




Data Driven Innovation: **Modern Data Platform as Foundation for Transformation and New Services**

Karsten Haldenwang, SAP

Public

September 2017



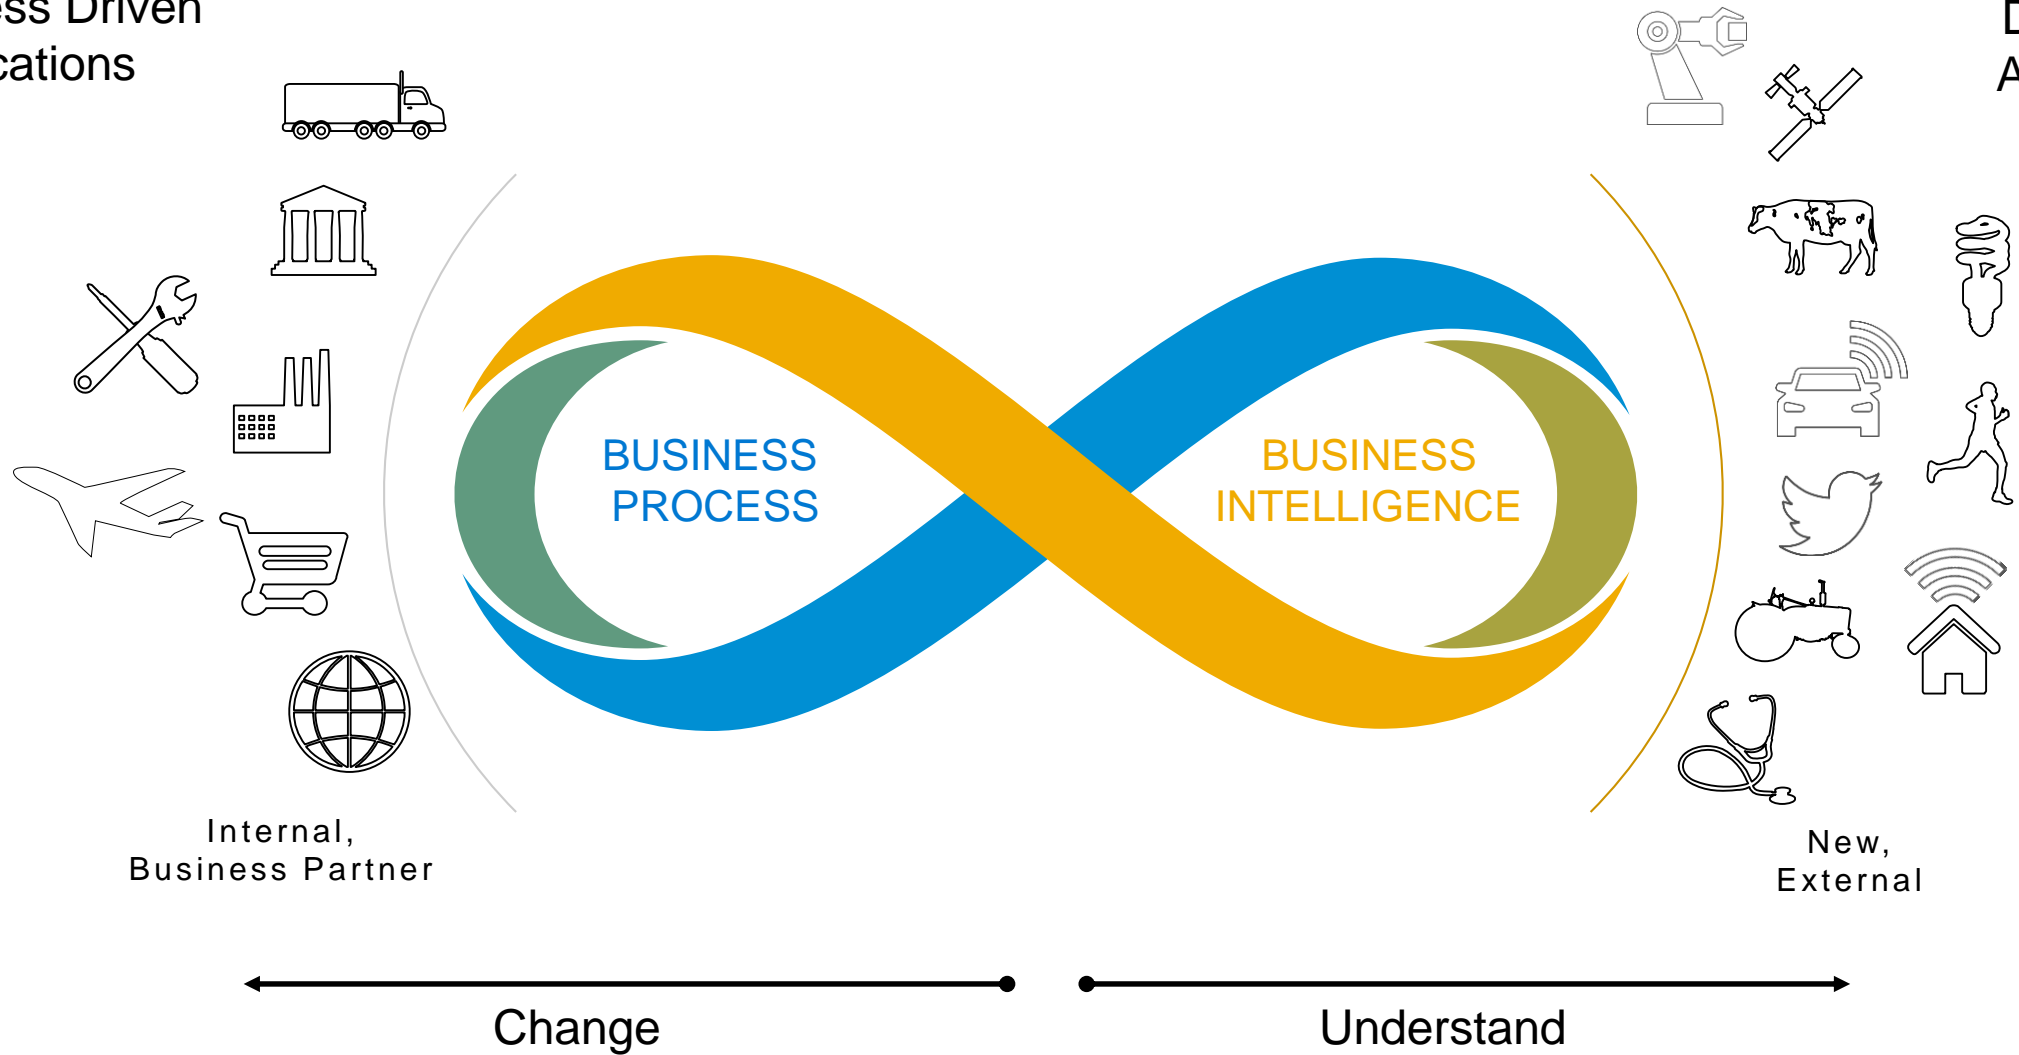
DATA

Connected Data as a Strategic Asset

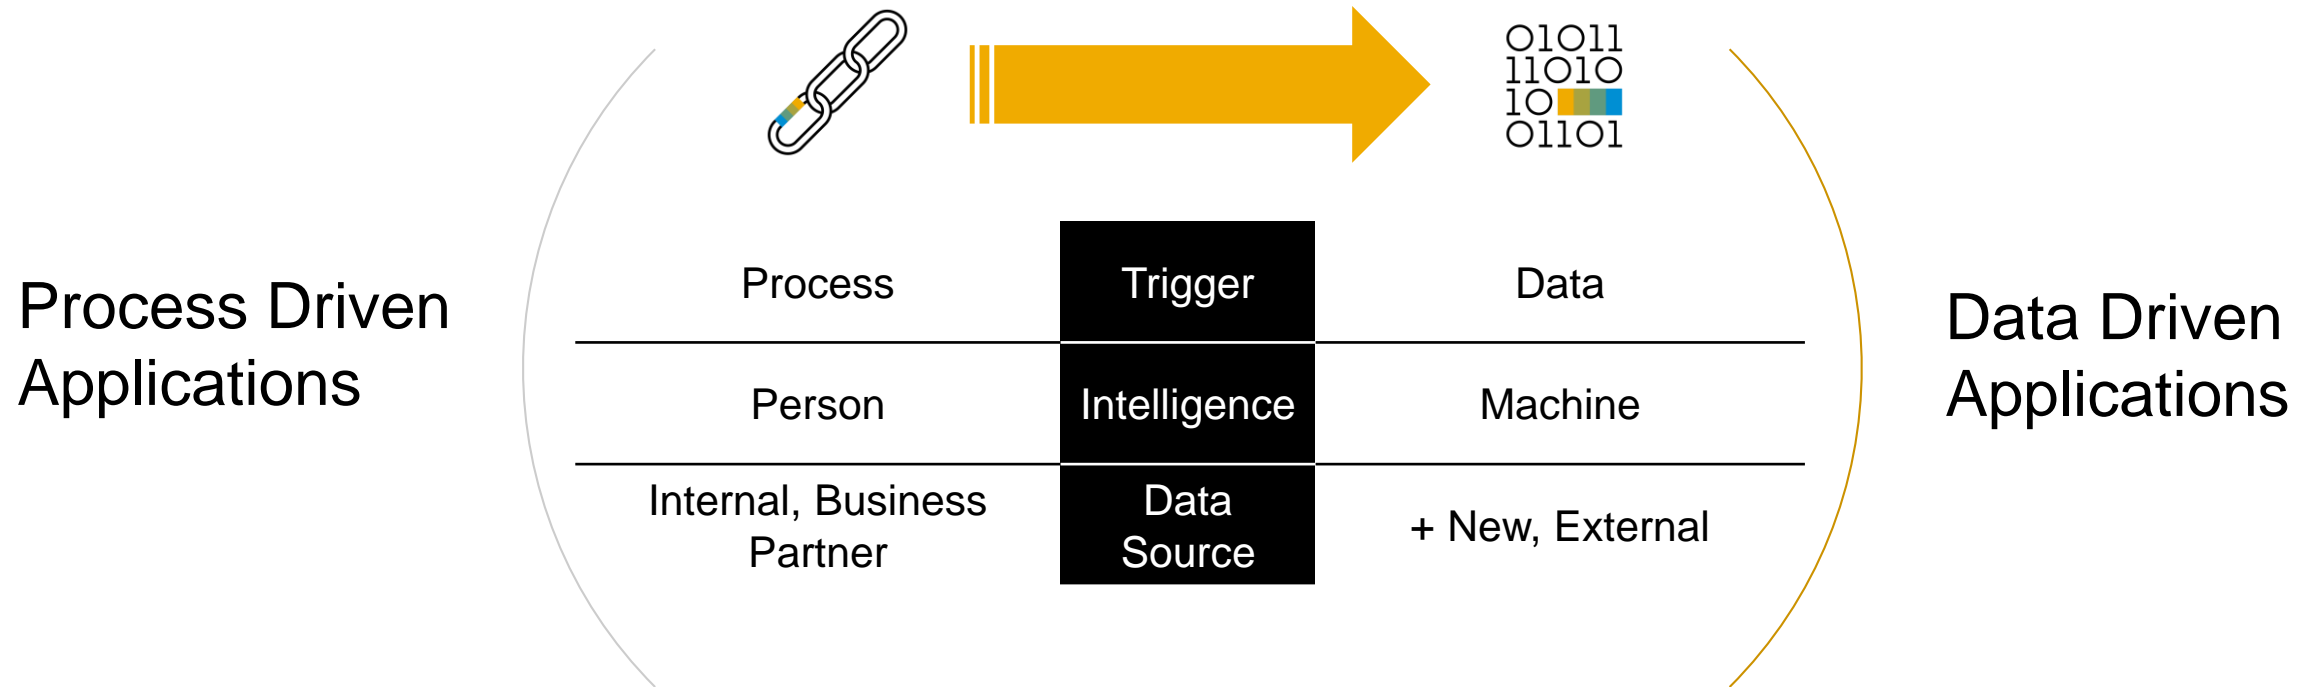
From Connected Data to Insight to Action

Process Driven
Applications

Data Driven
Applications



From Data to Data Driven Applications



DIGITAL / PHYSICAL

BLUR

UNDER ARMOUR CONNECTED FITNESS™

“When we started it was very much about the product. It was about the shirts and shoes. Now it’s much more about the digital product. The reality is that the physical product and the digital product are not distinguishable.”


Kurt Kendall, VP Analytics and Data Solutions at Under Armour




New Challenges New Technologies New Challenges

New challenges require new technologies


Distributed systems in a distributed landscape




Existing Systems




Hadoop/Spark/
NoSQL



Cloud Storage
(i.e. AWS S3)



Machine
Learning
(Python, Spark,
Tensorflow)




Containers
(Kubernetes, Docker)

Data volume is NOT the problem anymore

Data

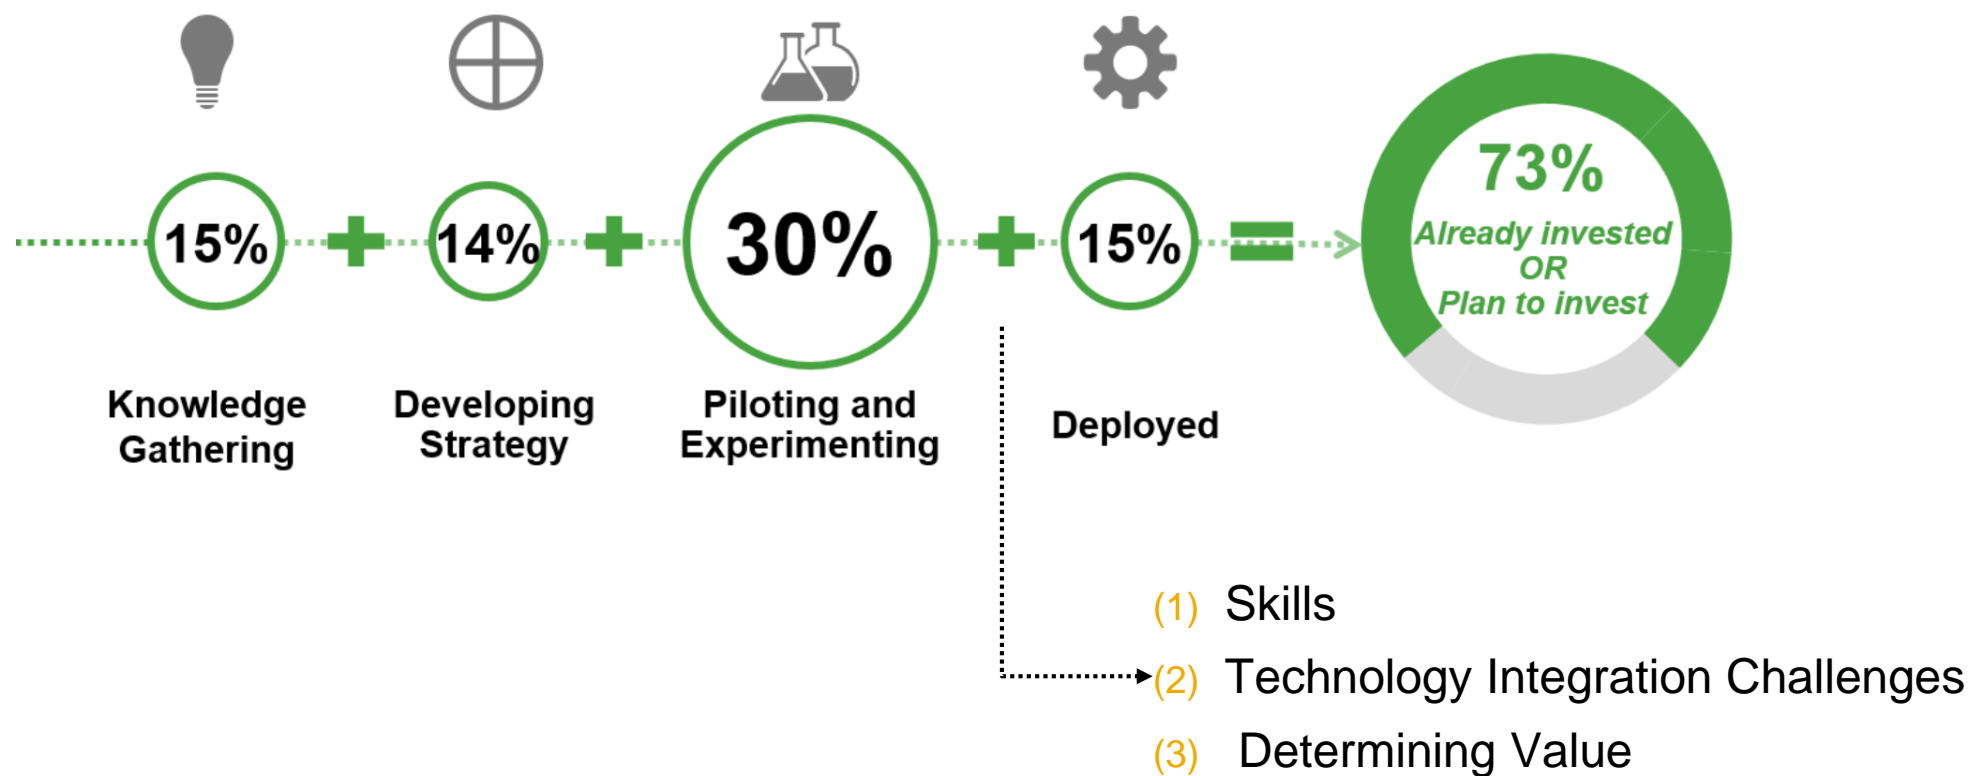
Budget



“**Delivering connected data across on-premises and cloud sources is not trivial**, especially when it involves large data volumes, complex data models, and high speed of ingestion.”

Noel Yuhanna,
Forrester

Gartner: „Still challenging to get Hadoop projects into production stage“



Supporting Data Innovation Projects...

From Initial Pilot to Large Scale Success

PIONEER



Show possibilities

Inspire business

Get buy-in

AUTOMATE



Automate & orchestrate

Schedule & monitor

Document

SCALE



Cloud deployment

Governance

Security

USER ACCEPTANCE



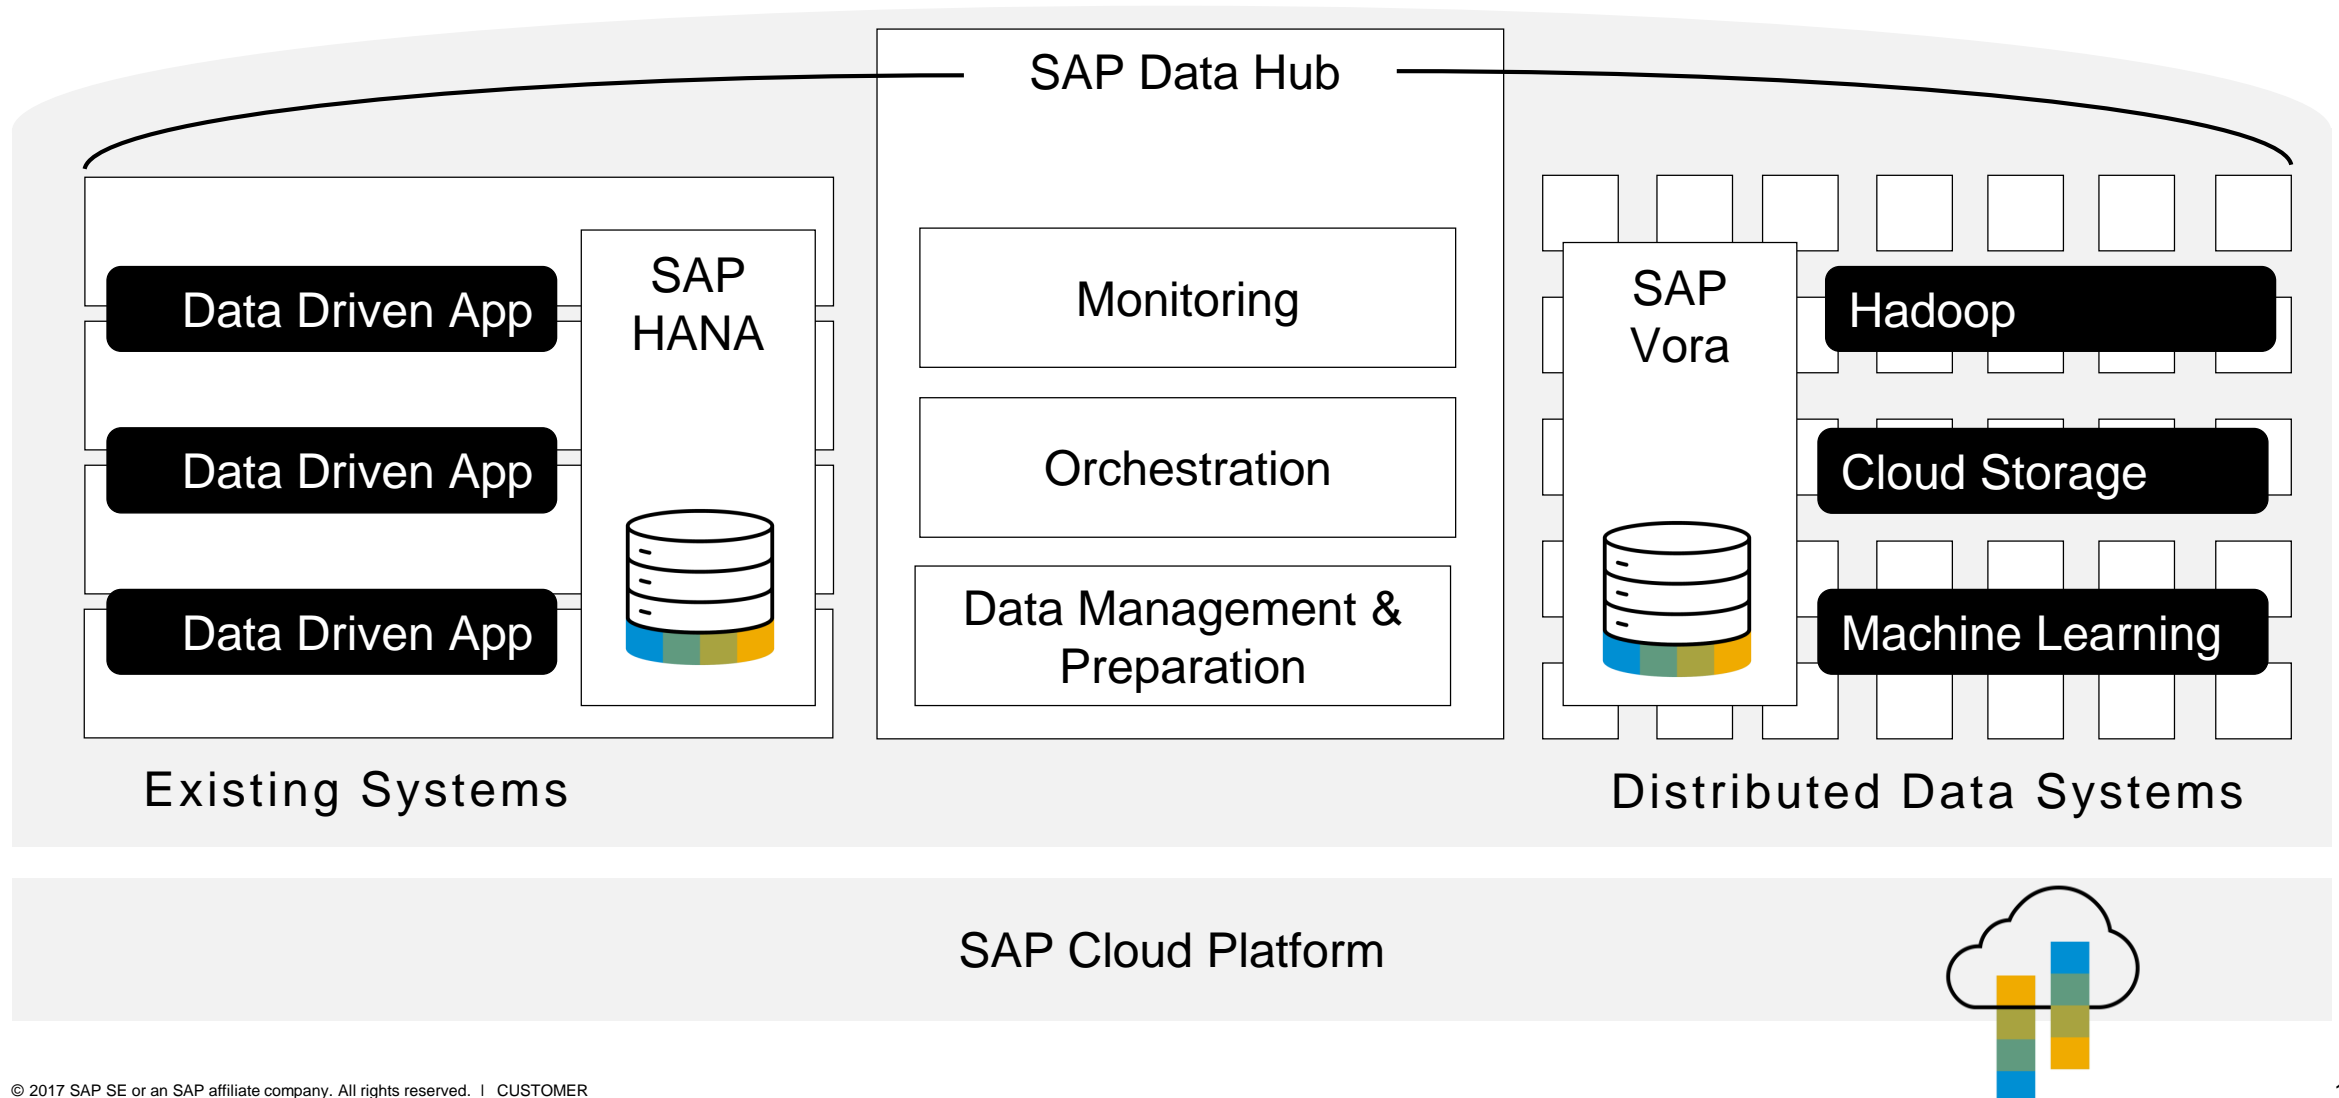
Self-Service Data
Preparation

Data driven apps

Use your own tools &
languages


Platform for Data Driven Applications

Simplify: Existing Skills, Technology Integration, Stable Operations




From Data to Data Driven Applications


LAND



REFINE



CONSUME

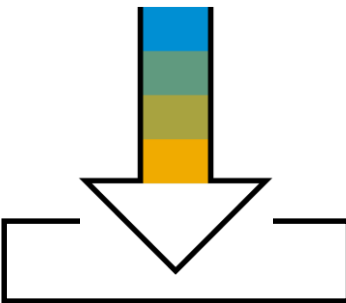


ACT



From Data to Data Driven Applications

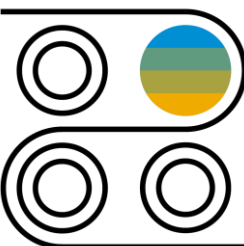
LAND



#123
\$
10101
%&?§

SAP DS | HDFS | S3

REFINE

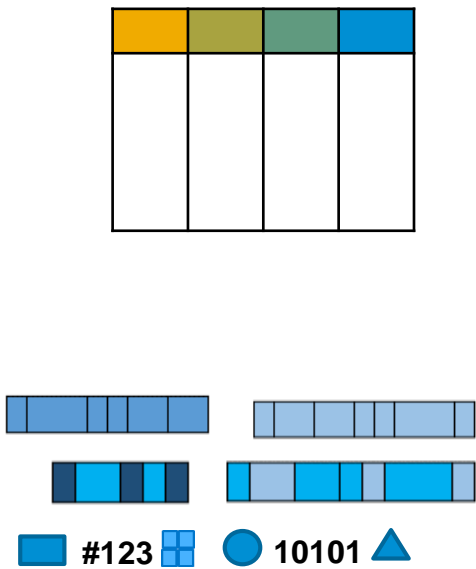


Join Cleanse Script Anonymize
Filter Look-up Mask Parse

Spark Job Python Script

Data Hub

CONSUME



#123 10101

Vora

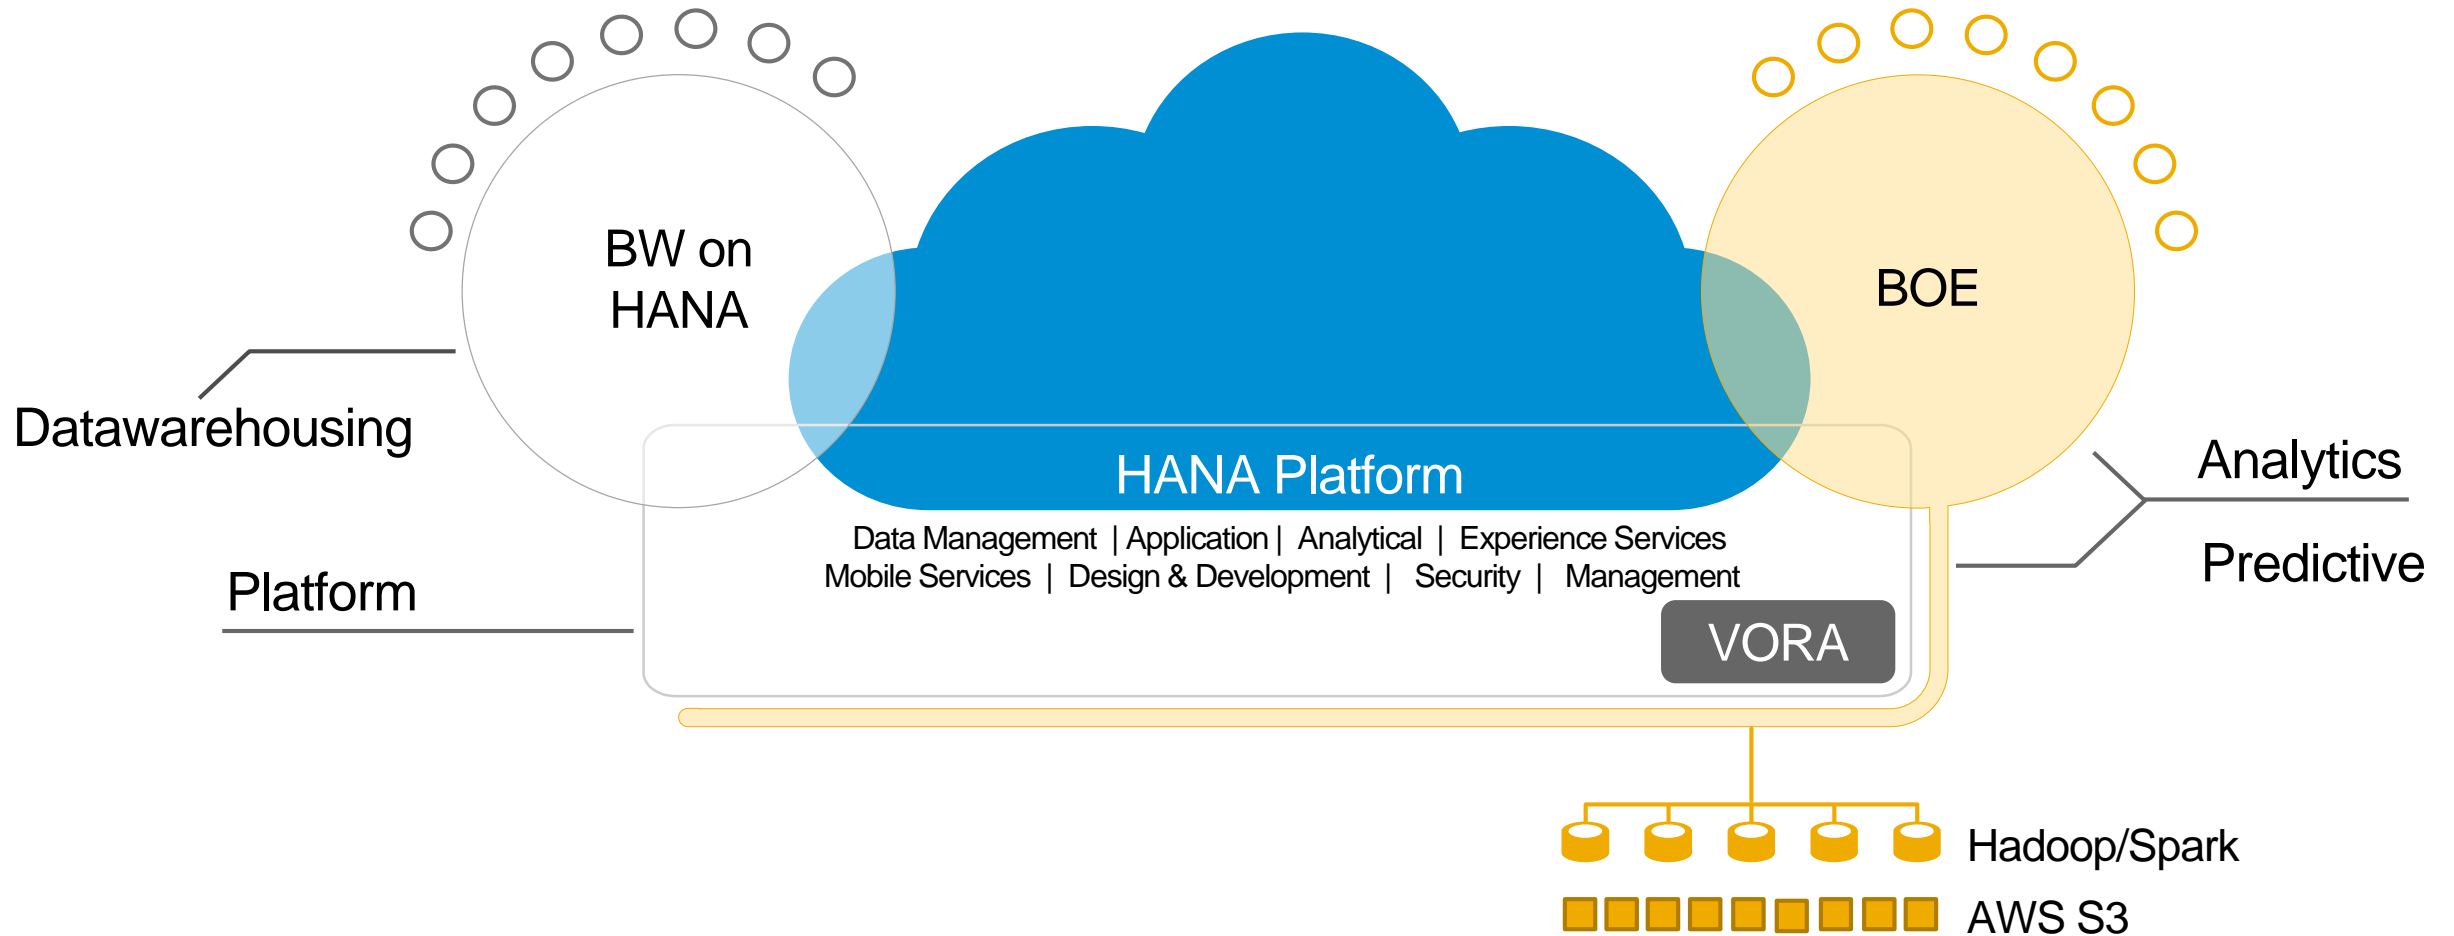
ACT



HANA | BW | Apps



Kundenbeispiel

High-Level Architecture




From Data to Data Driven Applications

LAND




SAP DS | Kafka | HDFS | S3

REFINE



Join Cleanse Script Anonymize
Filter Look-up Mask Parse





Spark Job

Python Script


Data Hub

CONSUME



Vora

ACT



HANA | BW | Apps



CIO/CDO/CTO


Business User



Business Analyst



Data Scientist



Data Engineer

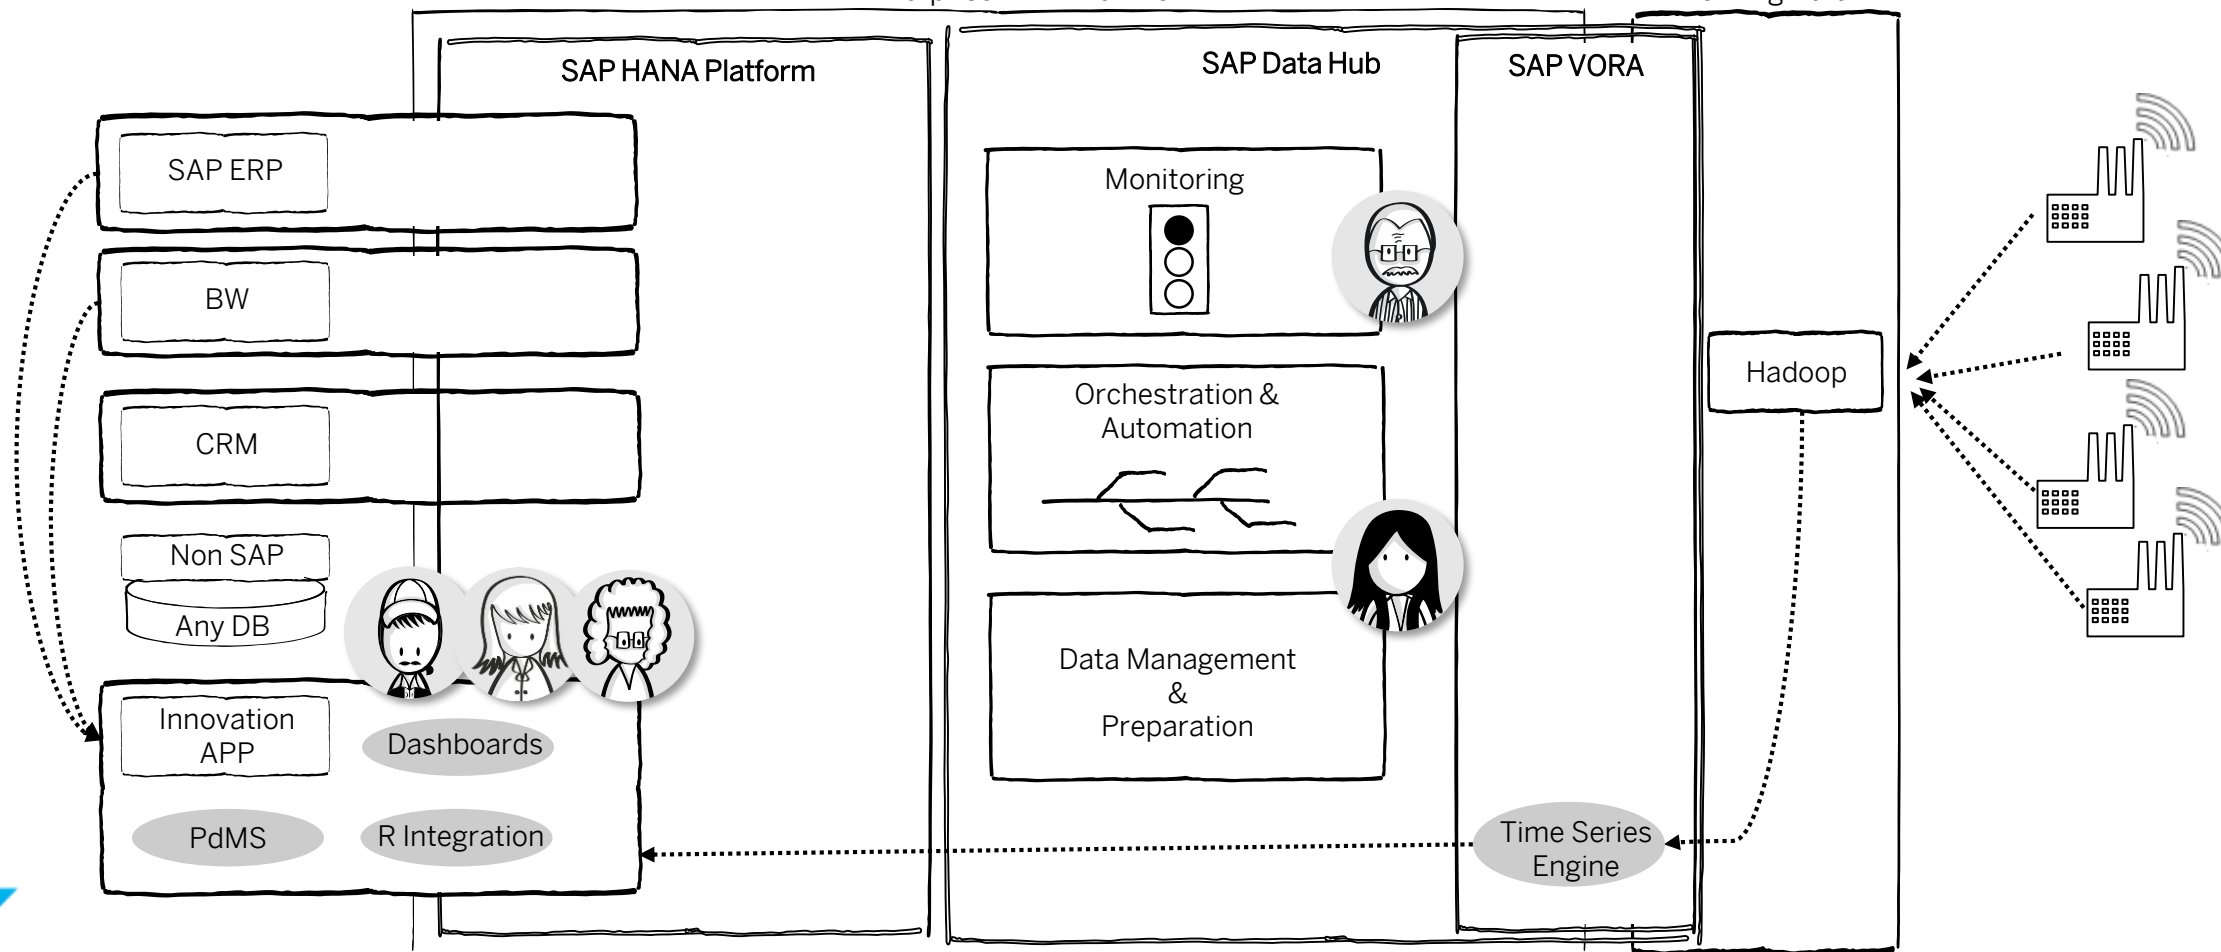


IT Operations

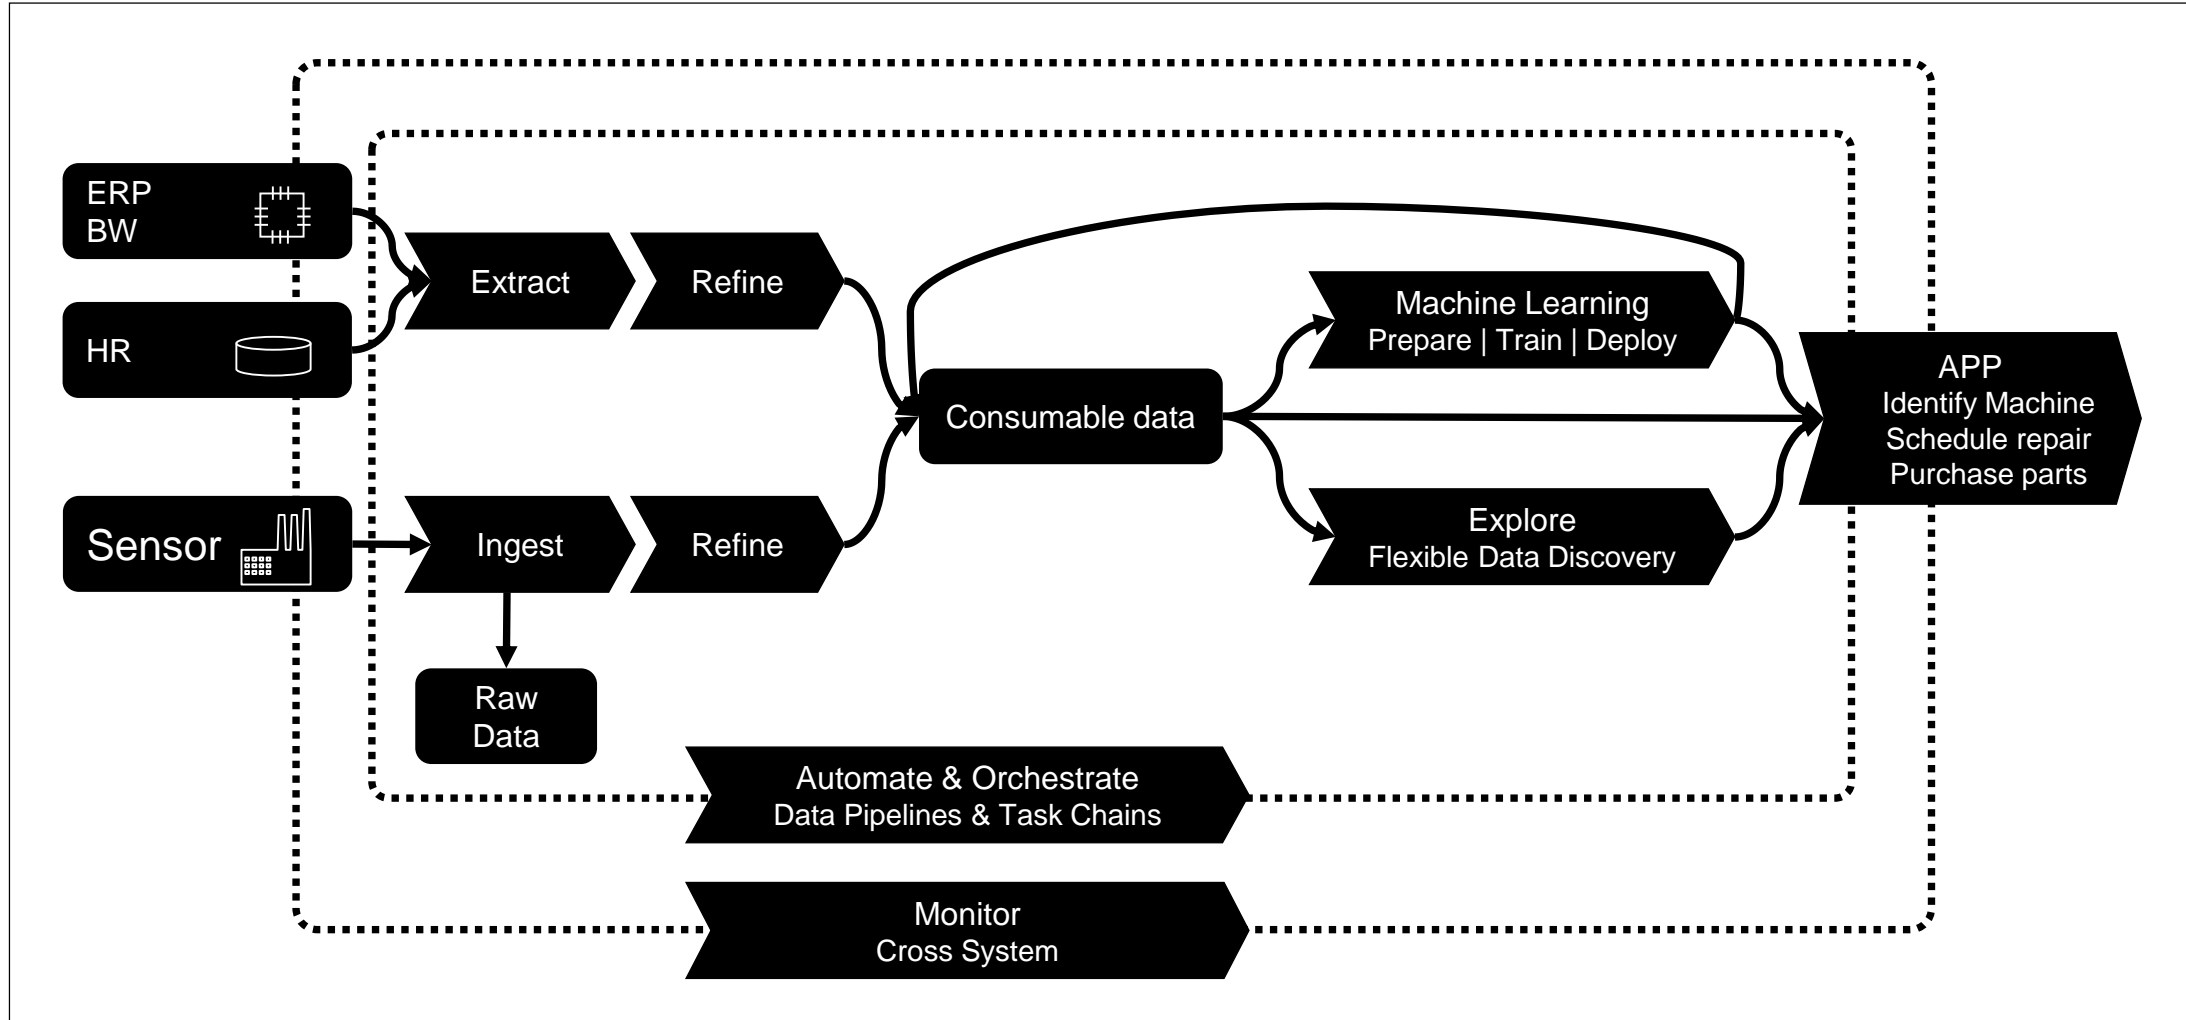


Enterprise IT Environment


New Big Data



Task Chain for IoT Scenario



SAP Data Hub



SAP Data Hub – Scope

SAP DATA HUB - Distributed Data Hub & Refinery

Refines enterprise big data across a heterogeneous data landscape to drive digital business decisions

Data Cockpit

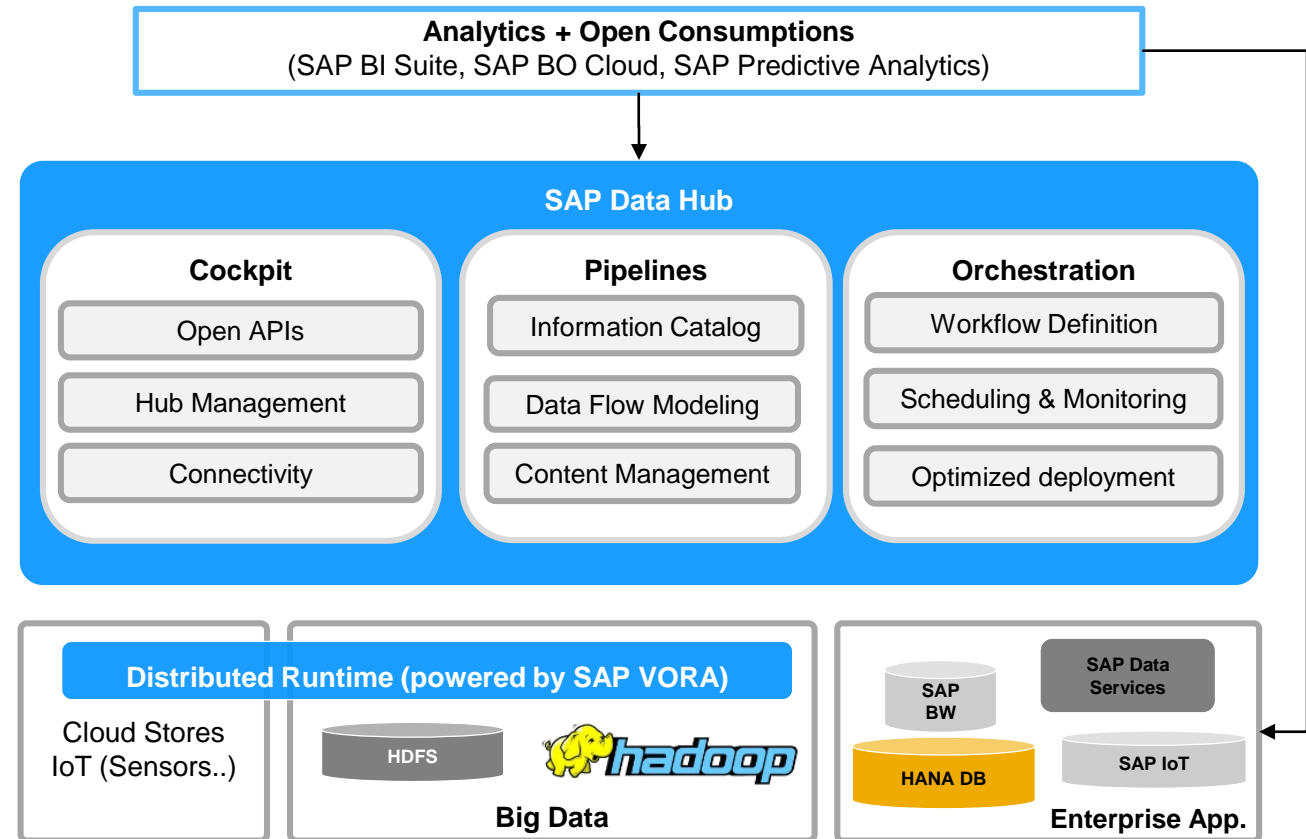
- Role-based Hub Experience with a modern UX
- Creation of logical hub structure e.g. zone management
- Connectivity to hybrid system landscapes
- Big Data Discovery & Security

Data Pipelines

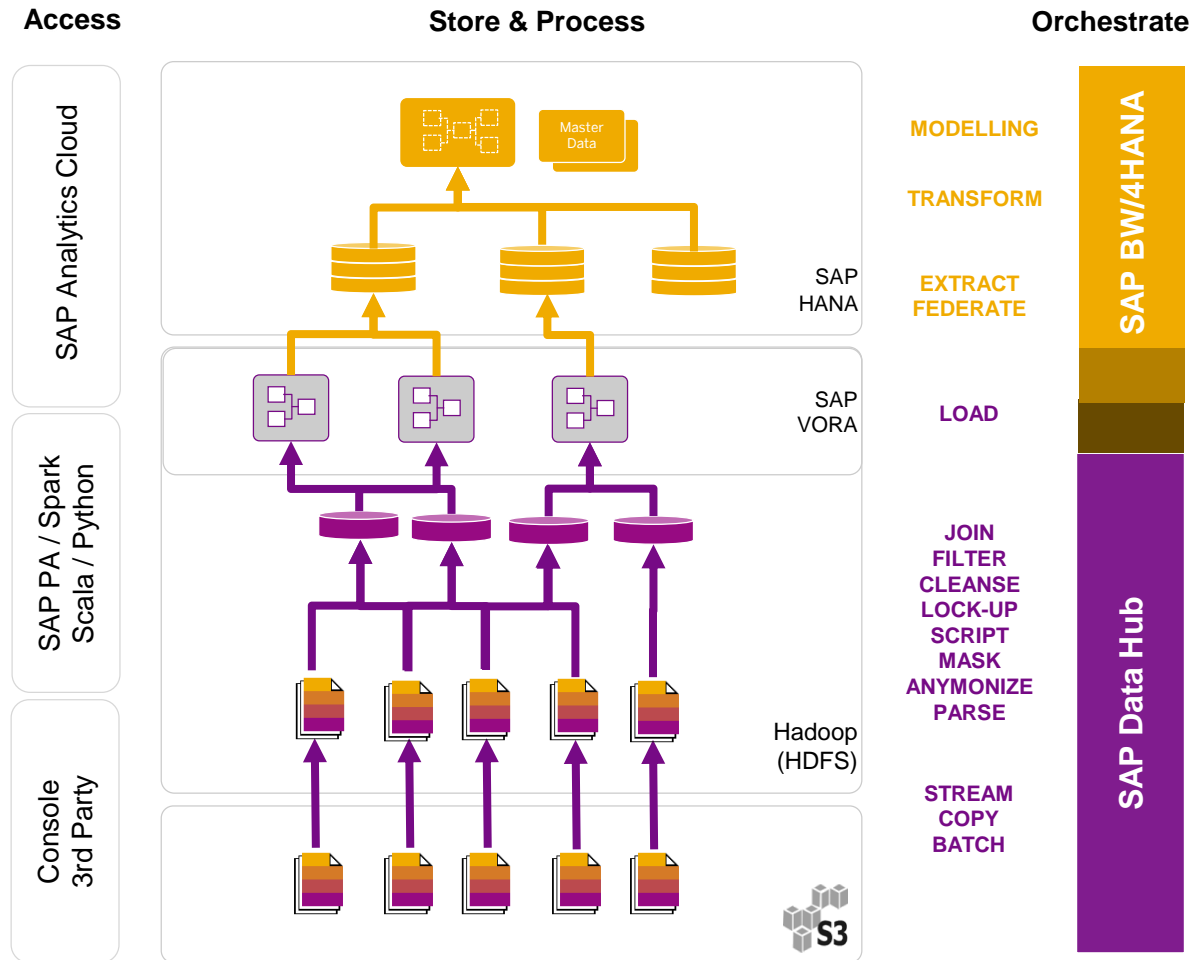
- Data Ingestion (Connectors) and Integration
- Data Pipelines suited for Server less Computing
- Data Quality Flow with predefined operations
- Meta Data Services across stores with Lineage

Orchestration

- Workflow scheduling
- Monitoring and alerting
- Pipeline deployment to SAP HANA and Vora



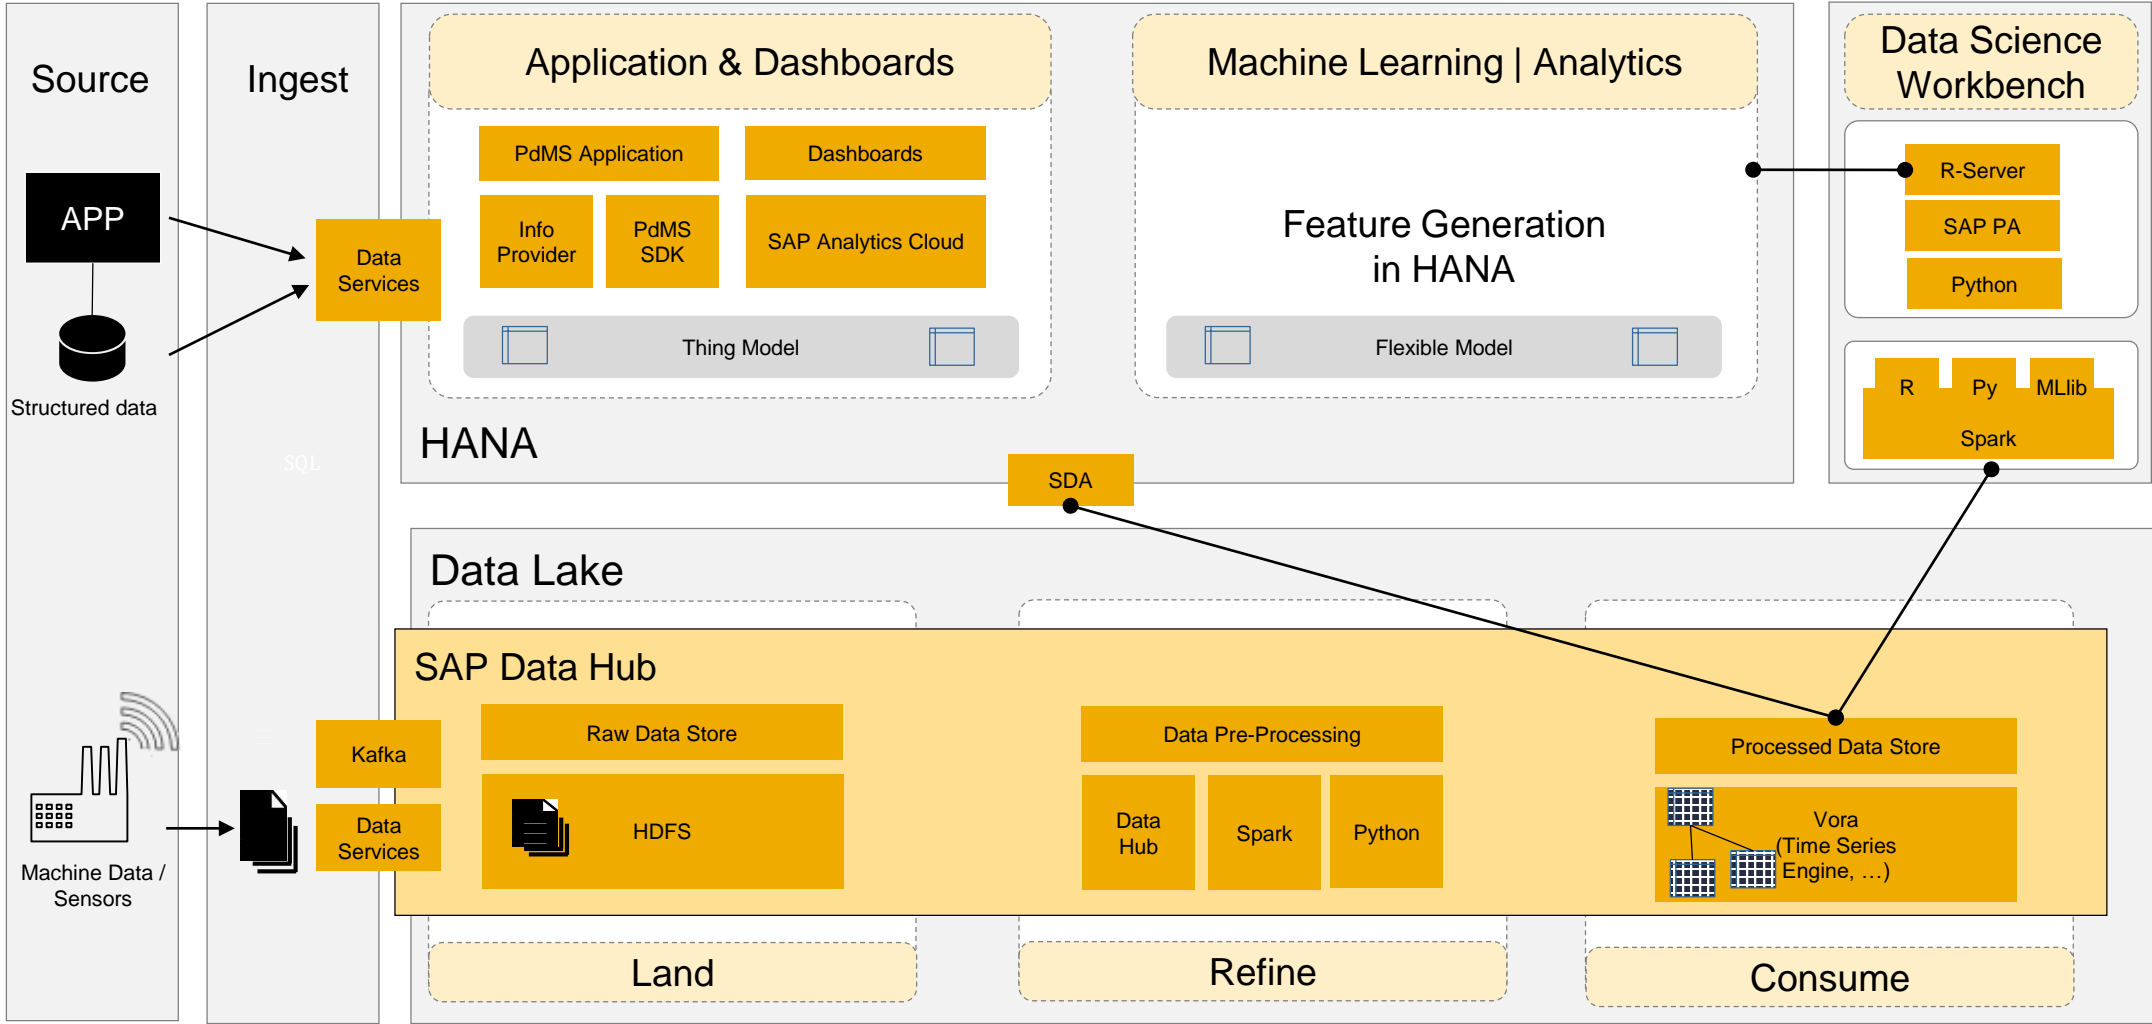
SAP Data Hub – Analyzing Web Log Data with Data Warehouse



Example Scenario

- Combine refined big data with enterprise data and corporate master data
- Extract or federate data into SAP BW/4HANA
- Ingest Data into S3 as Landing Zone for data
- Orchestrate and schedule all related processes
- Implement transformations and data pipelines
- Harmonize data structures and look up of reference data
- Execute operations on large data volumes

IoT Architecture for Predictive Maintenance



USER EXPERIENCE

Same thing – different perspective



A logistic regression takes me only a few lines of code.



Data Scientist

```
import org.apache.spark.mllib.classification.{SVMModel, SVMWithSGD}
import org.apache.spark.mllib.evaluation.BinaryClassificationMetrics
import org.apache.spark.mllib.util.MLUtils

// Load training data in LIBSVM format.
val data = MLUtils.loadLibSVMFile(sc, "data/mllib/sample_libsvm_data.txt")

// Split data into training (60%) and test (40%).
val splits = data.randomSplit(Array(0.6, 0.4), seed = 11L)
val training = splits(0).cache()
val test = splits(1)

// Run training algorithm to build the model
val numIterations = 100
val model = SVMWithSGD.train(training, numIterations)

// Clear the default threshold.
model.clearThreshold()

// Compute raw scores on the test set.
val scoreAndLabels = test.map { point =>
  val score = model.predict(point.features)
  (score, point.label)
}

// Get evaluation metrics.
val metrics = new BinaryClassificationMetrics(scoreAndLabels)
val auROC = metrics.areaUnderROC()

println("Area under ROC = " + auROC)

// Save and load model
model.save(sc, "target/tmp/scalaSVMWithSGDModel")
val sameModel = SVMModel.load(sc, "target/tmp/scalaSVMWithSGDModel")
```

- } Split data
- } Train one model
- } Apply the model on new data
- } Evaluate model quality

Source: <http://spark.apache.org/docs/latest/mllib-linear-methods.html>

Looks kinda complicated :-/



Business Analyst

```
import org.apache.spark.mllib.classification.{SVMModel, SVMWithSGD}
import org.apache.spark.mllib.evaluation.BinaryClassificationMetrics
import org.apache.spark.mllib.util.MLUtils

// Load training data in LIBSVM format.
val data = MLUtils.loadLibSVMFile(sc, "data/mllib/sample_libsvm_data.txt")

// Split data into training (60%) and test (40%).
val splits = data.randomSplit(Array(0.6, 0.4), seed = 11L)
val training = splits(0).cache()
val test = splits(1)

// Run training algorithm to build the model
val numIterations = 100
val model = SVMWithSGD.train(training, numIterations)

// Clear the default threshold.
model.clearThreshold()

// Compute raw scores on the test set.
val scoreAndLabels = test.map { point =>
  val score = model.predict(point.features)
  (score, point.label)
}

// Get evaluation metrics.
val metrics = new BinaryClassificationMetrics(scoreAndLabels)
val auROC = metrics.areaUnderROC()

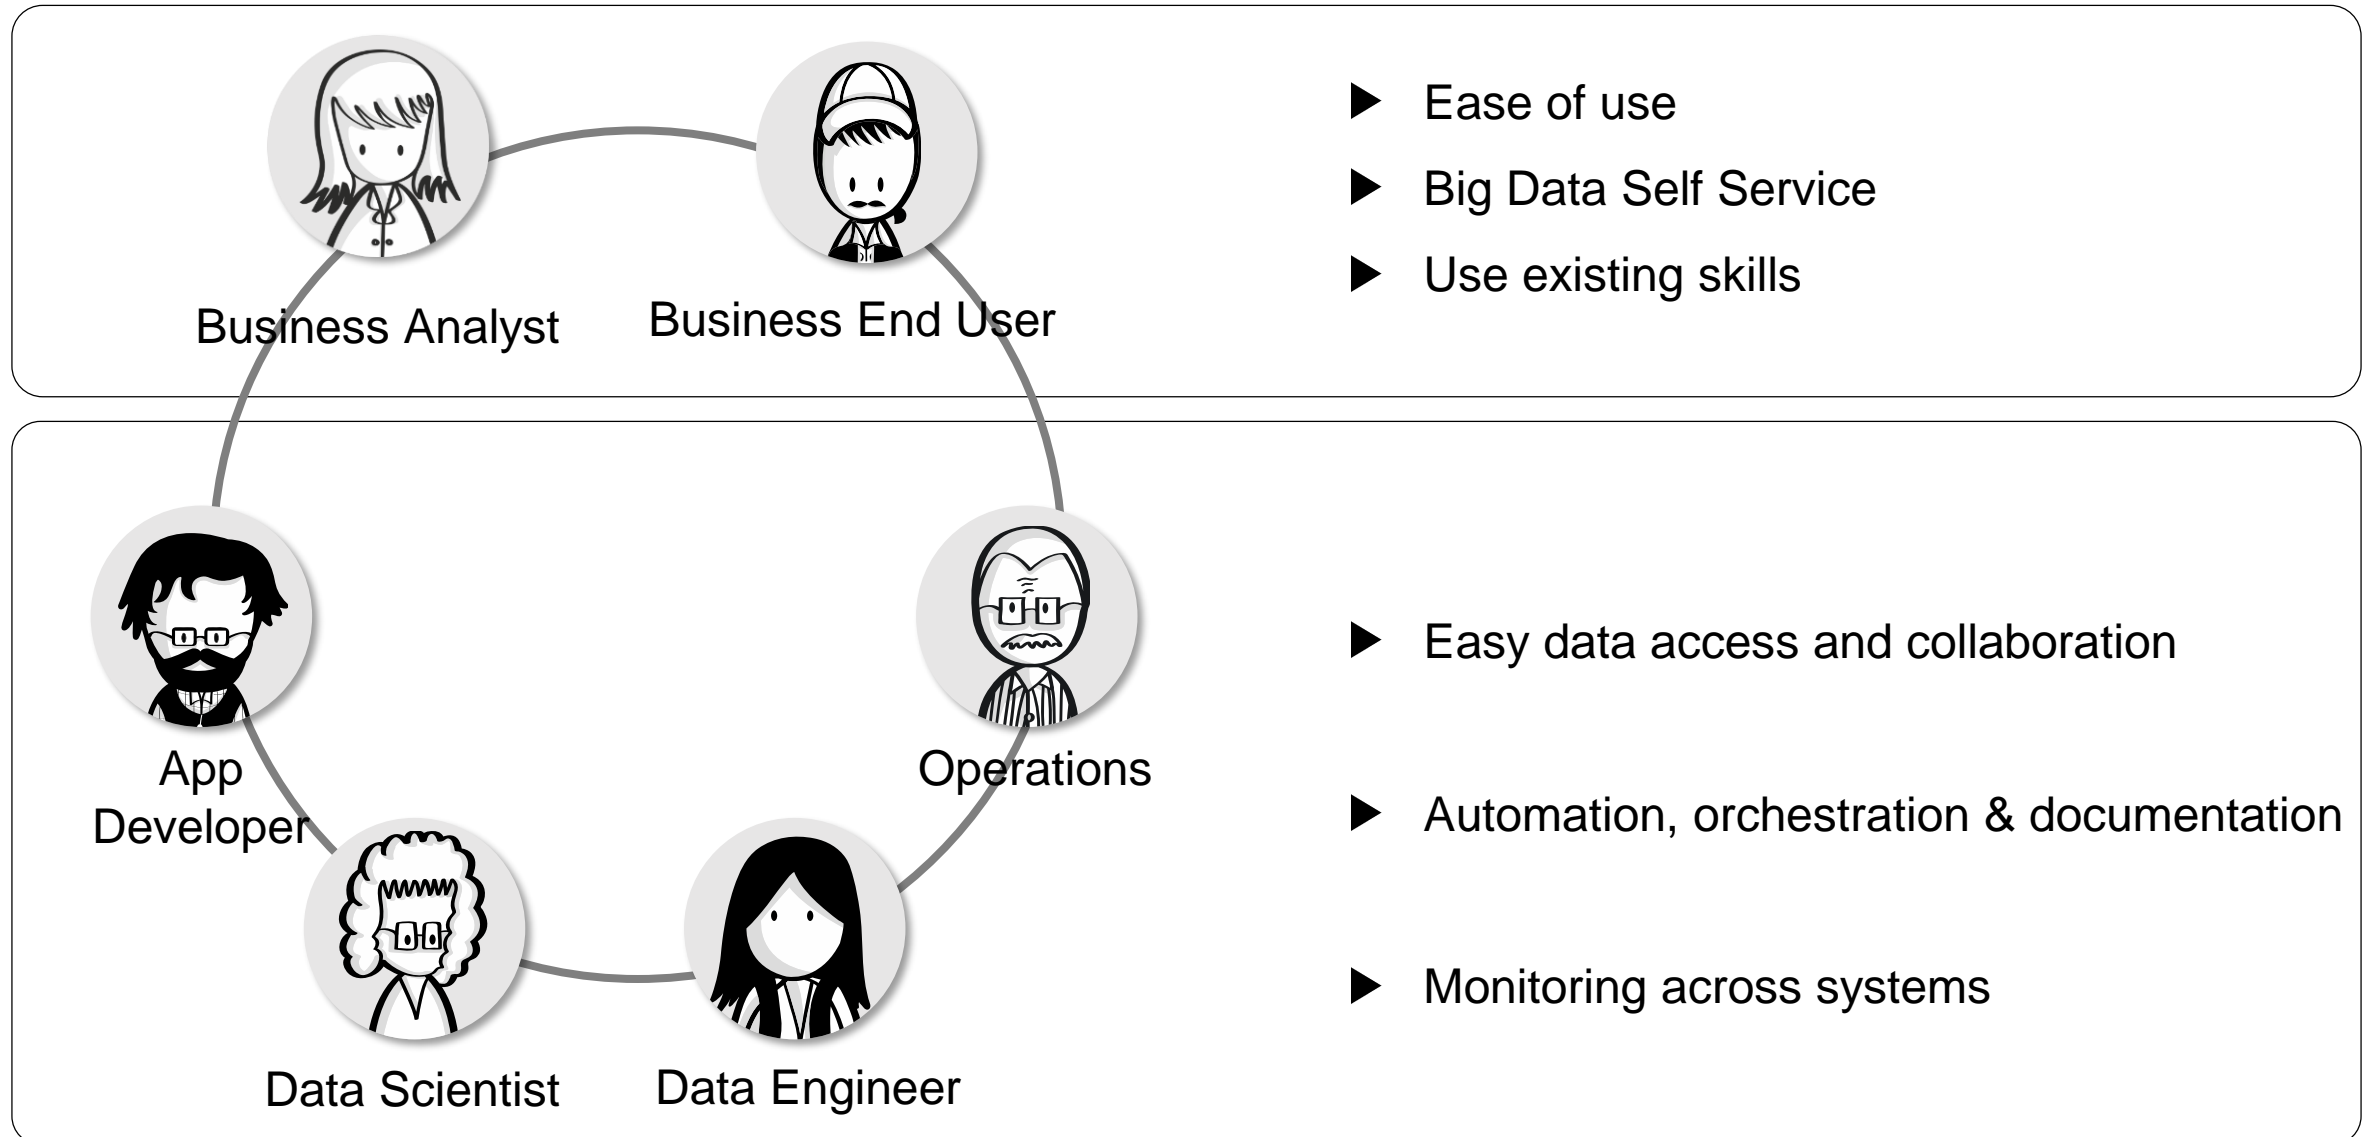
println("Area under ROC = " + auROC)

// Save and load model
model.save(sc, "target/tmp/scalaSVMWithSGDModel")
val sameModel = SVMModel.load(sc, "target/tmp/scalaSVMWithSGDModel")
```

- } Split data
- } Train one model
- } Apply the model on new data
- } Evaluate model quality

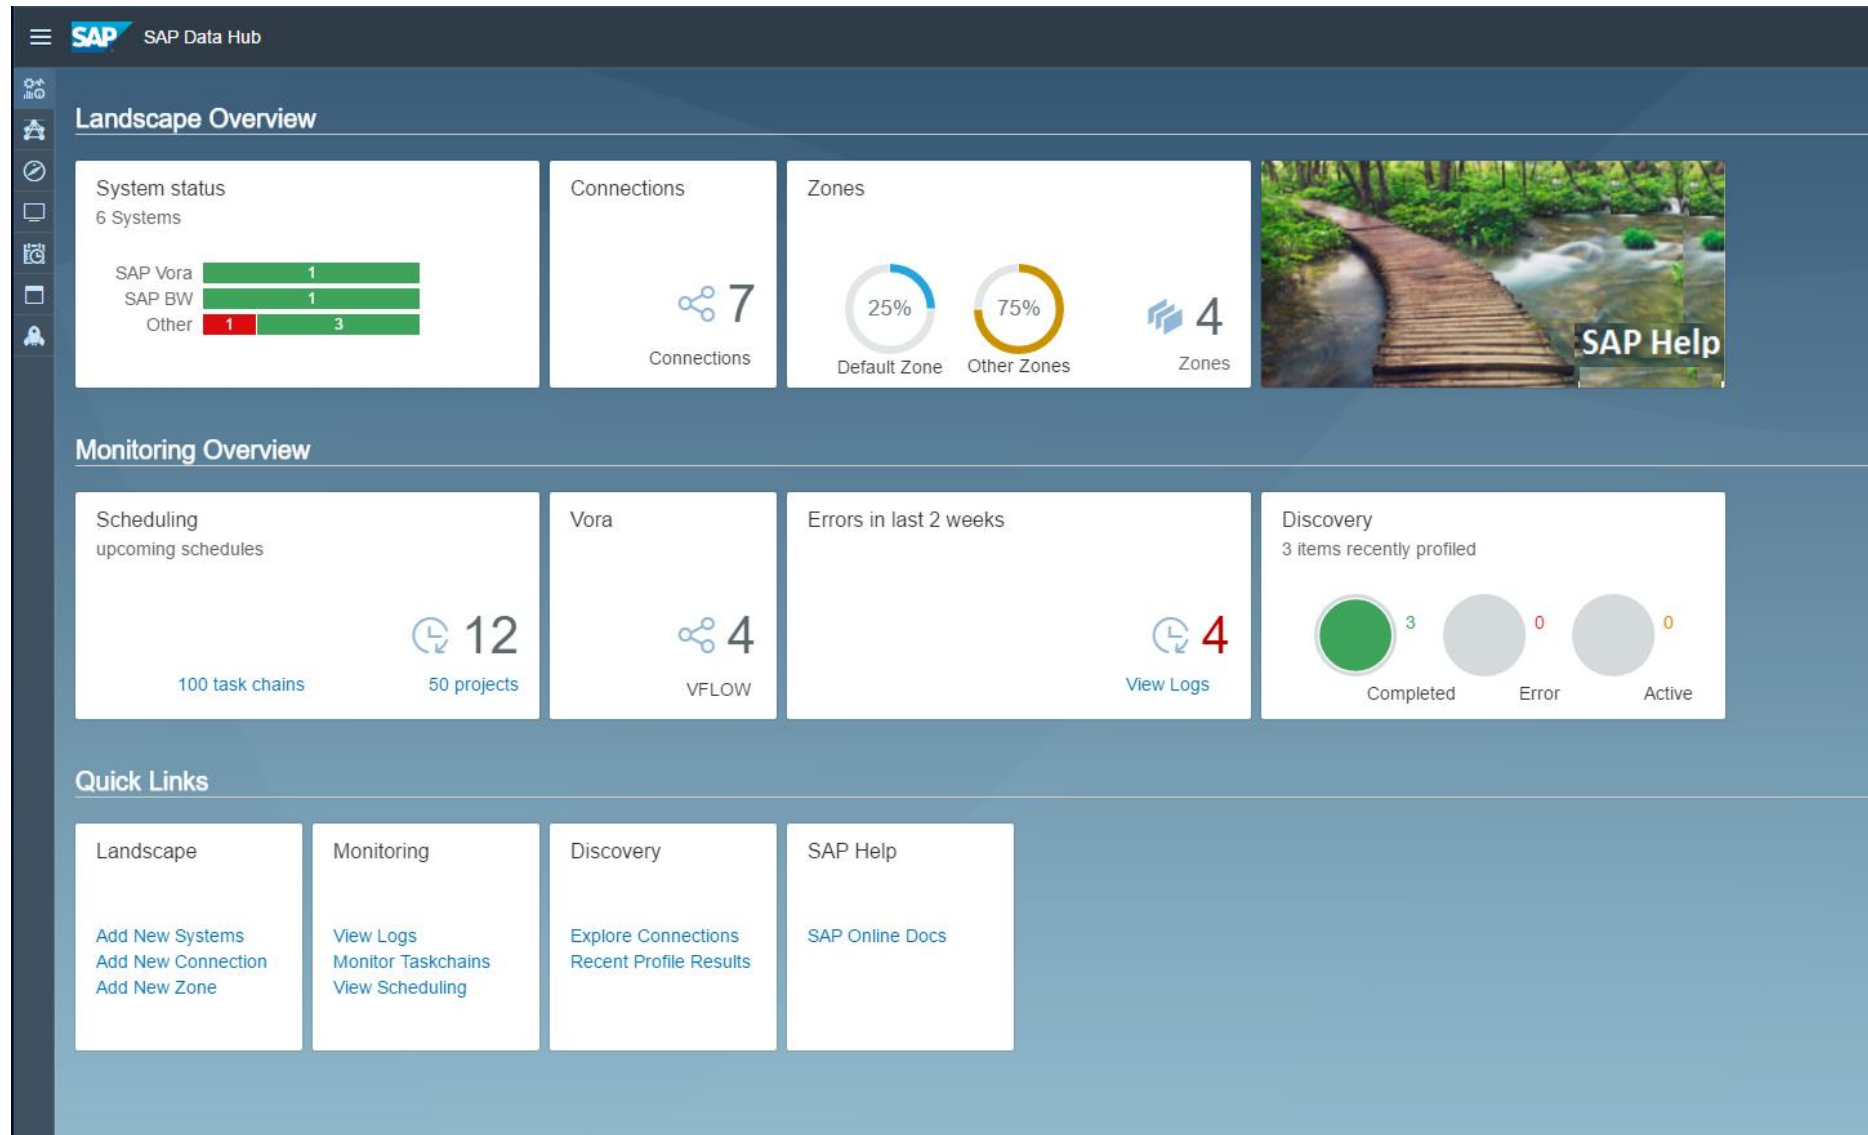
Source: <http://spark.apache.org/docs/latest/mllib-linear-methods.html>

Data Innovation is a Team Effort



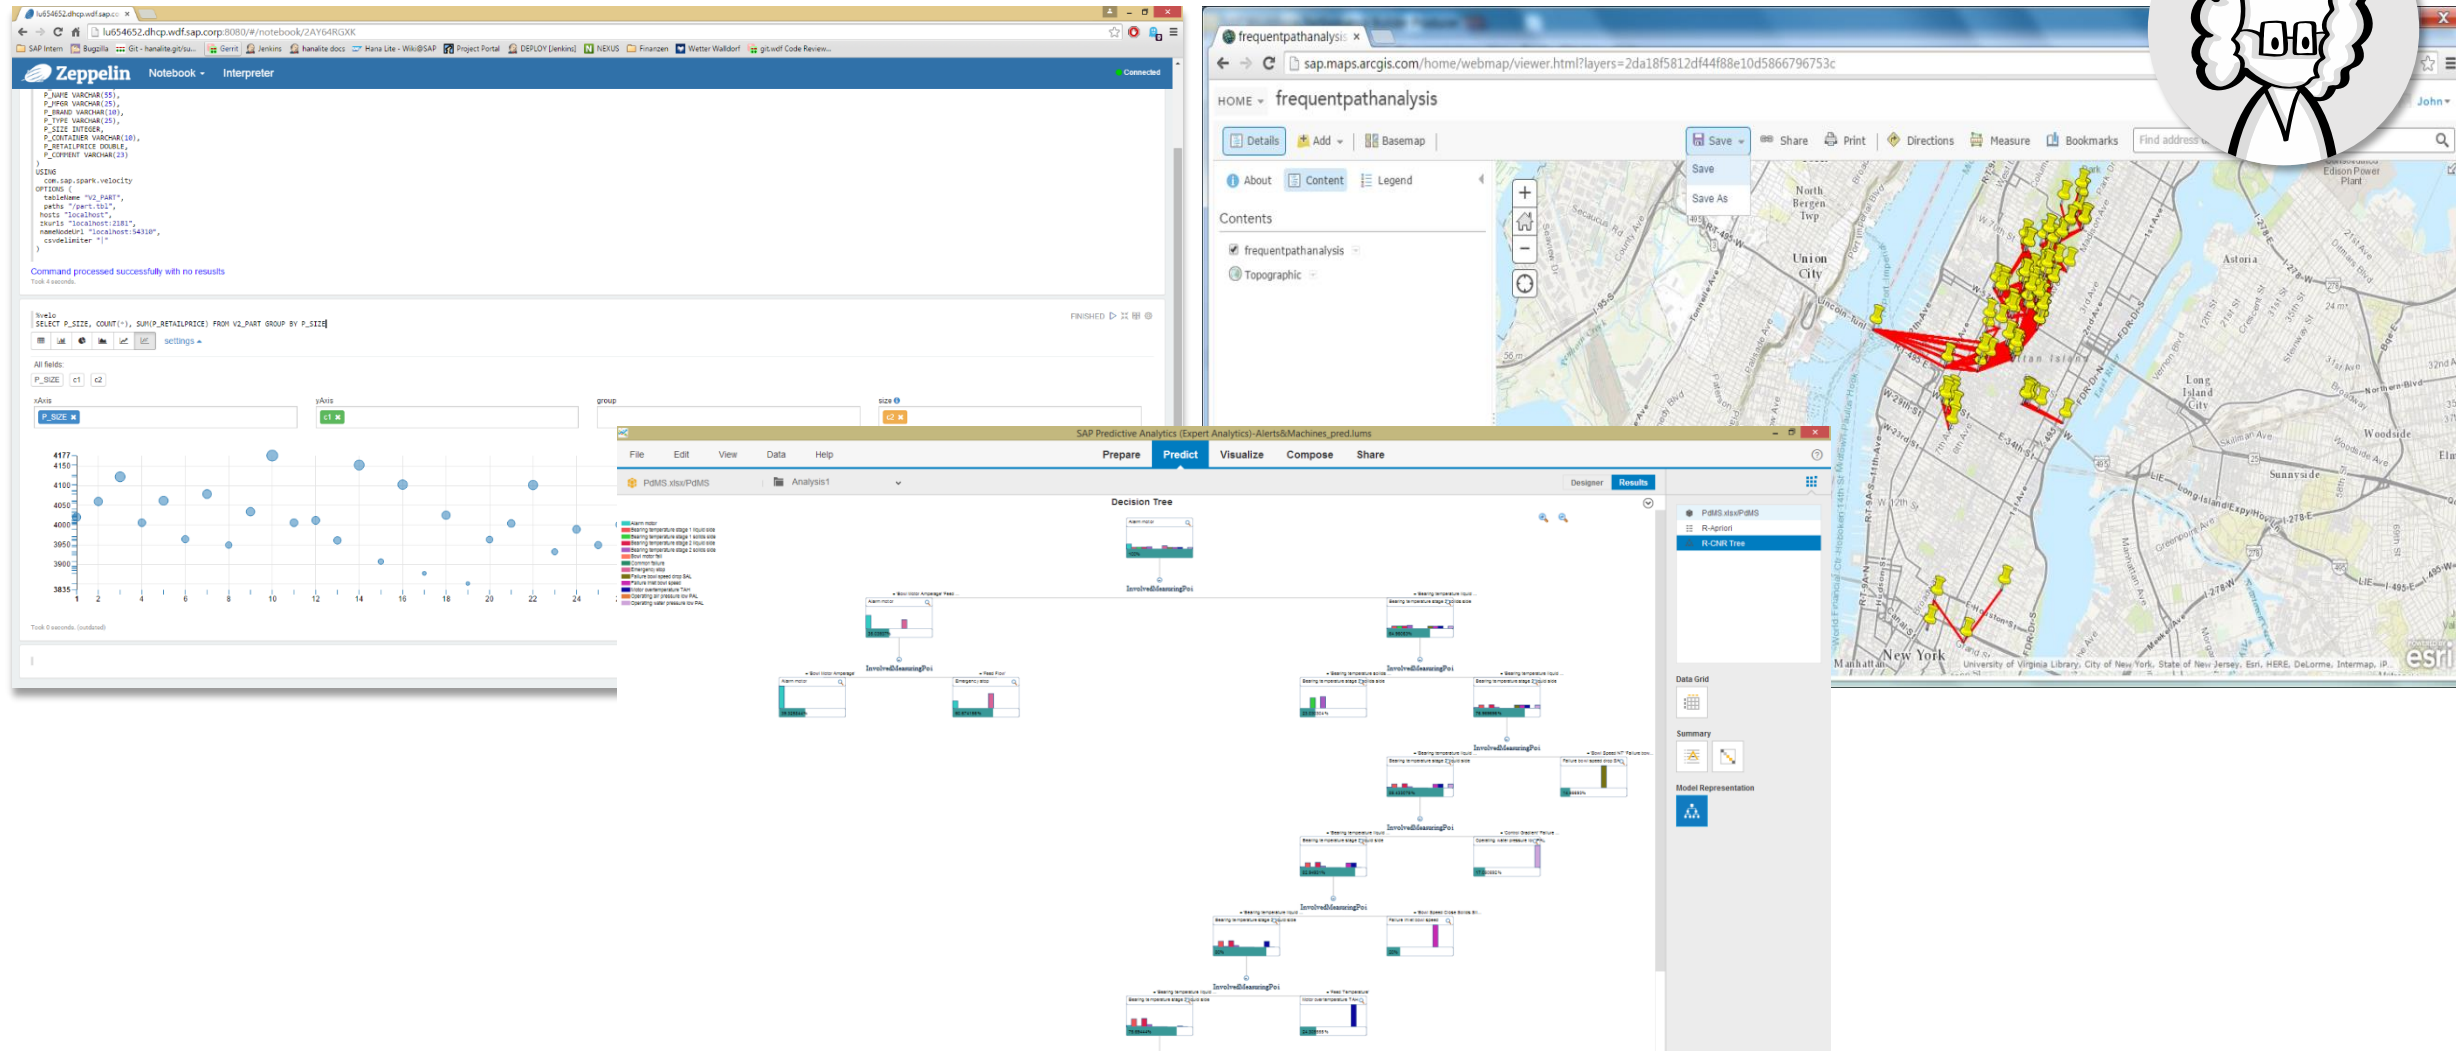
Cross System Monitoring & Management

IT Operations



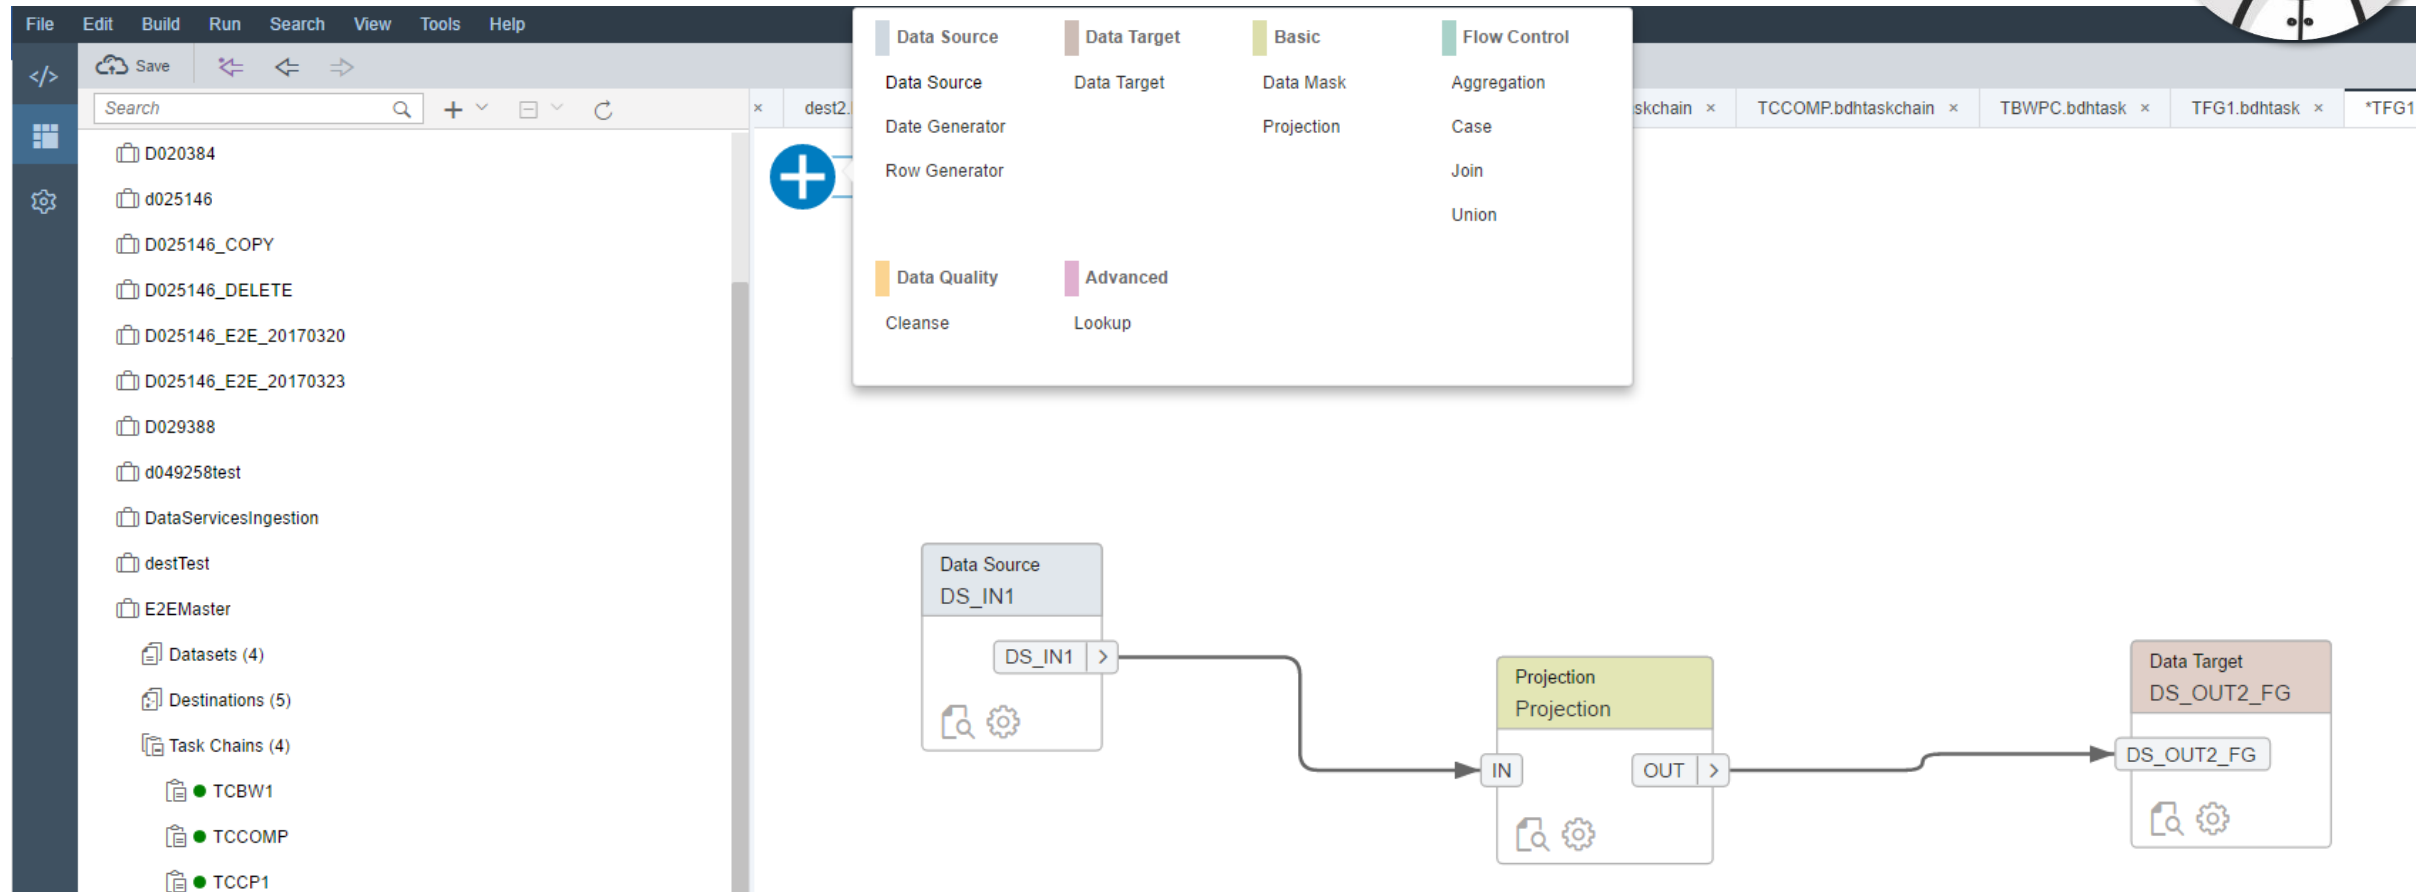
Data Scientist – Apache Zeppelin & SAP Predictive Analytics

Data Scientist



Big Data Self-Service Data Preparation

Business Analyst



Big Data Self-Service Data Modelling

Business Analyst



SAP HANA Vora

Filter

PART

PARTSUPP

PMDEMUNION

REGION

SUPPLIER

TEST1

TEST2

TEST3

V

V1

V2

V3

VV2

VVV2

V_STEST

XMYVIEW

XXVVV2

V_STEST

100%

CUSTOMER_1 (...)

- ☒ C_CUSTKEY INTEGER
- ☒ C_NAME STRING
- ☐ C_ADDRESS STRING
- ☒ C_NATIONKEY INTEGER
- ☐ C_PHONE STRING
- ☐ C_ACCTBAL DOUBLE
- ☐ C_MKTSEGMENT STRING
- ☐ C_COMMENT STRING

CUSTOMER

- ☒ C_CUSTKEY INTEGER
- ☒ C_NAME STRING
- ☐ C_ADDRESS STRING
- ☒ C_NATIONKEY INTEGER
- ☐ C_PHONE STRING
- ☐ C_ACCTBAL DOUBLE
- ☐ C_MKTSEGMENT STRING
- ☐ C_COMMENT STRING

ResultSet

- ☐ C_CUSTKEY INTEGER
- ☐ C_NAME STRING
- ☐ C_NATIONKEY INTEGER

Default ResultSet

- ☒ C_CUSTKEY INTEGER
- ☒ C_NAME STRING
- ☒ C_NATIONKEY INTEGER

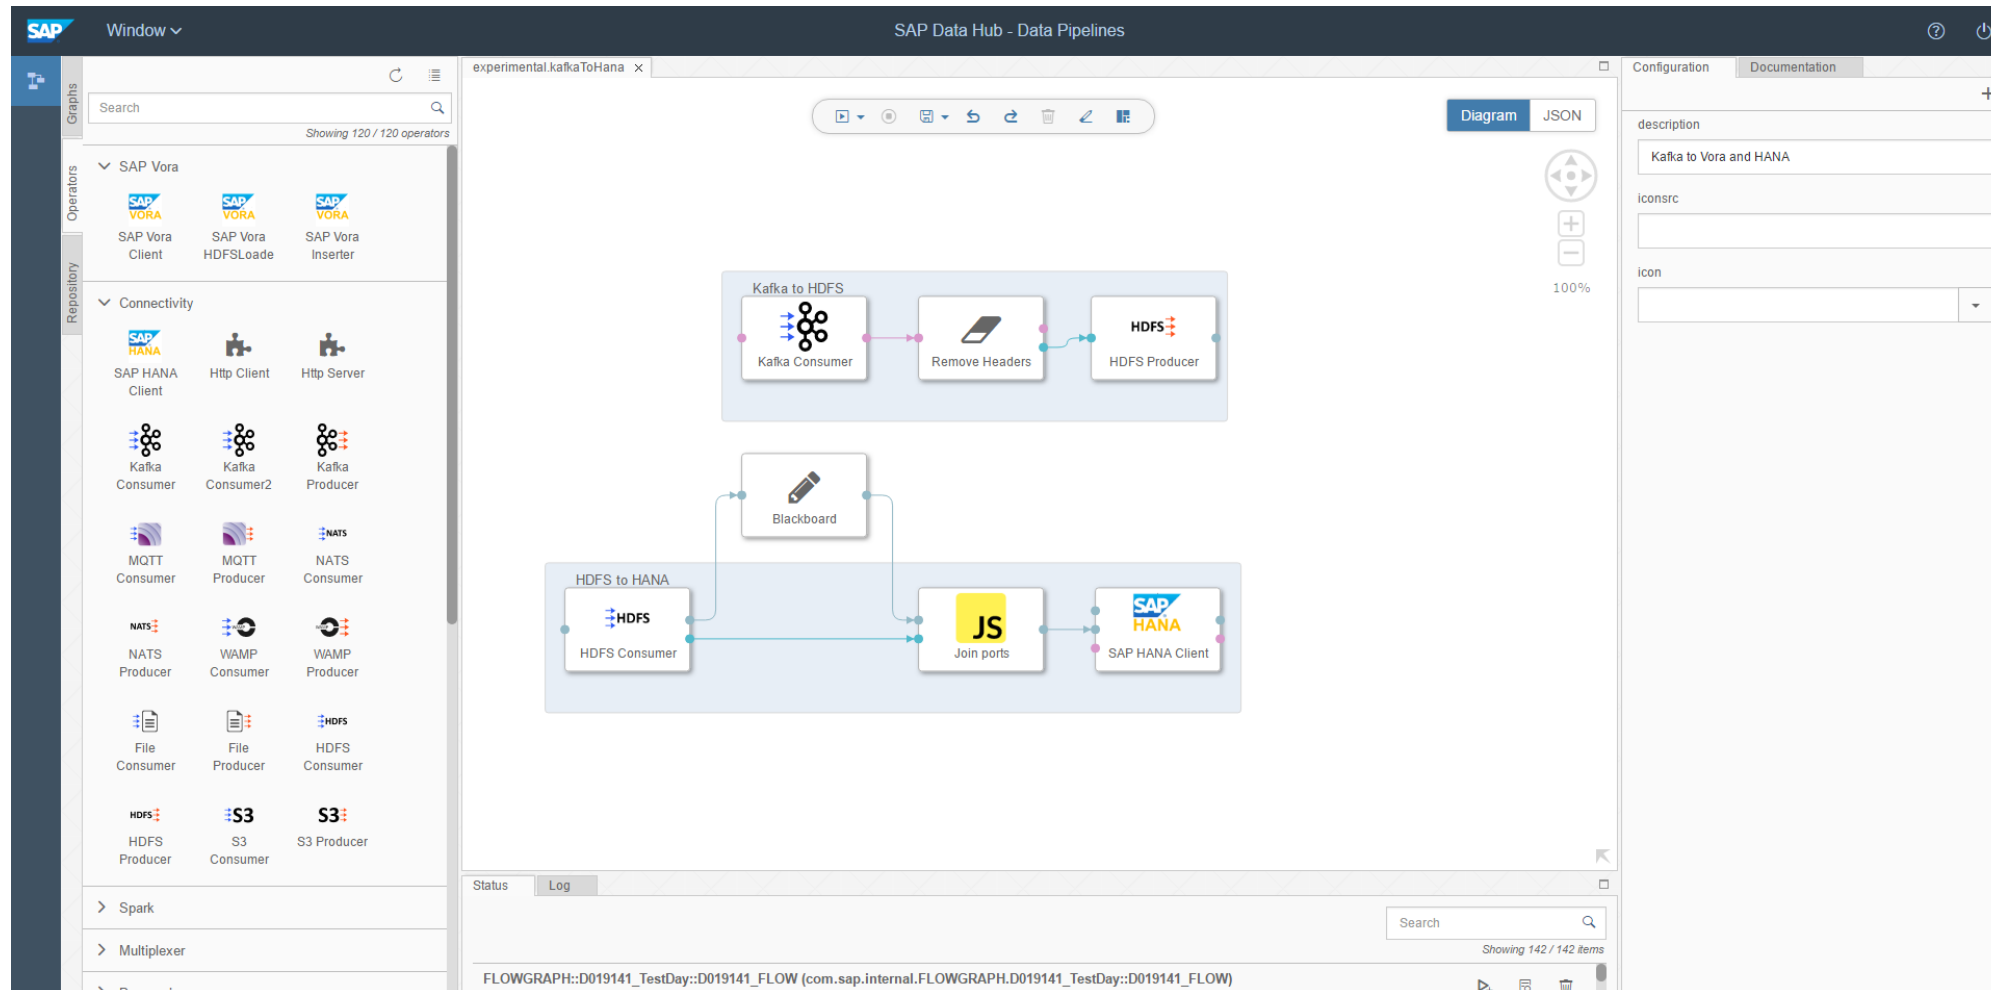
V_STEST

- ^ Data Source +
- ^ Columns +
- ^ Where Clause +
- ^ Group By +
- ^ Order By +
- ^ Having +




Design Cross-System Task Chains

Data Engineer



HANA Developer

Data Engineer



SAP HANA Web-based Development Workbench: Catalog

v 1.130.2 | Help | SYSTEM | VTA | VTA (ip-10-79-22-208 00)

Now editing: SQL CONSOLE: untitled0.sql - Current Schema: SY

1 3

untitled0.sql

```
select top 10 * from "SYSTEM"."V1_OMPLAINTS";
```

2

SAP HANA Web-based Development Workbench: Editor

Content

- 1
- 2
- 3

- New
- Import
- Export Package
- Undo (Command+Z)
- Redo (Command+Y)
- Cut (Command+X)
- Copy (Command+C)
- Paste (Command+V)
- Rename (F2)
- Delivery Unit
- Search Text
- Create Application
- Activate All
- Change Package Attributes
- Refresh
- Synchronize with Github
- Delete
- Force Delete
- Generate JSDoc
- Copy Shortcut

File (Command+Alt+N)

Package (Command+Alt+Shift+N)

Role

- Calculation View
- Analytic Privilege
- Flow Graph
- Replication Task
- HDB Procedure

Result1

ID	DATE_F
1	8/30/13
2	8/30/13
3	8/30/13
4	8/30/13
5	8/30/13
6	8/30/13
7	8/30/13
8	8/30/13
9	8/30/13
10	9/17/13

*calc_complaints.cal...

1

Semantics

Projection

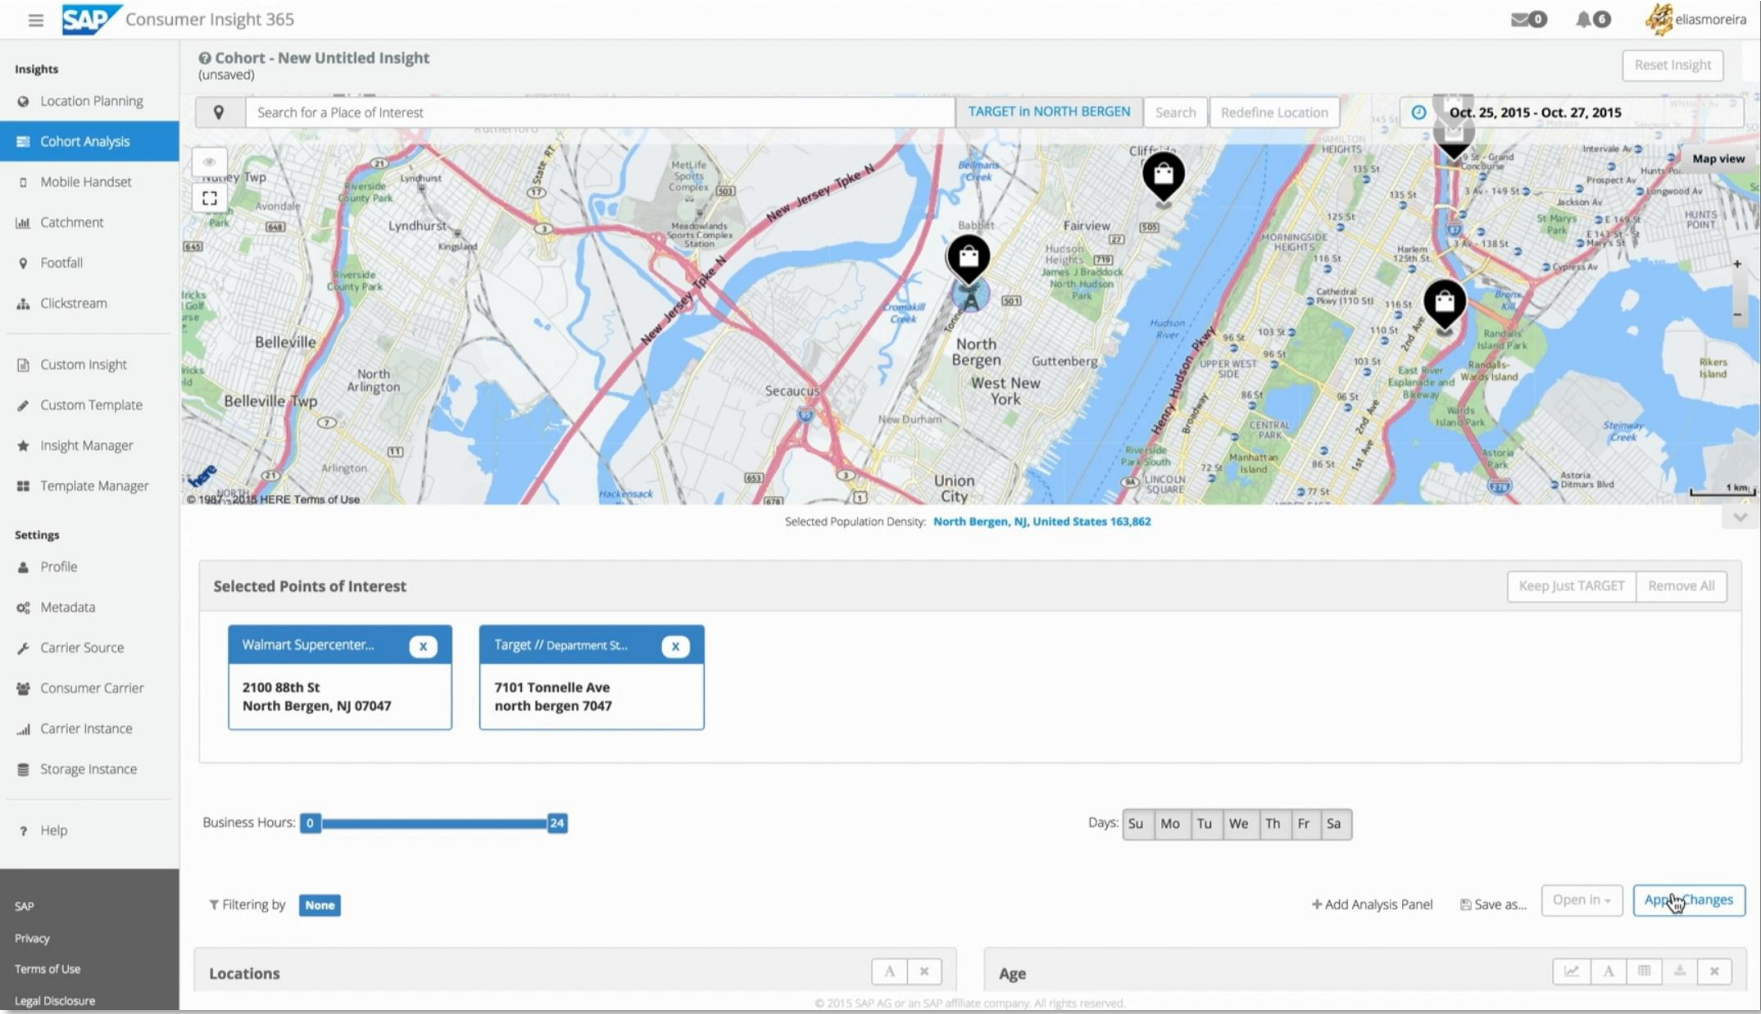
100%

Join_1

2

Consume Data Driven Application

Business User



Stable.Fast.Predictable.

SAP Cloud Platform Big data Service

IT Operations



Stable & Fast:

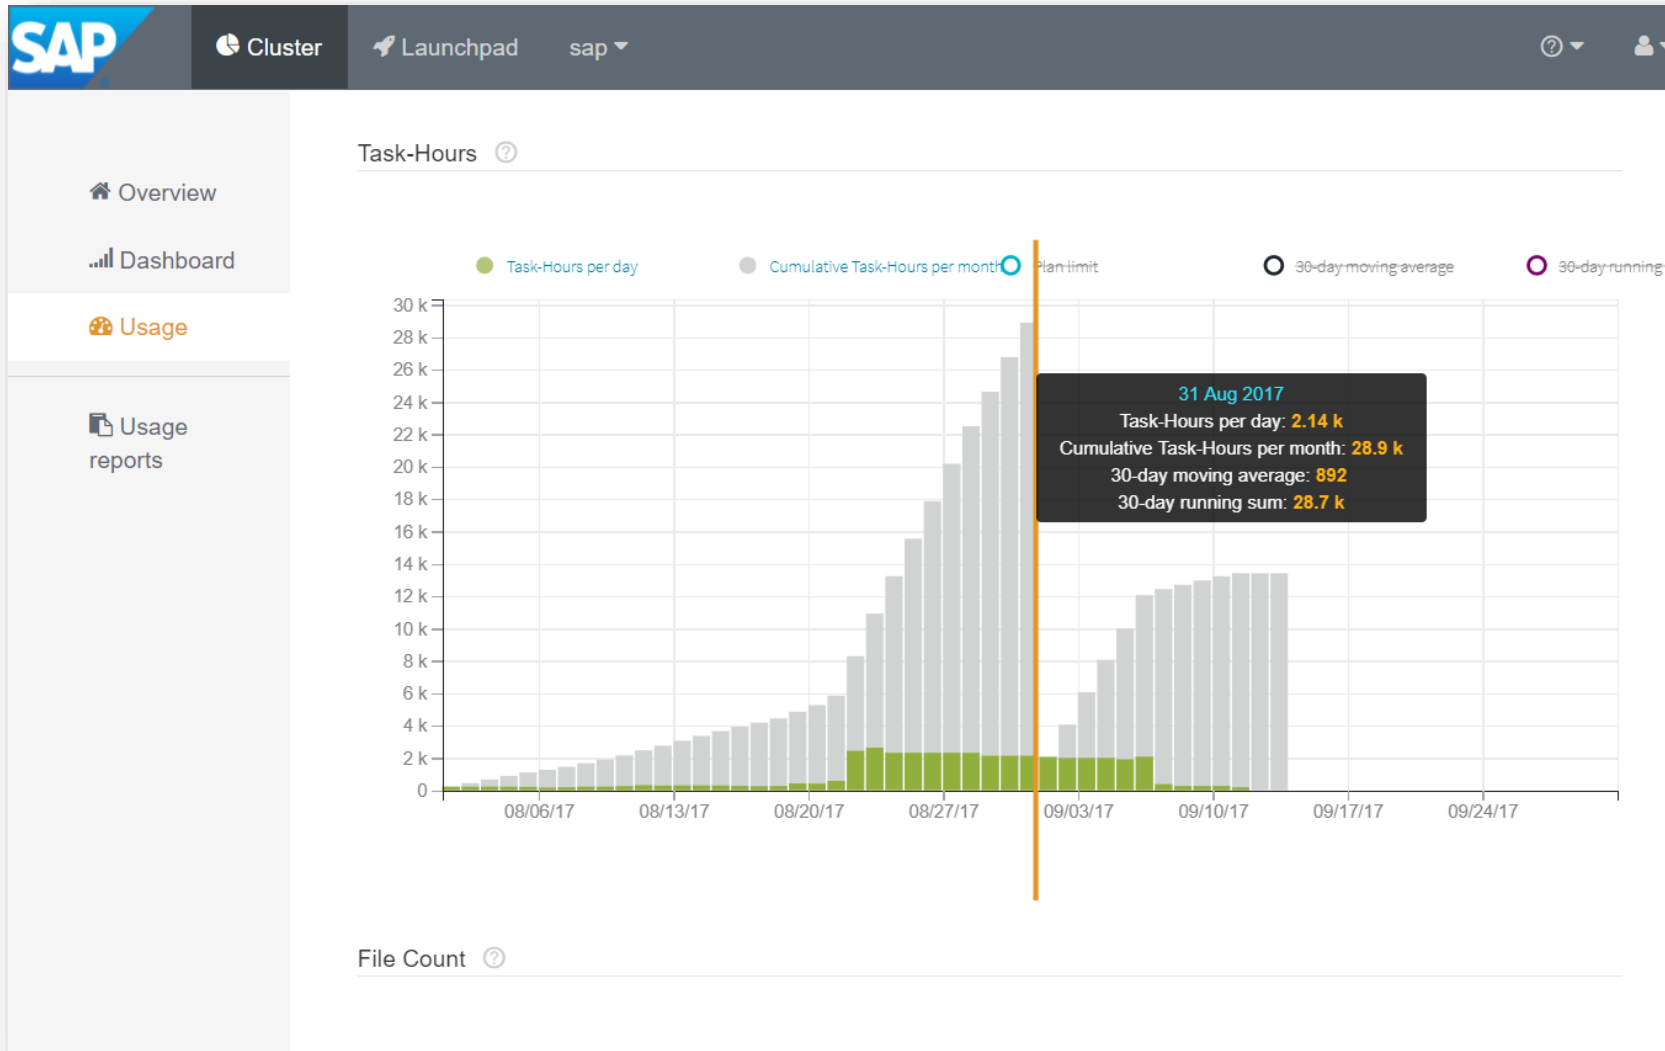
- Jobs execute and
- Execute faster

Autoscaling

- Always ON
- Compute Burst

Cost & Ease

- Predictable & Transparent
- Consumption NOT Hardware
- No sizing needed



And now what ... ?

openSAP

Enjoy! Good luck!
**Big Data with
SAP HANA Vora**
See you soon!

HANA Express ([here](#))
Vora Developer Edition ([here](#))
Vora GitHub Code Snippets ([here](#))
HANA Academy on [YouTube](#)
Open SAP Course ([here](#))

Self Learning



Design Thinking



SAP Tech Academy

Thank you.

Contact information:

Karsten Haldenwang

Center of Excellence Database and Data Management
Middle & Eastern Europe

Solution Advisor HANA, Vora and Hadoop

karsten.Haldenwang@sap.com
+49-151-53858452