



STREAM I

DATA ARCHITECTURE

Design of the data architecture is a strategic decision, reflecting an understanding of what matters for the performance of the revenue function.

Lecture #4

21 September 2021

Creating a data analytics function

This module will discuss how a professional data analytics function is introduced into a tax administration and the steps required to make it effective. Central to the organizational change is 1. *developing a strategy* for infusing data analysts and high-end skills of data science into taxation by:

2. deciding where to place the function
3. deciding on centralized versus decentralized data science capacity
4. establishing effective collaboration between data analysts, data science and tax specialists
5. educating tax professionals about data science and data scientists about taxation
6. providing space for experimentation, learning, and mistakes

Presented by:

Inland Revenue Authority of Singapore and SAS Institute



Our Objective

The objective of this stream is to equip participants with a practical and theoretical understanding of the critical success factors in establishing a data architecture by providing the appropriate tools and the vocabulary needed to establish or improve the tax administration's data architecture.

"Without facts and principles, data is useless."

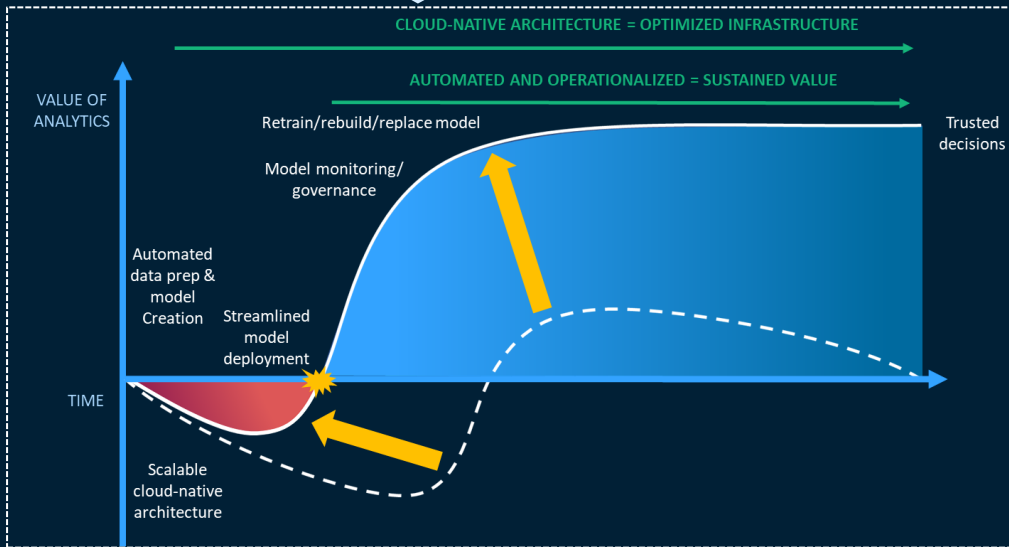
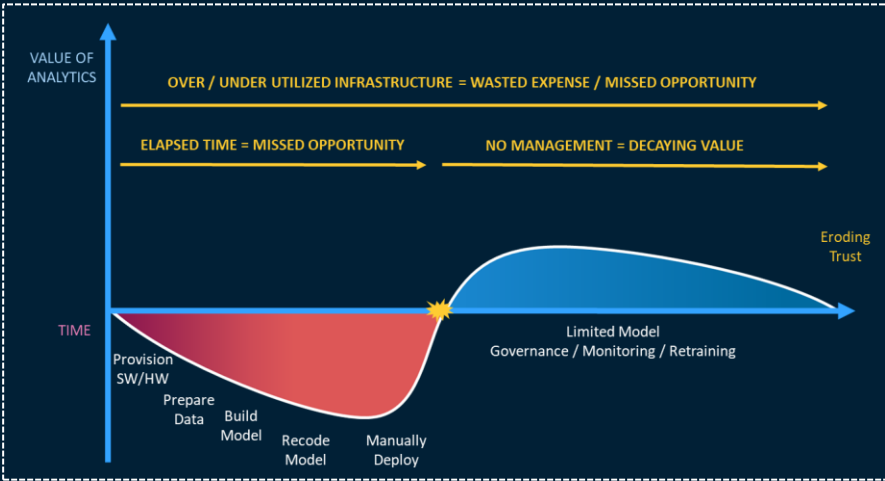
Bob Hoffman, Writer, Speaker and Type A Group Partner

Even though tax agencies are awash in data, they can struggle to turn raw data into actionable intelligence.

How can you apply the right analytics to produce intelligent decisions?



Using Analytics to Move Towards “Trusted Decisions”



Creating a Data Analytics Function

Overview

1

Developing a data strategy

2

Where to locate the data analytics function

3

Centralized vs de-centralized

4

Stakeholder collaboration

5

Capabilities and education

6

Experimentation & learning



Leveraging Analytics

Increase voluntary compliance, minimize loss and administer tax policy fairly and efficiently

Creating a Data Analytics Function

1

Developing a data strategy

2

Where to locate the data analytics function

3

Centralized vs de-centralized

4

Stakeholder collaboration

5

Capabilities and education

6

Experimentation & learning

Common Issues – Sound Familiar?

“I can’t wait days or even weeks for IT to have my data ready for me”

“Dated data = dated insights”

“We didn’t hire data scientists to prepare data all day ”

“I want to explore, experiment, and try new things by myself to spark new ideas and approaches for innovation.”

CASE STUDY:

Problem: Every project at the tax agency addressed data issues as one-off, built-from scratch activities. The agencies IT team had 17 projects underway (new applications, application enhancements, new reports, etc.).

Each project required access to customer data, and each had overlapping tasks and resources.

- Every project included a source data inventory and analysis activity because there was *no way to know where specific data resided*.
- New data extracts (subsets of data copied for use by other systems) had to be built because *IT had no way of determining if the data was already available*.
- *No teams shared their source extract data*. Each had their own copies to support their integration and database build activities (which tied up storage for this transient content).
- Each team's *integration logic was custom built and individually maintained*, because the logic and rules weren't identified or documented to be shared.

Departments and divisions do not have a common method for collecting and integrating raw data, cleansing it to ensure data quality, and preparing it for analytics.



Audit had to *continually update* its campaign system to adjust to frequent (and uncommunicated) changes occurring to the layouts of the extracts it received.



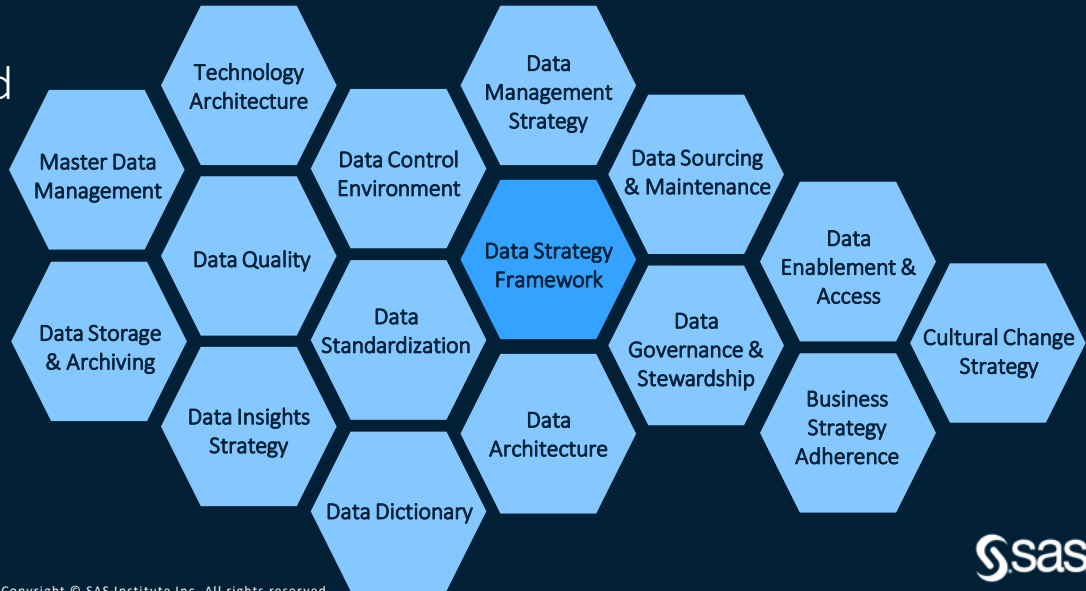
Managers always had questions about KPI reports because *titles and labels varied across reports* (even though they contained common data).

Business unit users often *built their own reports instead of using the standard reports from finance*, because there was no way to determine the origin of standard report data.



Common questions about the data *did not have a common version of the “truth”* so non-compliance detection and reporting was usually ad-hoc.

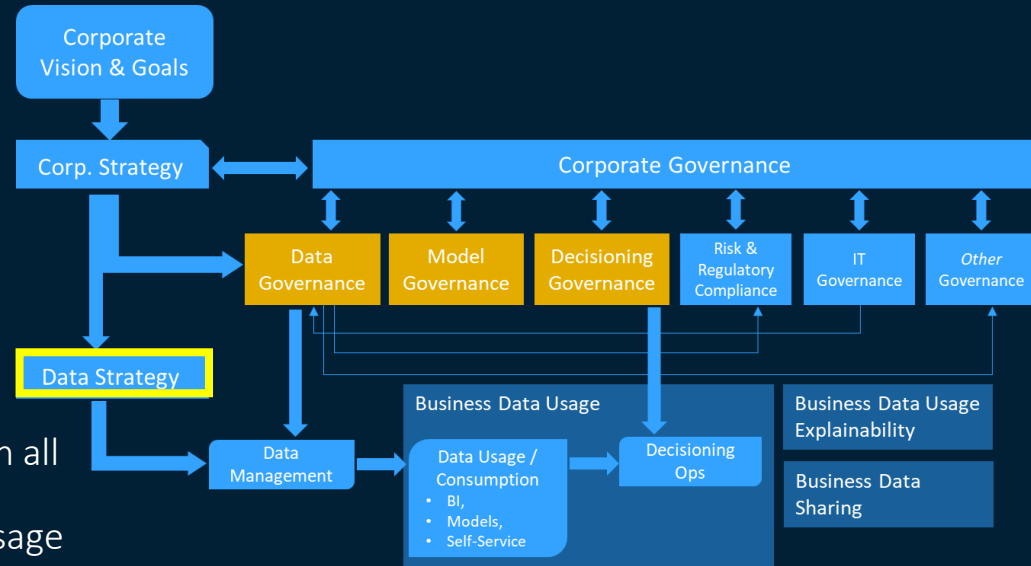
The data warehousing team lacked a **Data Strategy Framework** and had to continually chase data problems because data issues weren't managed like other business support activities.



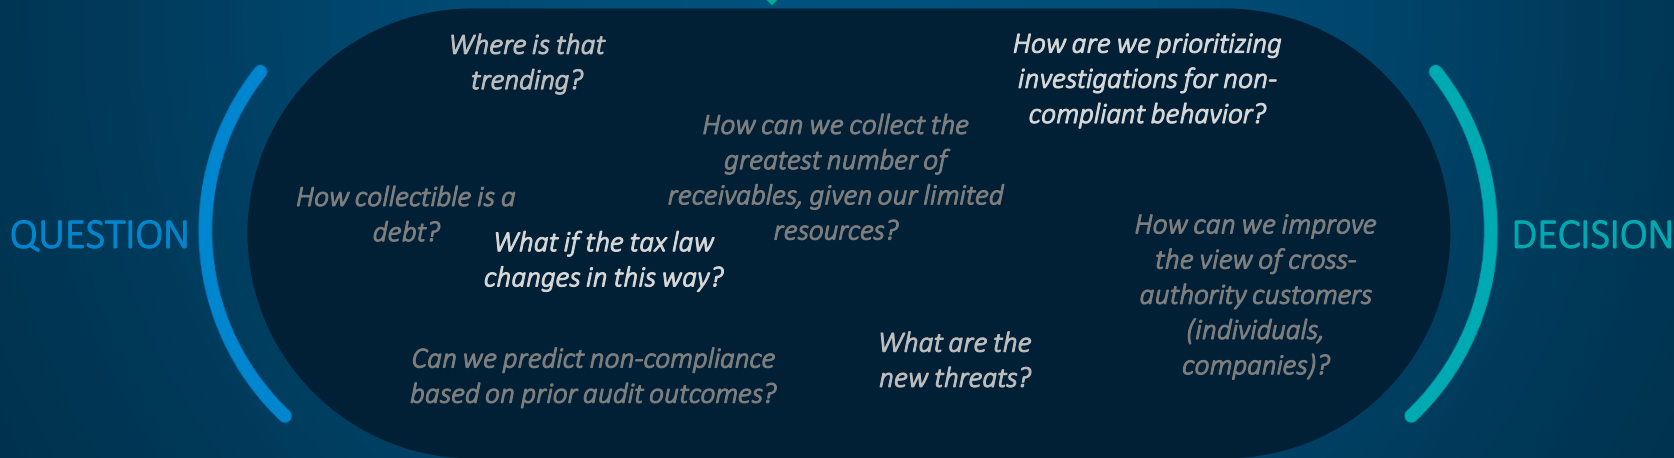
Where does my data strategy fit?

Governance...

- Aligned with Corporate Vision
- Depends on Corporate Strategy
- Define **DataOPs/ModelOPs** Implementation
- **Data Strategy:**
 - Processes and practices aligned with all necessary stakeholders
 - Assets (re)usage, from request to usage to retirement
 - Measures to assess Data Source Trustability



A data strategy *with analytics in mind* is a plan designed to improve all the ways you acquire, store, manage, share and use data to go from **QUESTIONS** to **DECISIONS**



A Day in the Life

1.



Commissioner

"Our accounts receivable balance is growing. How can we get better at prioritizing which cases are most/least collectable?"

"Cases with characteristics A, B, and C are 85% likely to resolve without us taking any action."

"Cases with characteristics X, Y, and Z have less than 5% chance of collection."



2.



Director of Research



"I have a lot of data on past cases that we have collected. How can I use historical data to determine collectability?"

3.



Business Analyst



Domain knowledge

Data Science

Coding

"What characteristics make a case collectable? How does time impact collectability?"

4.



Data Scientist



Domain knowledge

Data Science

Coding

"How do I explain the results of my model to a business user?"

5.



Director of Research



A Day in the Life

1.




Commissioner

"Our accounts receivable balance is growing:
 1. *Which cases are most/least collectable?*
 2. *What's the best action to take?"*
 3. *How can I maximize revenue with my limited people and budget?"*

"Here is a customized collection plan for each debtor. I can update it daily, based on new data received."

"X is your biggest bottleneck in the collections process. If you increase X by 25%, you can increase revenues by \$7.9 million in 6 months."



2.



Director of Research



"I have a lot of data on past cases that we have collected and actions we took. How can I use historical data to optimize collections, given the resource constraints?"

3.



Business Analyst



Domain knowledge



Data Science



Coding



"Collections is a process, not an isolated action. How do I create a customized action plan for each debtor?"

5.



Director of Research



"What algorithms best support the idea of 'collections is a process'? What are the most important constraints that prevent us from collecting more?"

4.



Data Scientist



Domain knowledge



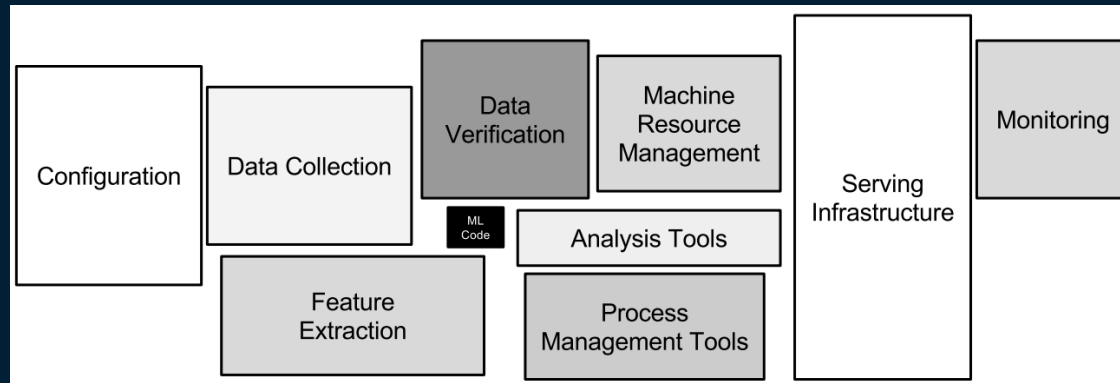
Data Science



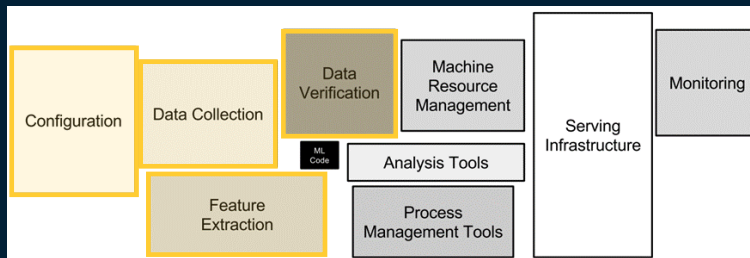
Coding



“ Only **a small fraction** of real-world Machine Learning systems **is composed of the Machine Learning code**, as shown by the small black box in the middle. The required surrounding infrastructure is vast and complex. ”



Ref: “Hidden Technical Debt in Machine Learning Systems”, Google Inc.



Ref: *"Hidden Technical Debt in Machine Learning Systems"*, Google Inc.

Developing a Data Strategy



Design with the end in mind

QUESTION

Where is that trending?

How collectible is a debt?

Can we predict non-compliance based on prior audit outcomes?

How can we collect the greatest number of receivables, given our limited resources?

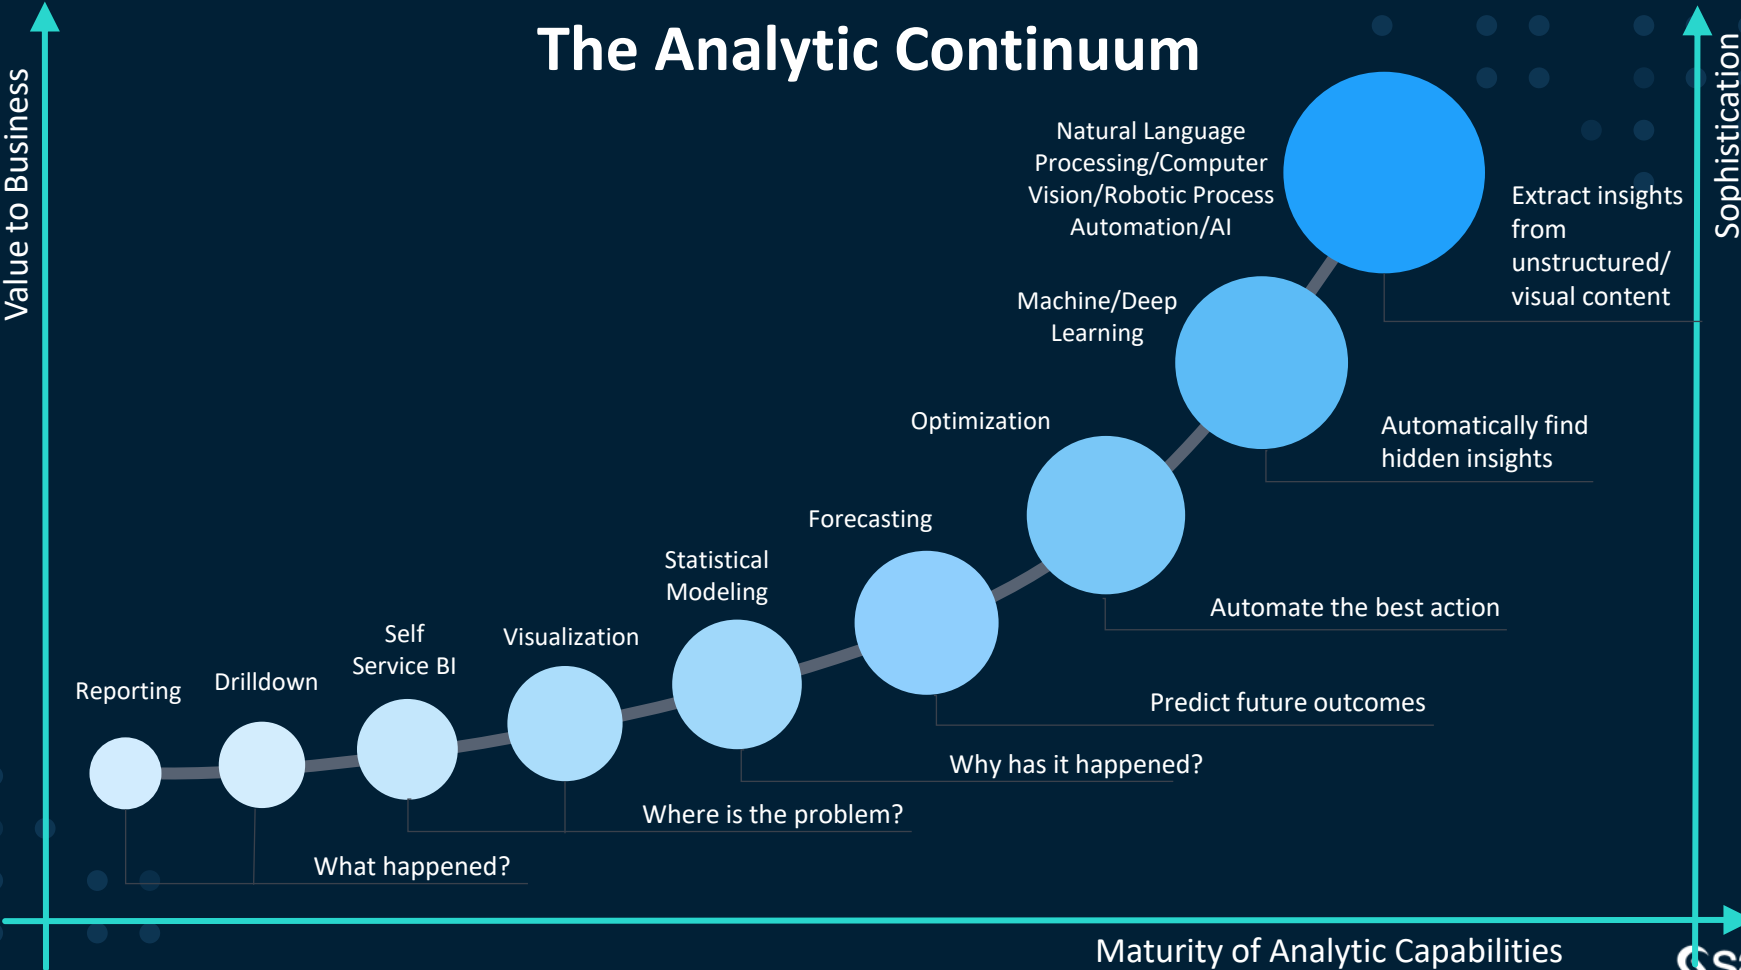
What are the new threats?

How are we prioritizing investigations for non-compliant behavior?

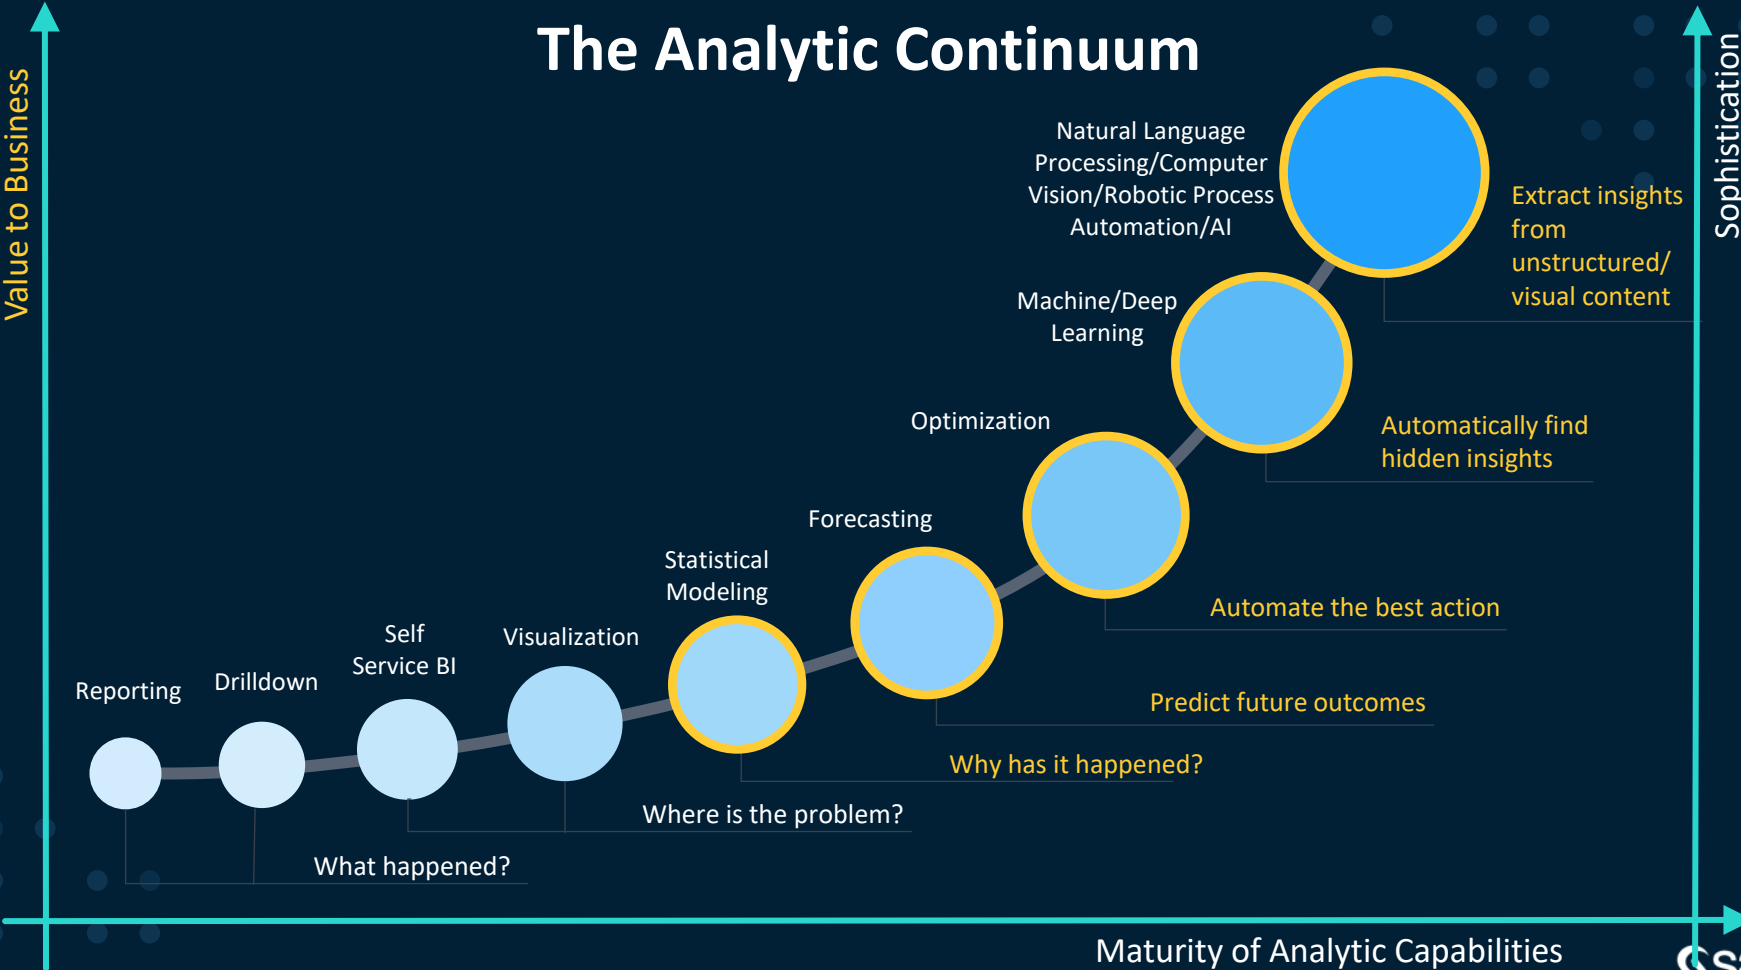
How can we improve the view of cross-authority customers (individuals, companies)?

DECISION

The Analytic Continuum



The Analytic Continuum



Introducing: The “Analytics Life Cycle”



The Analytics Life Cycle

Borrowing from agile software development practices, **DataOps** provides an agile approach to data access, quality, preparation, and governance. It enables greater reliability, adaptability, speed and collaboration in your efforts to operationalize data and analytic workflows.



Access

Access data, regardless of size or complexity



Prepare

Transform raw data, including AI powered suggestions



Visualize

View important relationships in data and share insights

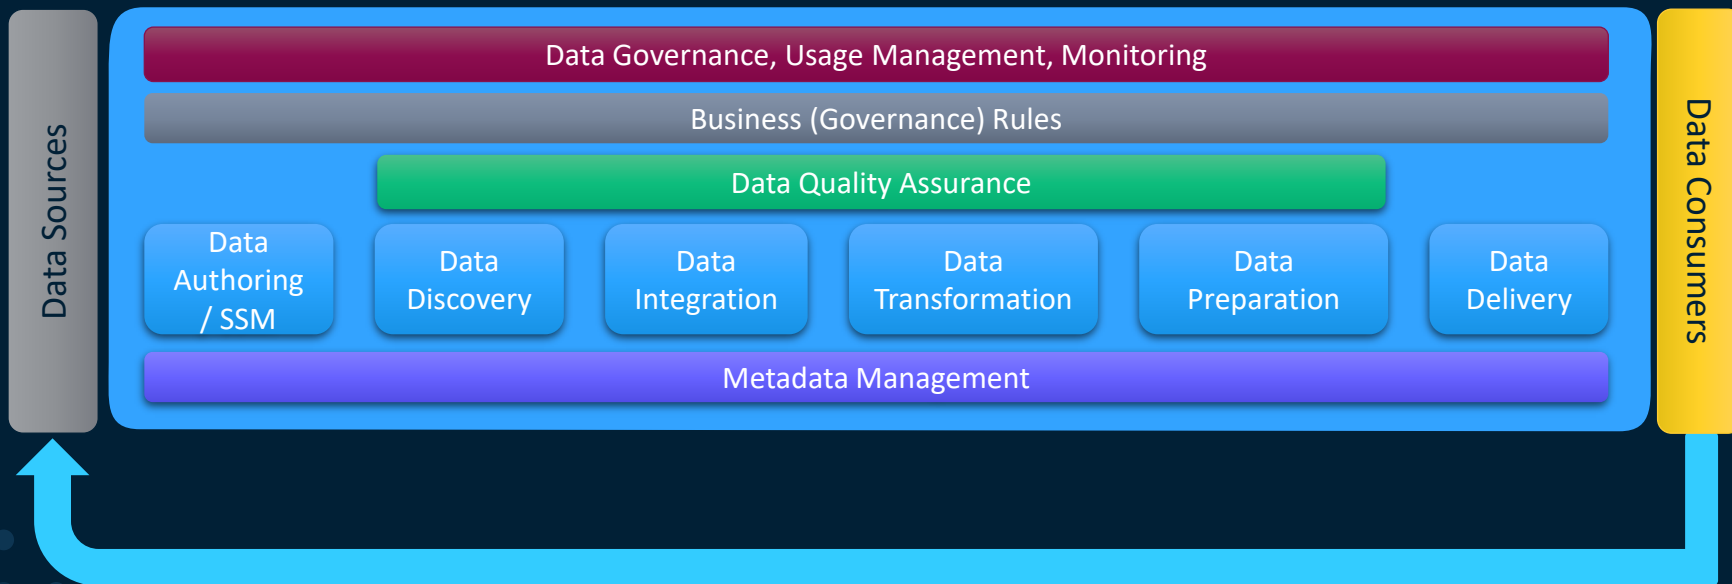


Govern

Build trust in data, understand lineage and gain transparency


DataOPs – “Operationalizing Analytics”

The **How** and the **What** ... From the Source to the Consumer




Data Preparation

The Need to “Flip the Script”



Data Preparation

The Need to “Flip the Script”



20%

80%

Discovery



Business Analyst
Citizen Data Scientist

Data Preparation



Self-service



Wrangling



Blending

flexible



Data



ETL Developer

Data Integration & Quality

Data Governance



IT Users



Extract,
Transform,
Load



Managed



Monitored



Glossary



Auditing



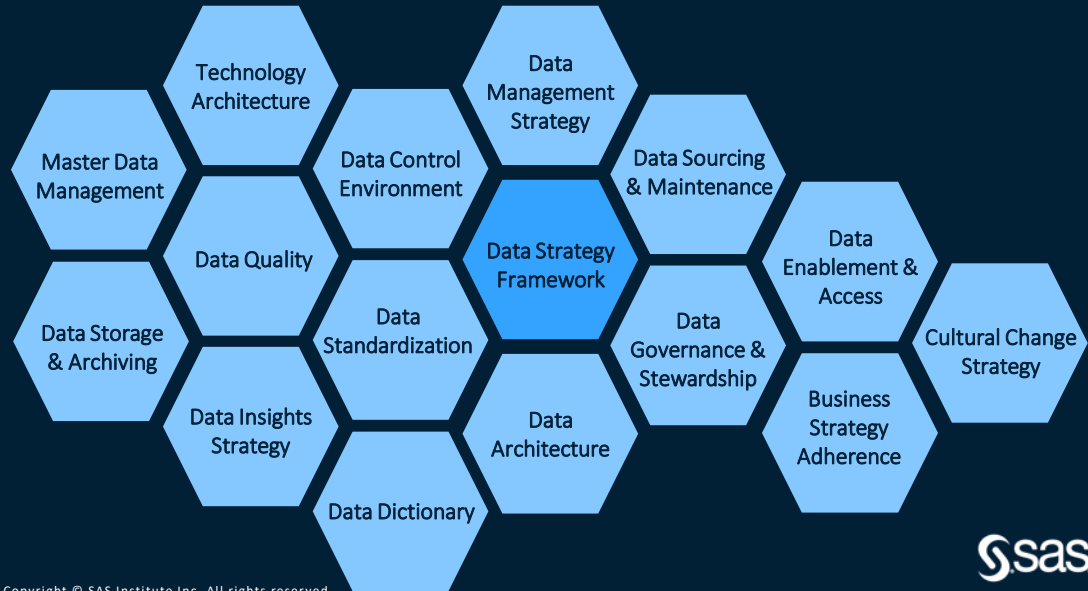
Lineage

Highly managed data
and optimized



Data Steward

A data strategy is a plan designed to improve all the ways you acquire, store, manage, share and use data



https://www.sas.com/content/dam/SAS/en_au/doc/whitepaper1/five-essential-components-data-strategy.pdf

AI Adoption in the Enterprise 2021


<https://www.oreilly.com/radar/ai-adoption-in-the-enterprise-2021/>

- In the past, **company culture** has been the most significant barrier to AI adoption. While it's still an issue, culture has dropped to fourth place.
- The biggest problems in this year's survey are lack of skilled people and difficulty in hiring (19%) **and data quality (18%)**. It's no surprise that the demand for AI expertise has exceeded the supply, but it's important to realize that it's now become the biggest bar to wider adoption. The biggest skills gaps were ML modelers and data scientists (52%), understanding business use cases (49%), and **data engineering (42%)**. The need for people managing and maintaining computing infrastructure was comparatively low (24%), hinting that companies are solving their infrastructure requirements in the cloud.

O'REILLY Learning

<https://learning.oreilly.com/topics/enterprise-data-strategy/>

The screenshot shows the O'Reilly Learning website interface for the 'Enterprise Data Strategy' topic. The top navigation bar includes the O'Reilly logo, a search bar, and a user profile icon. The main content area is titled 'Enterprise Data Strategy' with a 'Follow this topic' button. Below the title, there are filters for 'All Topics > Data > Data Science > Enterprise Data Strategy', a search input, and dropdown menus for 'Sort order: Popularity' and 'Formats: All Formats'. A grid of four items is displayed: 'The Self-Service Data Roadmap' (book), 'Creating a Data-Driven Organization' (book), 'The Chief Data Officer Management Handbook' (book), and 'Learning Path: Developing a Modern Enterprise Data Strategy' (video). A 'Topics You Follow' section is partially visible at the bottom.



WHITE PAPER

Top Trends: Why Tax Administrators Are Adopting New Data and Analytics Strategies

presented by SAS

<https://www.sas.com/en/whitepapers/top-trends-tax-administrators-111564.html>

Paper SAS4617-2020
Smarter and Faster Self-Service Data Preparation
Atrin Assa, SAS Institute Inc.

ABSTRACT

In this paper, you'll learn about the latest and greatest self-service data preparation capabilities of SAS® Visual Analytics. You will understand how smart suggestions can help you improve the quality of your data, how the new interface can help you work faster, and how better-prepared data can help you build better visualizations, better reports, and tell a more compelling data story.

<https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4617-2020.pdf>

Creating a Data Analytics Function

1

Developing a data strategy

2

Where to locate the data analytics function

3

Centralized vs de-centralized

4

Stakeholder collaboration

5

Capabilities and education

6

Experimentation & learning



Making it easier for
everyone to use analytics
to ask and answer the
questions that drive
business forward

Where to Locate...*Fit for Purpose*

Why are we building a model?

- How does that model start generating value?
 - What are your business requirements?
 - What are your metrics for success?

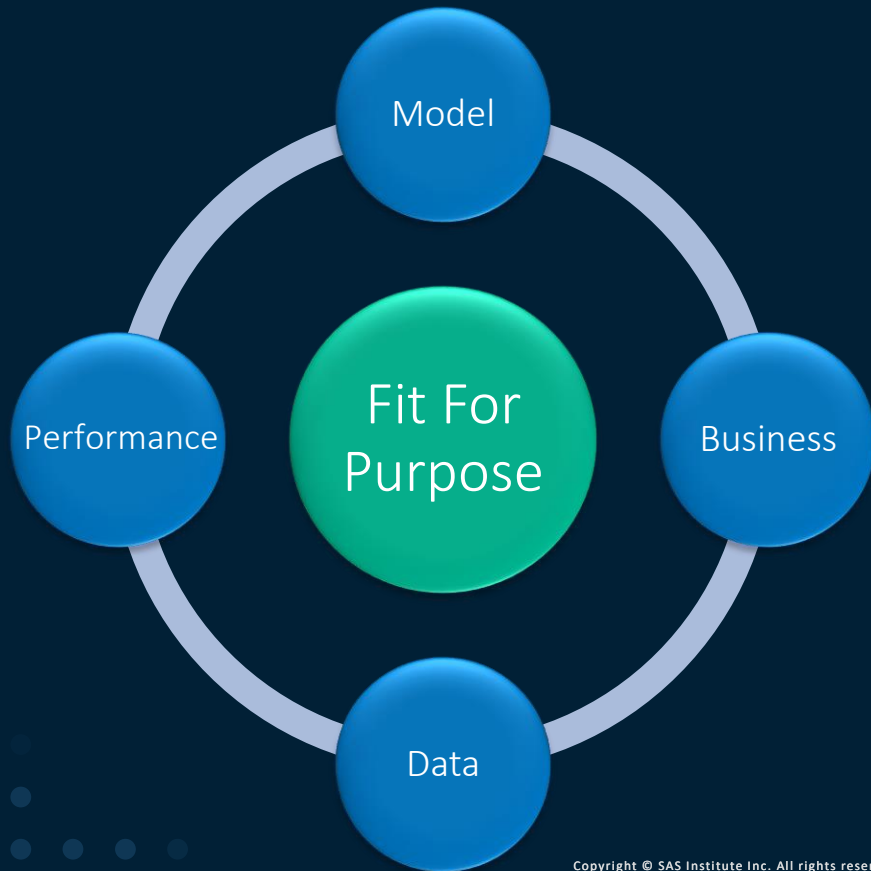
Model deployment (or serving) is operationalization of the model and where it starts to generate value.

- The statistically best model might be the model you should deploy, but what about:
 - How fast does it execute?
 - What data does it use?
 - Is it reflecting your business goals?

“Fit for Purpose” means considering your business objectives and technical requirements alongside model accuracy.

Where to Locate...*Fit for Purpose*

What makes a model fit for its purpose?



Define all requirements before embarking on an analytical model building project.

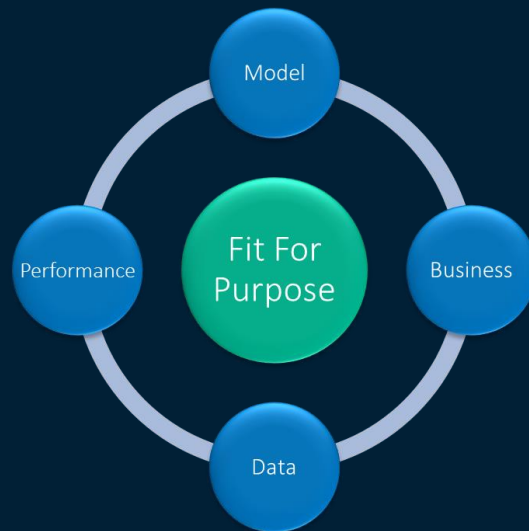
This means it is a collaborative process to define the requirements of the project, involving businesses, analytics and IT.

Where to Locate...*Fit for Purpose*

What makes a model fit for its purpose?

The Machine Learning Canvas helps to identify business, data, model, execution challenges before the start of the project (i.e., fail fast principle).

Data Sources Which raw data sources can we use (internal and external) and what is the effort to access them? How much (historical) data do we have?		Value Propositions What objectives are we serving? What is the business potential (in €'s, savings, efficiency, etc.)? Who are the stakeholders? Value of correct decisions versus cost of incorrect decisions?	Collecting Data Which data sources are available in an operational context? Streaming data or off line data? Which calculation should be done in stream / off line?
Data pre-processing & featurization What is the quality of the data (e.g., missing values)? Which features are extracted from the raw data sources?	Analytical task Inputs? Output? Type of problem?	Making Predictions When do we make predictions on new inputs? How much time do we have to make a prediction? Volume of predictions to make?	Decisions How are data and predictions used to make decisions that provide the proposed value to the end user? How much time do we have to make a decision?
Offline Evaluation Methods and metrics to evaluate the system before deployment? Which are the success criteria to deploy the system?	Live Evaluation and Monitoring Methods and metrics to evaluate the system after development, and how to quantify the value creation.	Building Models When do we create/update models with new training data? How long do we have to featurize training inputs and create a model?	



Fit for Purpose Environment

INTERACTIVE
DASHBOARDS



REGISTER

Provide oversight and governed analytics.



DEPLOY

Embed and execute models, leverage different engine based on the business need



MONITOR

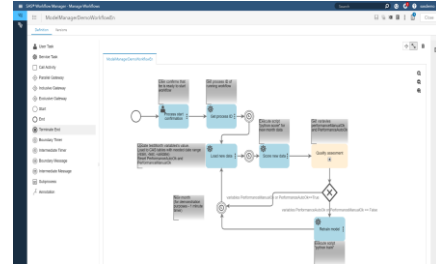
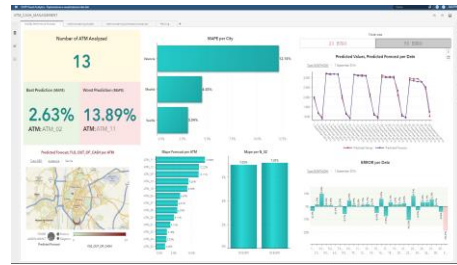
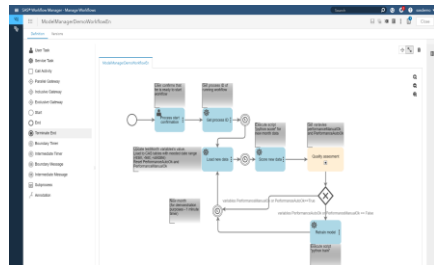
Gain visibility into analytic outcomes and approach.



RETRAIN

Models using pre-defined workflow To execute training jobs

Event	Project name	Model	Deployment	Modified
Deployment	1.1	model01	Aug/11, 2017 10:26 PM	active
Model build	1.1	Deployment01	Aug/11, 2017 10:26 PM	active
Model build	1.1	model01	Aug/11, 2017 10:26 PM	active
Model build	1.1	Deployment01	Aug/11, 2017 10:26 PM	active
Deployment	1.1	model01	Aug/11, 2017 10:26 PM	active
Deployment	1.1	model01	Aug/11, 2017 10:26 PM	active



MODEL
REPOSITORY



WORKFLOW



ModelOps (MLOps) White Paper

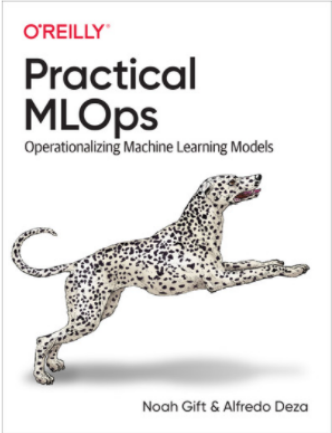


BOOK

Practical MLOps

★★★★☆ 6 reviews

By [Noah Gift, Alfredo Deza](#)



O'REILLY
Practical MLOps
Operationalizing Machine Learning Models
Noah Gift & Alfredo Deza

TIME TO COMPLETE:
11h 38m

TOPICS:
[Machine Learning](#)

PUBLISHED BY:
[O'Reilly Media, Inc.](#)

PUBLICATION DATE:
September 2021

PRINT LENGTH:
458 pages

Continue

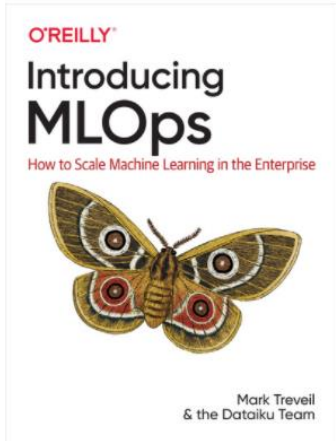
<https://learning.oreilly.com/library/view/practical-mlops/9781098103002/>

BOOK

Introducing MLOps

★★★★☆ 10 reviews

By [Mark Treveil, Nicolas Omont, Clément Stenac, Kenji Lefevre, Du Phan, Joachim Zentgraf](#)



O'REILLY
Introducing MLOps
How to Scale Machine Learning in the Enterprise
Mark Treveil & the Dataiku Team

TIME TO COMPLETE:
5h 9m

TOPICS:
[Machine Learning](#)

PUBLISHED BY:
[O'Reilly Media, Inc.](#)

PUBLICATION DATE:
November 2020

PRINT LENGTH:
183 pages

Start

<https://learning.oreilly.com/library/view/introducing-mlops/9781492083283/>

ModelOPs White Paper

SAS Global Forum: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2019/3368-2019.pdf>

SAS3368-2019

Deploying Models Using SAS® and Open Source

Jared Dean, SAS Institute Inc.

ABSTRACT

In the excitement and hype around machine learning (ML) and artificial intelligence (AI) most of the time is spent in the model building. Much less energy is expended on how to take the insights from models and deploy them efficiently to create value and improve business outcomes.

This paper will show a complete example using DevOps principals for building models and deploying them using SAS® in conjunction with opens source projects including Docker, Flask, Jenkins, Jupyter, and Python. The reference application is a recommendation engine on a web property with a global user base. This use case forces us to confront security, latency, scalability, repeatability. The paper will outline the final solution but also include some of the problems encountered along the way that informed the final solution.

ModelOPs White Paper

SAS Global Forum: <https://www.sas.com/content/dam/SAS/support/en/sas-global-forum-proceedings/2020/4402-2020.pdf>



SAS[®] GLOBAL FORUM 2020

#SASGF

Paper SAS4402-2020

Open-Source Model Management with SAS[®] Model Manager

Glenn Clingroth, Hongjie Xin, and Scott Lindauer, SAS Institute Inc.

ABSTRACT

Open-source models that are developed in Python, R, TensorFlow, and so on, are increasingly important to organizations that produce and deploy analytical and machine learning models. Not only are the models created using open-source tools, they are deployed to open-source environments that use Docker and Kubernetes in place of more traditional environments. SAS[®] Model Manager is evolving to be a management platform that handles traditional SAS models and open-source models as equal partners. This paper discusses strategies for managing the life cycles of Python, R, and TensorFlow models using SAS Model Manager.

Creating a Data Analytics Function

Overview

1

Developing a data strategy

2

Where to locate the data analytics function

3

Centralized vs de-centralized

4

Stakeholder collaboration

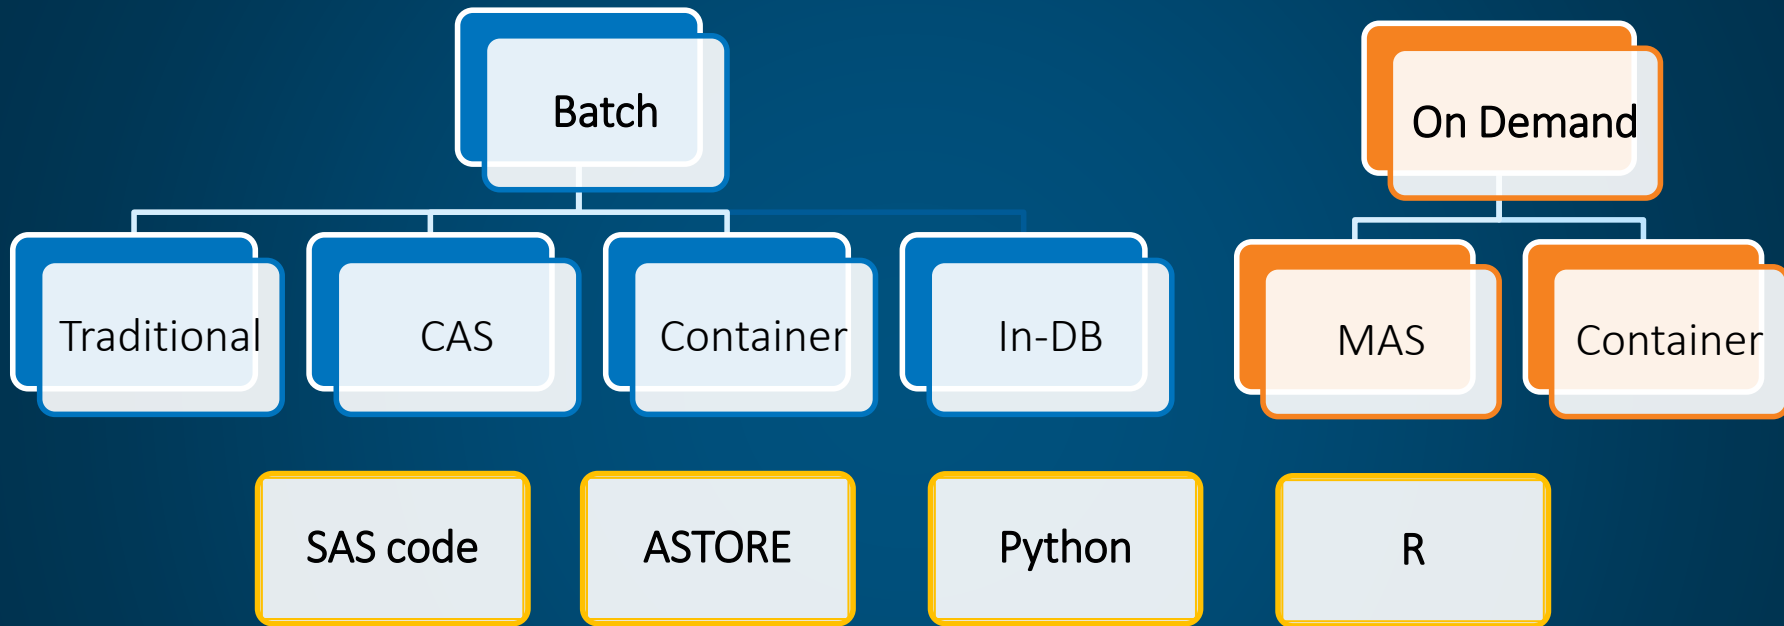
5

Capabilities and education

6

Experimentation & learning

The deployment method is driven by the *consuming application*



CASE STUDY:

Problem: Every project at the tax agency addressed data issues as one-off, built-from scratch activities. The agencies IT team had 17 projects underway (new applications, application enhancements, new reports, etc.).

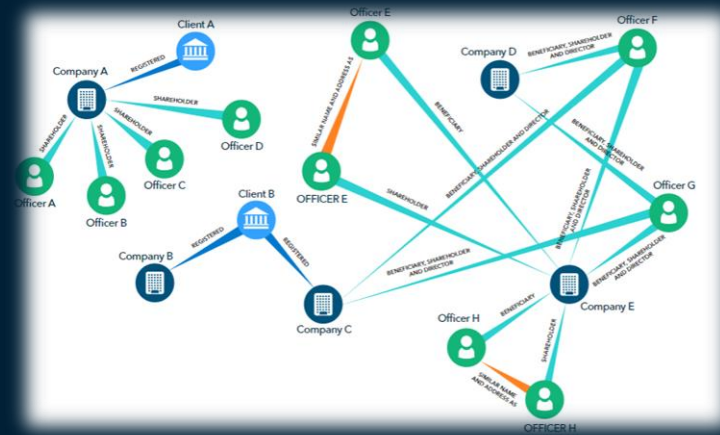
How does an integrated “centralized” analytics platform help?

- A single, integrated analytics platform incorporates data management techniques to ingest disparate data sets from internal sources (tax data) and external sources (other agencies, third parties, corporate data, etc.) – then the data is blended and cleansed before it’s used with analytics.
- All tax agents use the same data-based foundation for making decisions – with key information displayed visually in a way that’s easy to configure, search and consume
- Behavioral analytics identifies key types of fraud entities such as “shell company” – as well as patterns in behavior – to uncover connections among entities

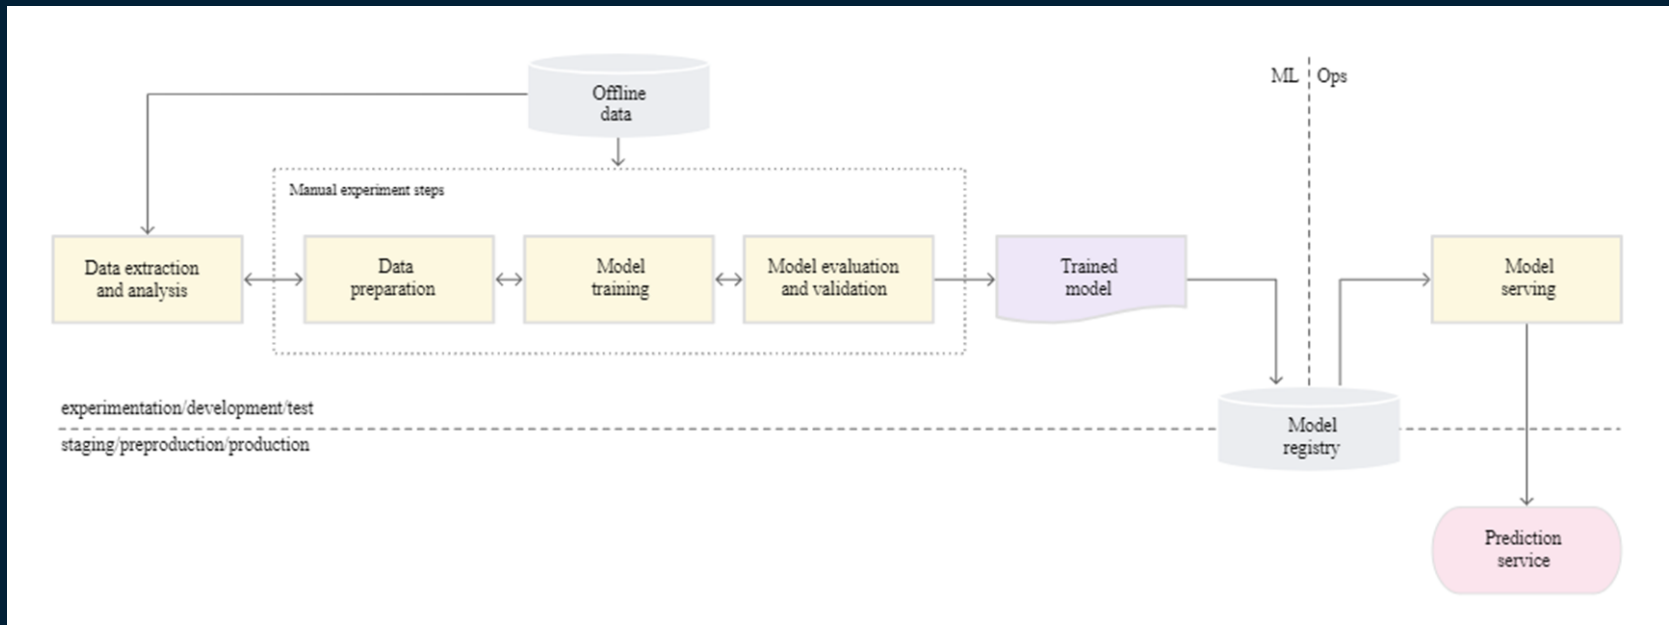
CASE STUDY:

Problem: Every project at the tax agency addressed data issues as one-off, built-from scratch activities. The agencies IT team had 17 projects underway (new applications, application enhancements, new reports, etc.).

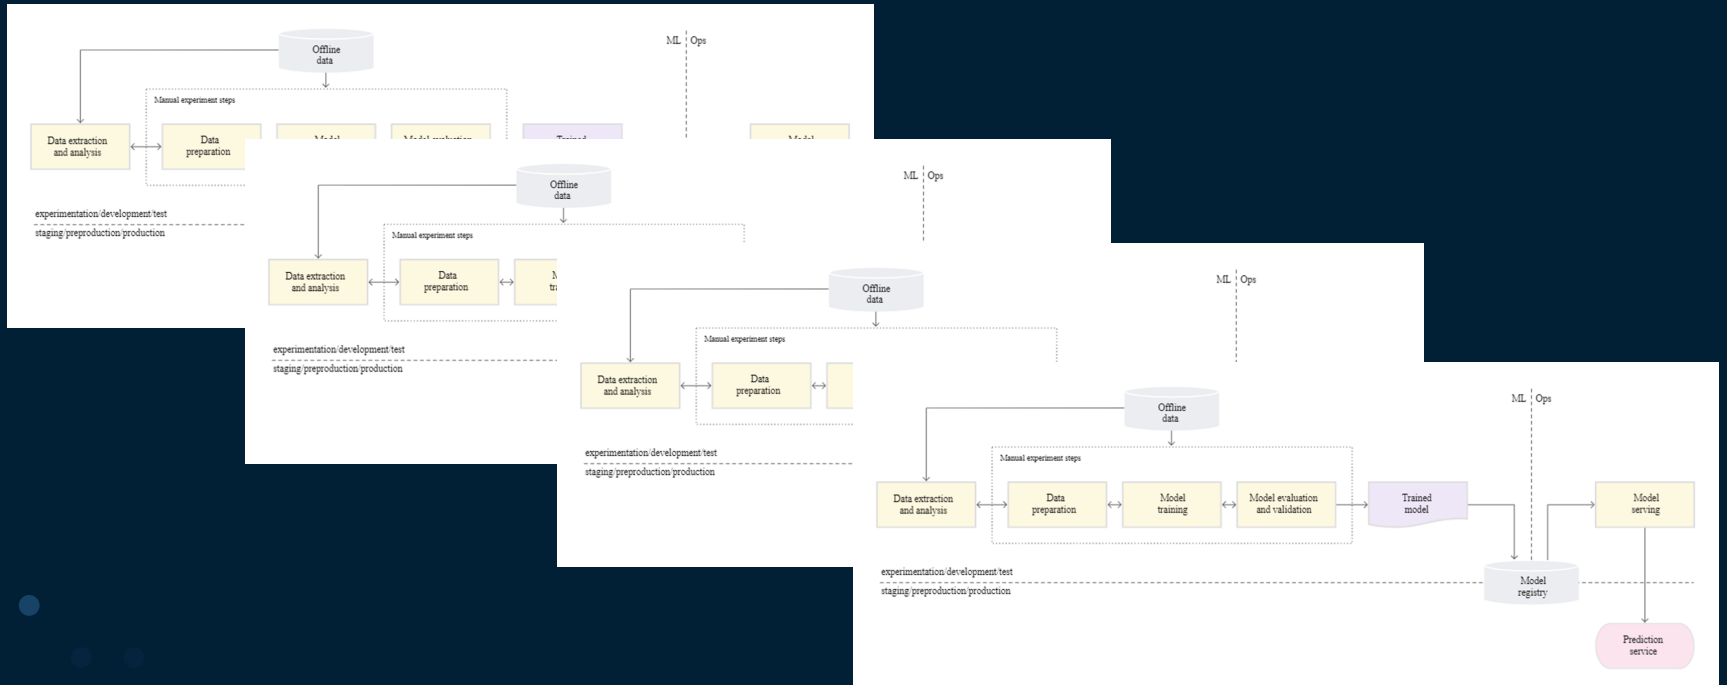
- Alert generation, scoring and risk analysis target the most appropriate cases for tax agents to investigate.
- Hybrid analytics approach blends social network analysis, anomaly detection and other analytical techniques to deliver the best possible results



High-Level “*De-Centralized*” Operating Model for Developing Data Analytic Solutions

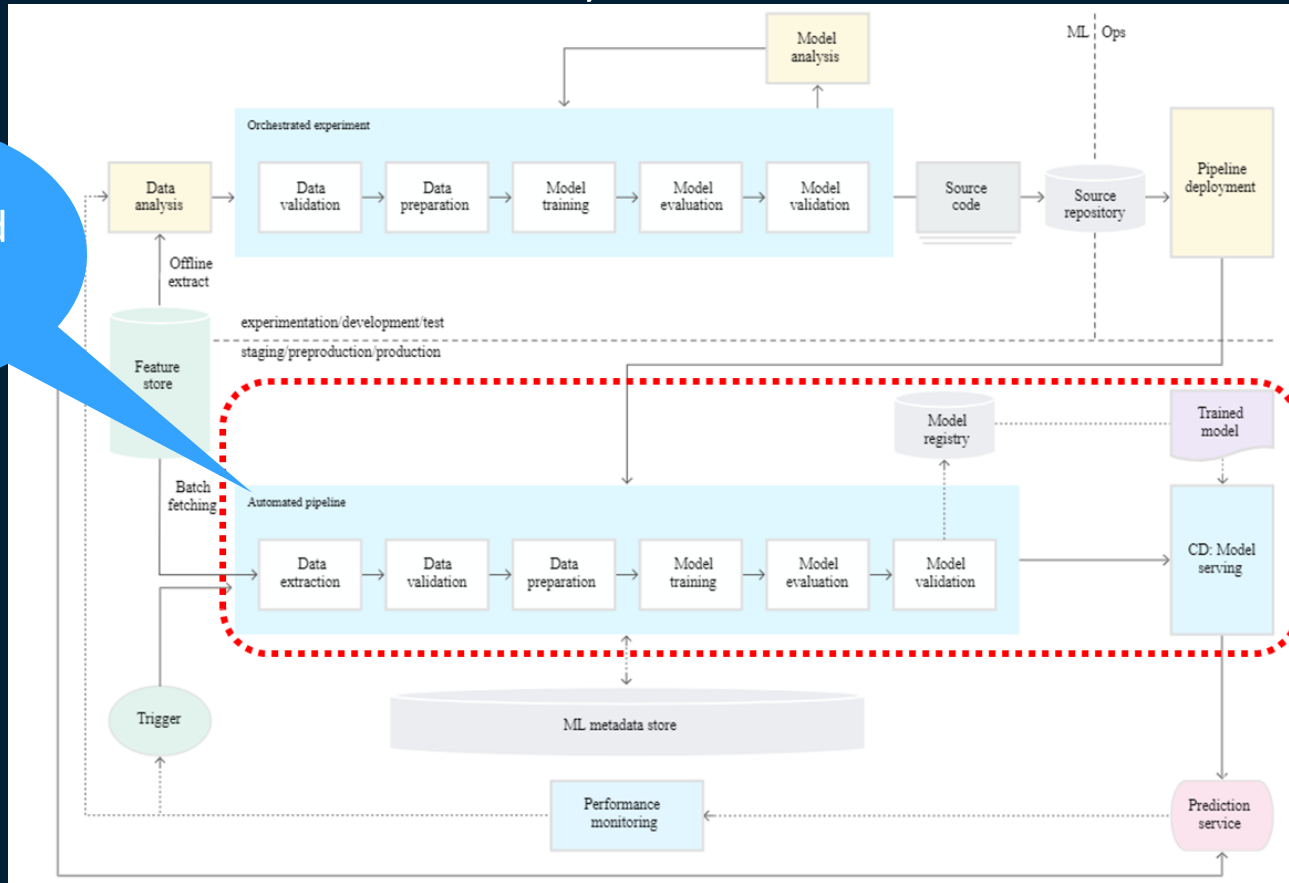


High-Level “*De-Centralized*” Operating Model for Developing Data Analytic Solutions



High-Level “*Centralized*” Operating Model for Developing Data Analytic Solutions

Automated Pipelines



Operationalizing Analytics

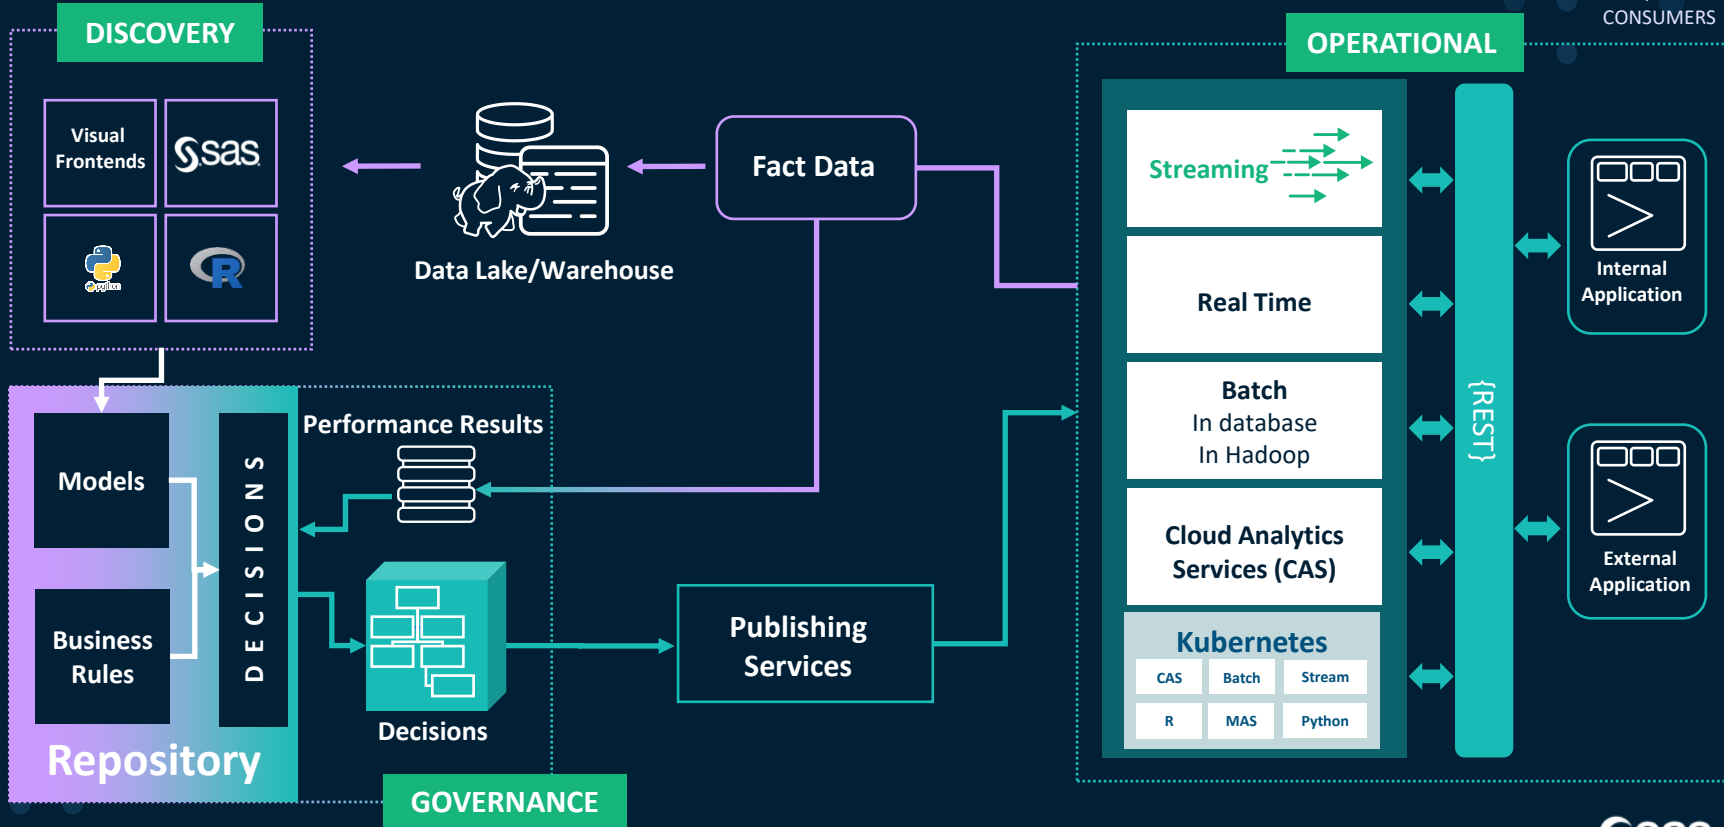
An Integrated “*Centralized*” Analytics Platform



MODEL/DECISION PRODUCERS

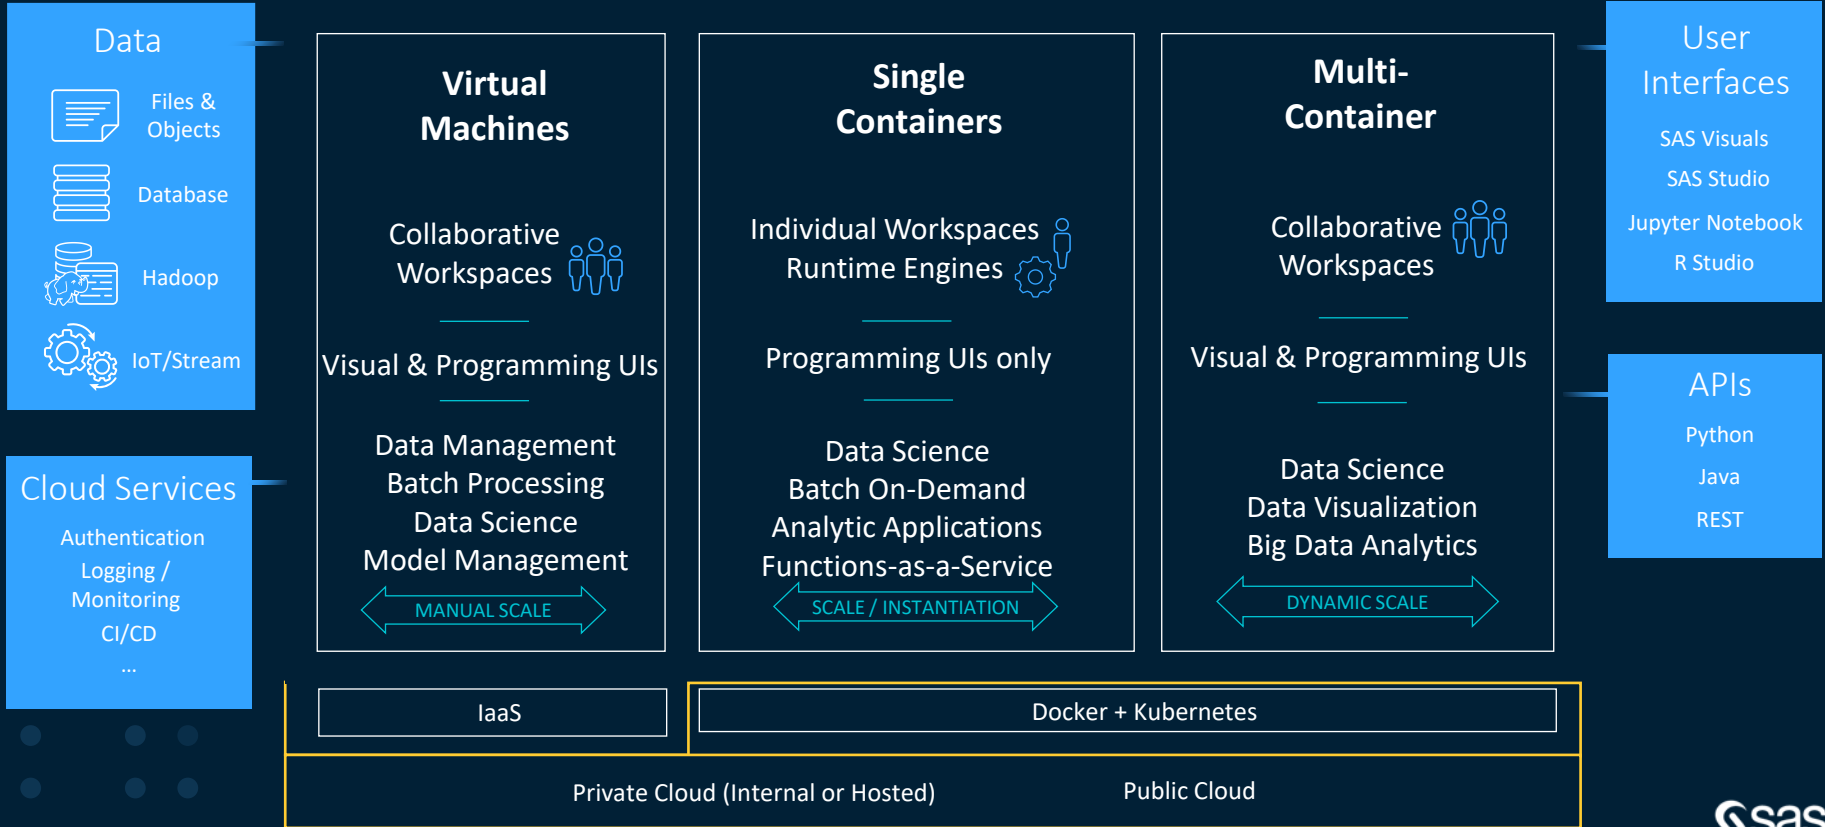



MODEL/DECISION CONSUMERS

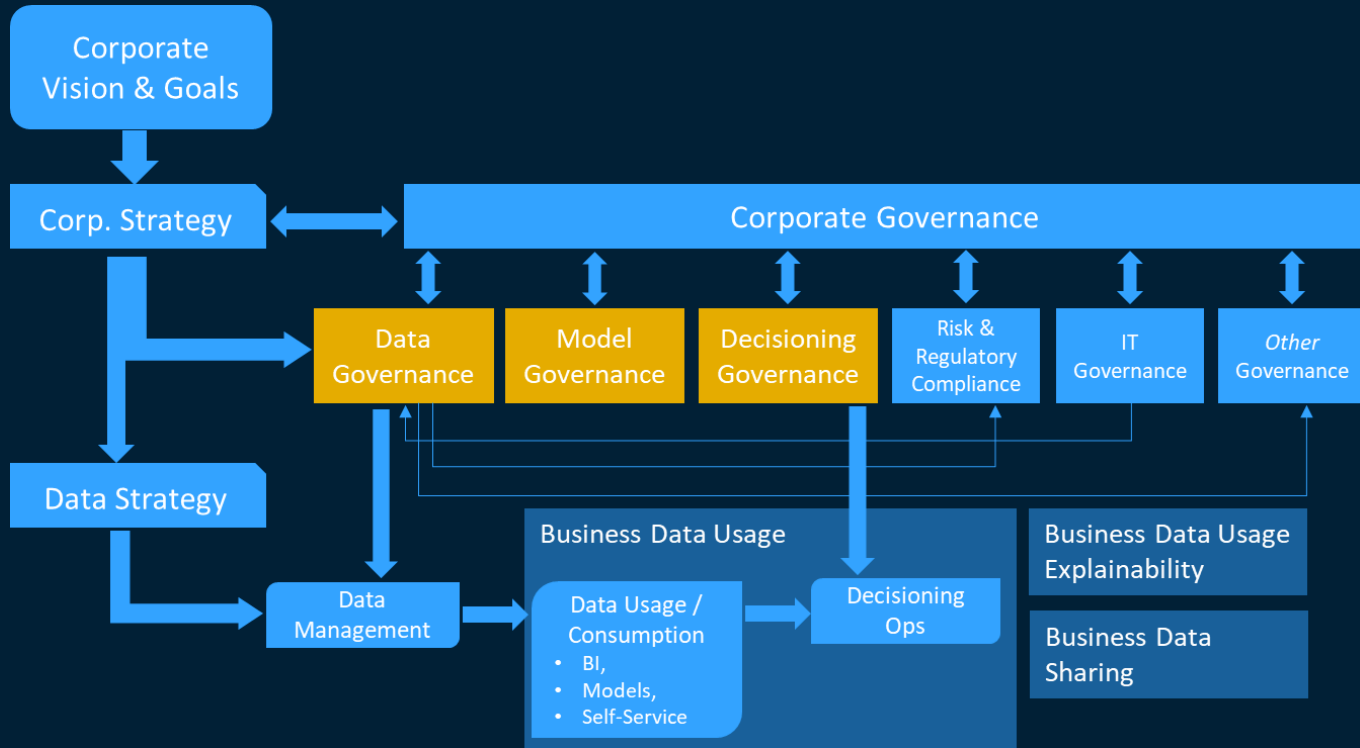


MODEL/DECISION GOVERNORS

Using Cloud / Containers to *Centralize*




“Centralized” Analytics Platform...Needs Governance



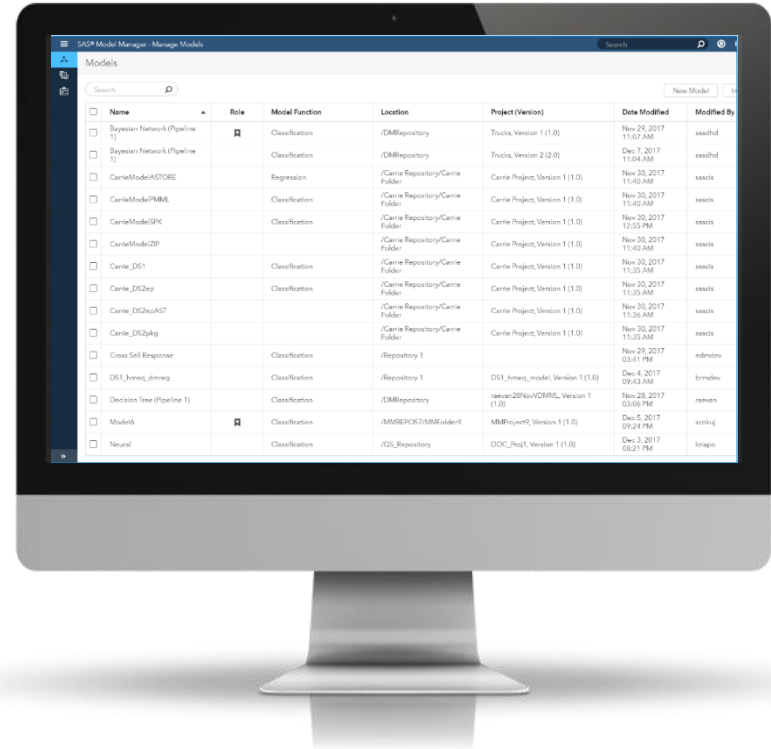
Analytical Model Governance

Centrally Manage ALL Analytic Models



Interoperability with third-party modeling tools to form a model management hub

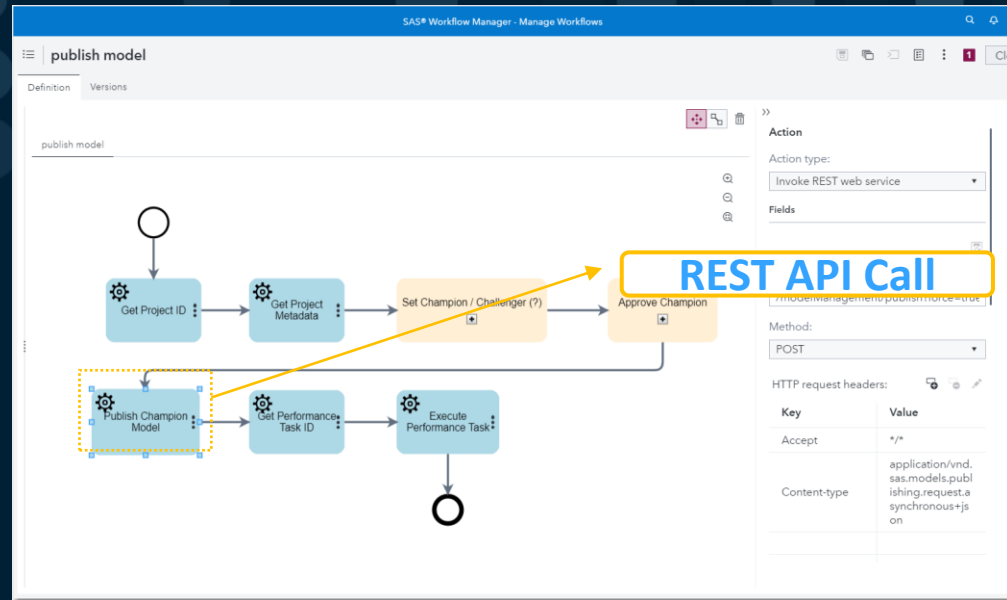
- Monitor ALL models in production to maximize business value
- Govern ALL analytic models
- Share & collaborate
- Trace & prove with workflow and versioning



Centralized can “push” and “pull” with REST API integration

Executing steps in different backends wherever it makes the most sense

- Execute business logic or rules
- Communicate with other processes using web services or Job Execution tasks
- Integrate the process with third-party applications



DASHBOARD: Continuous view of the model repository to the stakeholders

Models

Hide Dashboard


83

Total Number of Models




26 Published Models

Models per Score Code Type



- DATA step
- DS2 embedded process
- DS2 multi-type
- DS2 package
- Not specified
- Python
- SAS program

Published Models per Destination



- AWSTest_0526
- AzureTest_0526
- CASDest
- HadoopDest
- HarborTest0
- HarborTest0WF
- maslocal
- PrivateDocker0526
- TeradataDest

Search name

Name	Role	Project (Version)	Model Attributes	Model Function	Modified By	Date Modified
<input type="checkbox"/> Neilsen_VA_LinearRegression2			DATA step	Prediction	dishaw	May 29, 2020 06:18 PM
<input type="checkbox"/> Neilsen_VA_LinearRegression			DATA step	Prediction	dishaw	May 29, 2020 05:50 PM
<input type="checkbox"/> Decision Tree (Intermediate Template Approach)		Fleet Management Report_Version 1(1.0)	DATA step	Classification	dishaw	May 29, 2020 12:02 PM
<input type="checkbox"/> FleetTest_XGBoostNew		Fleet Management Report_Version 1(1.0)	Published Python	Classification	dishaw	May 29, 2020 12:01 PM
<input type="checkbox"/> HMEQ_Python_ClassTree		Home Mortgage Equity_Version 1(1.0)	Python	Classification	dishaw	May 29, 2020 10:00 AM
<input type="checkbox"/> Linear_Regression - SAS		Fleet Trip Time_Version 1(1.0)	DATA step	Prediction	dishaw	May 29, 2020 09:39 AM
<input type="checkbox"/> QS_Reg1		Lisa Project DS1_Version 1(1.0)	Published DATA step	Classification	edmddev	May 29, 2020 01:39 AM
<input type="checkbox"/> QS_Tree1		Lisa Project DS1_Version 1(1.0)	DATA step	Classification	edmddev	May 29, 2020 01:36 AM
<input type="checkbox"/> QS_Reg_PyModel		Michael python_Version 1(1.0)	Published Python	Classification	edmddev	May 28, 2020 11:08 PM
<input type="checkbox"/> QS_Reg_PyModel		azuretesting_Version 1(1.0)	Python	Classification	edmddev	May 28, 2020 01:39 PM
<input type="checkbox"/> QS_Reg_PyModel_P1		Shane Project 1_Version 1(1.0)	Published Python	Classification	brmddev	May 28, 2020 12:19 PM
<input type="checkbox"/> Gradient Boosting (Pipeline 1)		Project Michael 0527_Version 1(1.0)	Published DS2 multi-type	Classification	scnxic	May 27, 2020 03:21 AM
<input type="checkbox"/> hpsvm		Project Michael 0527_Version 1(1.0)	DS2 multi-type	Classification	scnxic	May 27, 2020 03:16 AM



Analytic / Data Architecture Governance

O'REILLY : <https://learning.oreilly.com/home/>

analytic architecture governance

1 - 10 of 691 search results for "analytic architecture governance"

All Books Videos Learning Paths

Topics

BOOK

Information Governance, 2nd Edition
By Robert F. Smallwood
GOVERNANCE

★★★★★ 1 review

Wiley December 2019

INFORMATION GOVERNANCE

Titles in the Wiley CIO series include: The Agile **Architecture** Revolution: How Cloud Computing, REST-Based SOA, and Mobile Computing Service Models (SaaS, PaaS, and IaaS) by Michael Kavis Big Data, Big Analytics: Emerging Business Intelligence and...

BOOK

Self-Service Data Analytics and Governance for Managers
By Nathan E. Myers and Gregory Kogan
GOVERNANCE

Write the first review

Wiley June 2021

CHAPTER 4: Self-Service Data Analytics Project Governance

"Analytic Models Controls and Tests - Auditors can perform numerous tests to provide assurance on analytics controls used 129/Accenture-Data-is-the-New-Capital-POV.pdf Banks and bank holding companies, for example, are under keen pressure to

BOOK

Performing Information Governance: A Step-by-Step Guide to Making Information Governance Work
By Anthony David Giordano
GOVERNANCE

OREILLY

Data Governance: The Definitive Guide
People, Processes, and Tools to Operationalize Data Trustworthiness

Evren Eryurek, Uri Gilad, Valliappa Lakshmanan, Anita Kibunguchi-Grant & Jessi Ashdown

BOOK

MAR 2021

Data Governance: The Definitive Guide
By Evren Eryurek, Uri Gilad, Valliappa Lakshmanan, Anita Kibunguchi-Grant and Jessi Ashdown

OREILLY

VIDEO

JAN 2016

Data Governance
By John Adler

BOOK

NOV 2012

Data Governance
By John Ladley

BOOK

DEC 2010

MASTER DATA MANAGEMENT AND DATA...
By Alex Berson and Larry Dubov

DATA GOVERNANCE
Creating Value from Information Assets

OREILLY

Understanding Data Governance

THE DATA GOVERNANCE IMPERATIVE
A business strategy for corporate data

IBM

Information Governance Principles and Practices for a Big Data Landscape

Creating a Data Analytics Function

10-minute break



4

Stakeholder collaboration

5

Capabilities and education

6

Experimentation & learning



Creating a Data Analytics Function

Overview

1

Developing a data strategy

2

Where to locate the data analytics function

3

Centralized vs de-centralized

4

Stakeholder collaboration

5

Capabilities and education

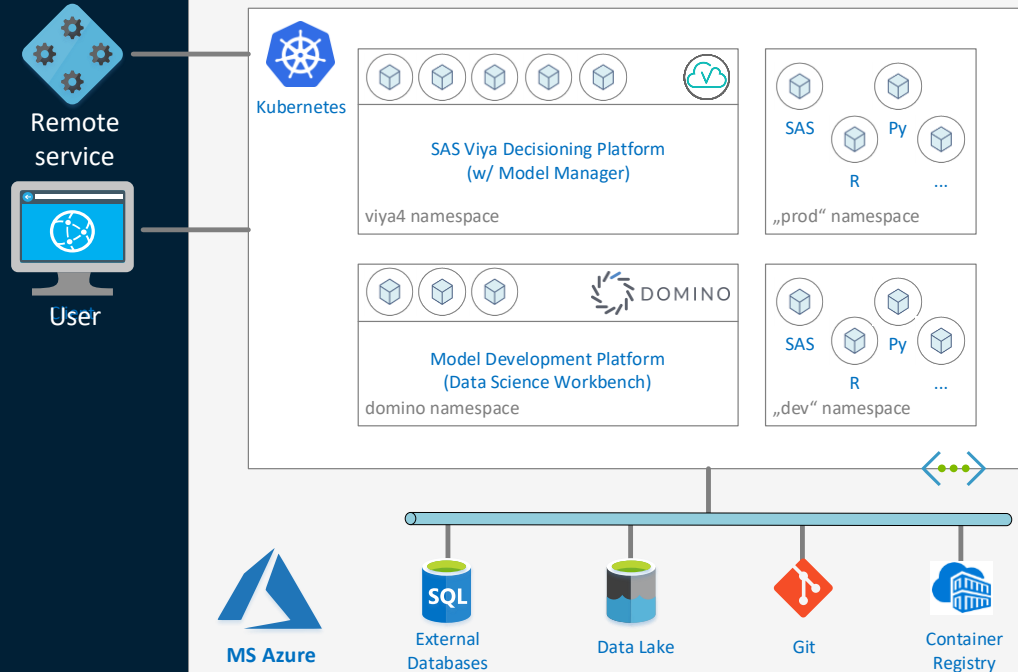
6

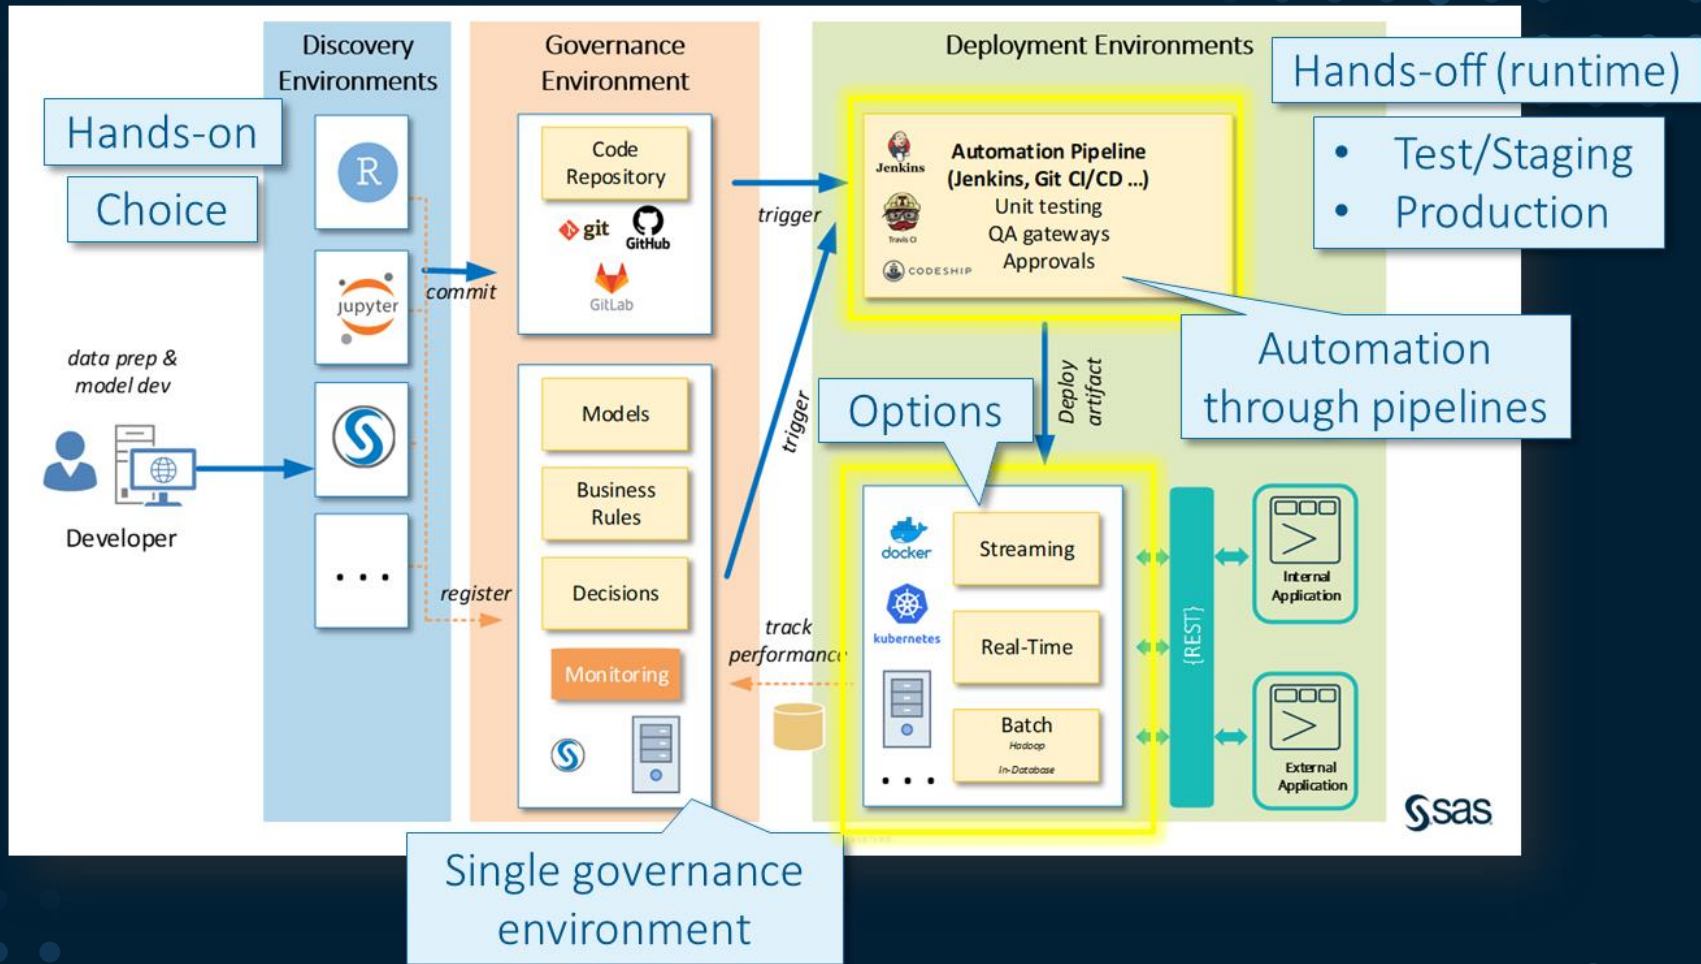
Experimentation & learning

Data Science Workbench

Example

- Cloud deployment (MS Azure)
 - All components deployed on managed Kubernetes cluster (AKS)
 - Connect to cloud data sources as needed
- Data Science Workbench provides Python, SAS, ... environments as Docker containers on demand
 - Scheduled to run in “dev” area
- From here, models are registered to Decisioning Platform and deployed to “prod” area using SAS Model Manager





Operationalizing Analytics

The Next Frontier

