

Markdown

Time Series Analysis on Aircraft Engine Sensor Data

Data scientists on Acme Aircraft's quality assurance team regularly analyze real world in-flight sensor data to improve the precision of design specifications for various types of flights.

In this demo, we'll show how to use a knowledge graph to record, analyze, and interpret time series data.

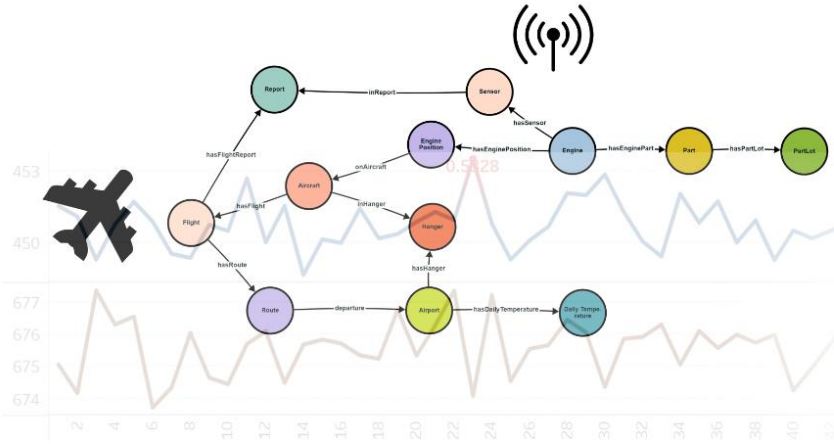

Scenario - Discover what patterns the sensor data can reveal about in-flight engine performance for different types of flights

- Acme corporation manufactures aircraft and publishes expected performance metrics in a design specification document
- Sensors in the engines record measurements at every minute of a flight
- The sensor data can be compared with the expected metrics to evaluate real-world engine performance

Our goal: Analyze time series data in the graph to determine whether engine performance was within the parameters of Acme's design specifications or was nonconformant

STARDOG

Graph Queries over Time Series

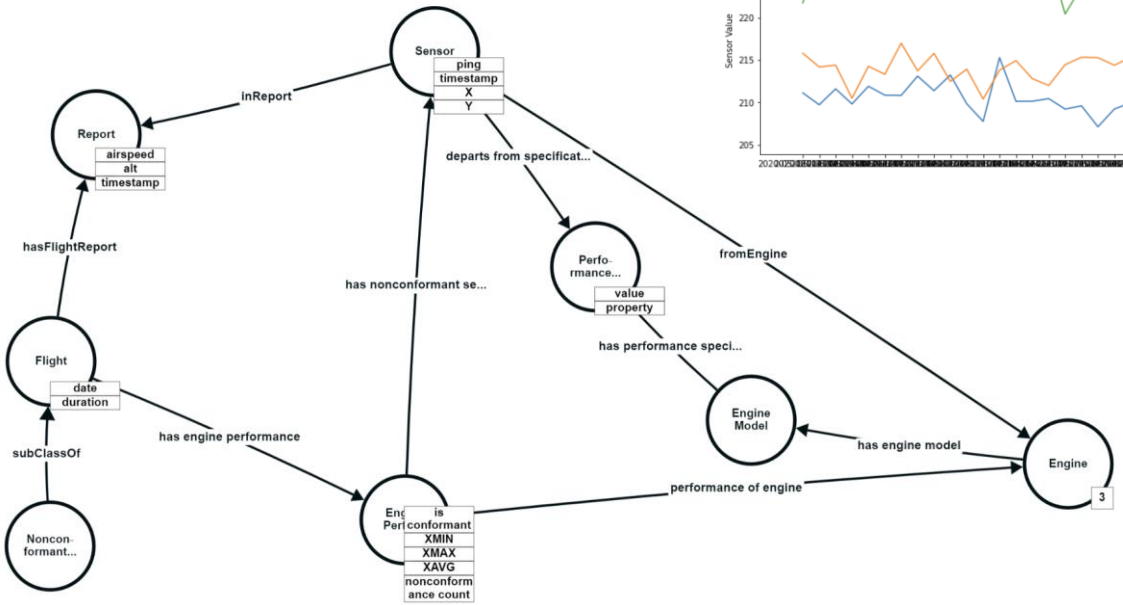

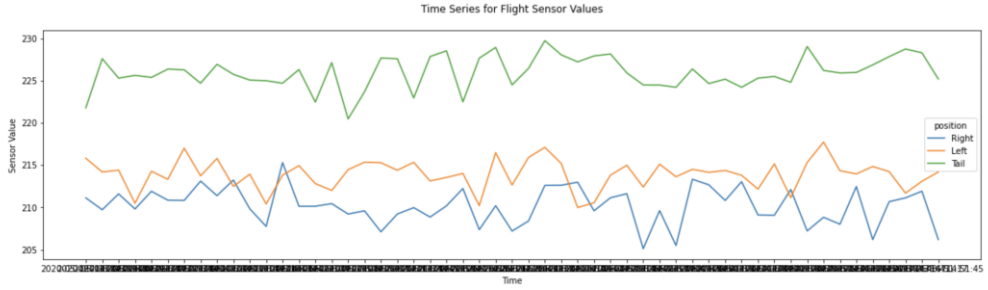

Visualize Graph Results in Stardog Explorer

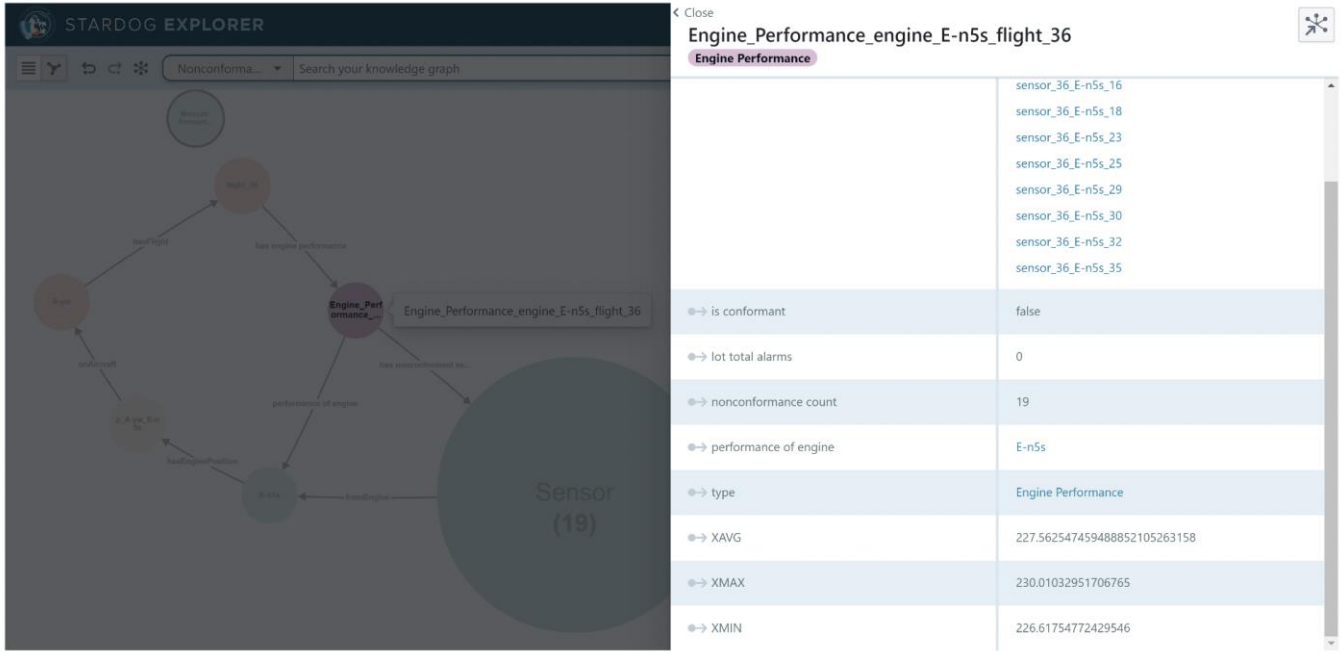

Visualize Graph Results in Stardog Explorer

...Or visualize other important information, for example that this aircraft flew 36 flights with 2/3 of its engines in nonconformance!

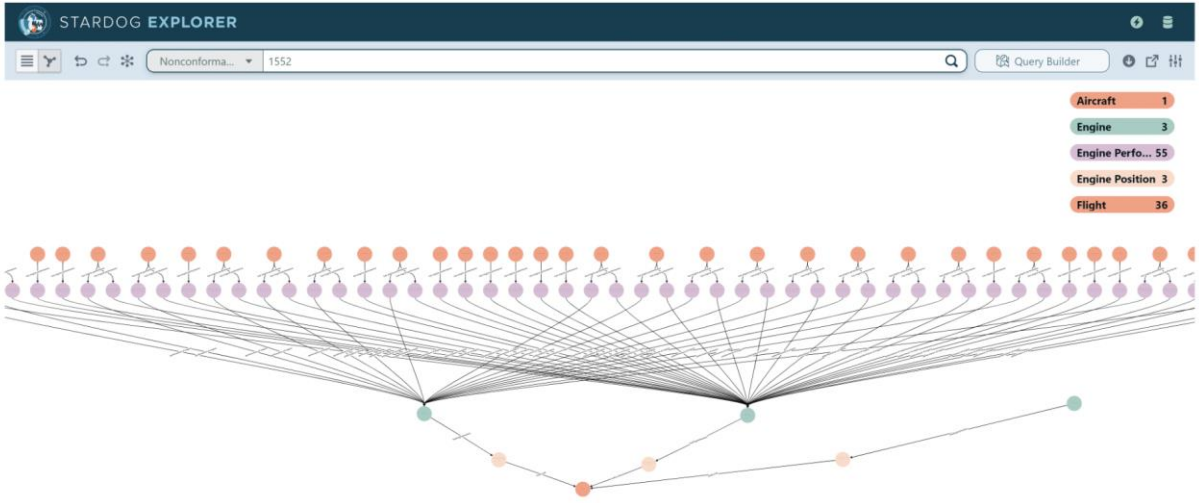

Quick-start Knowledge Kits

Find this kit at:
https://cloud.stardog.com/kits

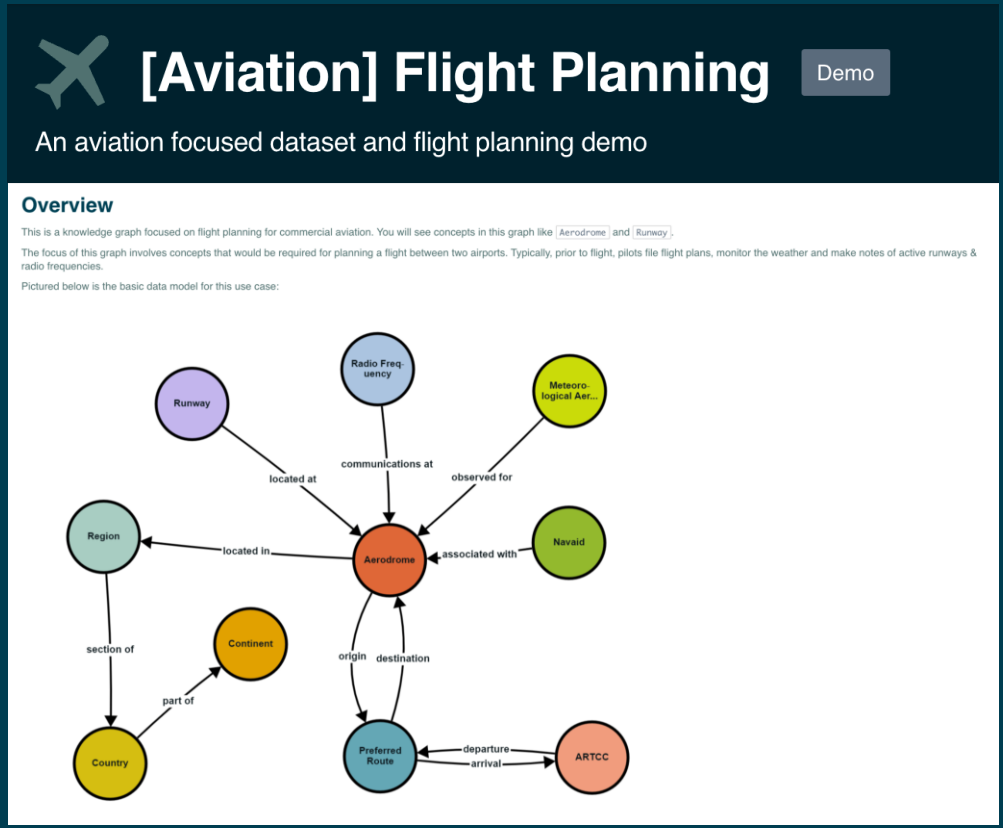

What is Industry 4.0?

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

How does Industry 4.0 support manufacturers' digital transformation initiatives?

What does it mean for manufacturing organizations to leverage data-driven insights to open up new business opportunities and accelerate positive outcomes within their supply chains?

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

How does digital transformation of manufacturers' supply chains depend on enterprise investments in modern data and AI technologies?

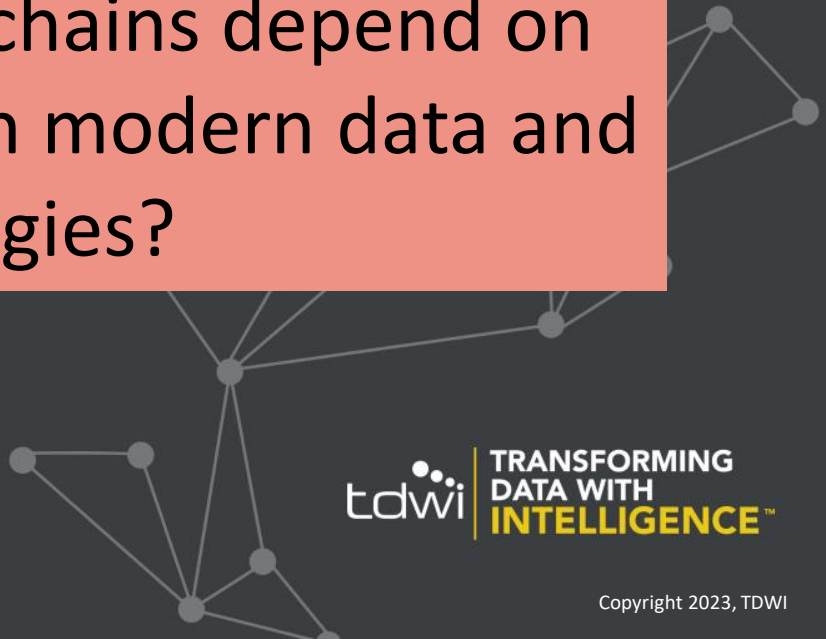

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

How can manufacturers boost the return on their digital transformation initiatives through investments in customer 360, semantic knowledge layers, predictive maintenance, edge computing, digital twins, the Internet of Things, and other pillars of “Industry 4.0”?

How is the payoff from investments in digital transformation boosted by providing self-service access to data and analytics for everyone in the organization regardless of their role or technical experience?

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

Questions?

tdwi TRANSFORMING DATA WITH INTELLIGENCE™

CONTACT INFORMATION

If you have further questions or comments:

James Kobielus, TDWI
jkobielus@tdwi.org

Bala Amavasai, Databricks
bala.amavasai@databricks.com

Al Baker, Stardog
al.baker@stardog.com

SPONSOR

databricks

STARDOG

Thank you for attending!

tdwi TRANSFORMING DATA WITH INTELLIGENCE™