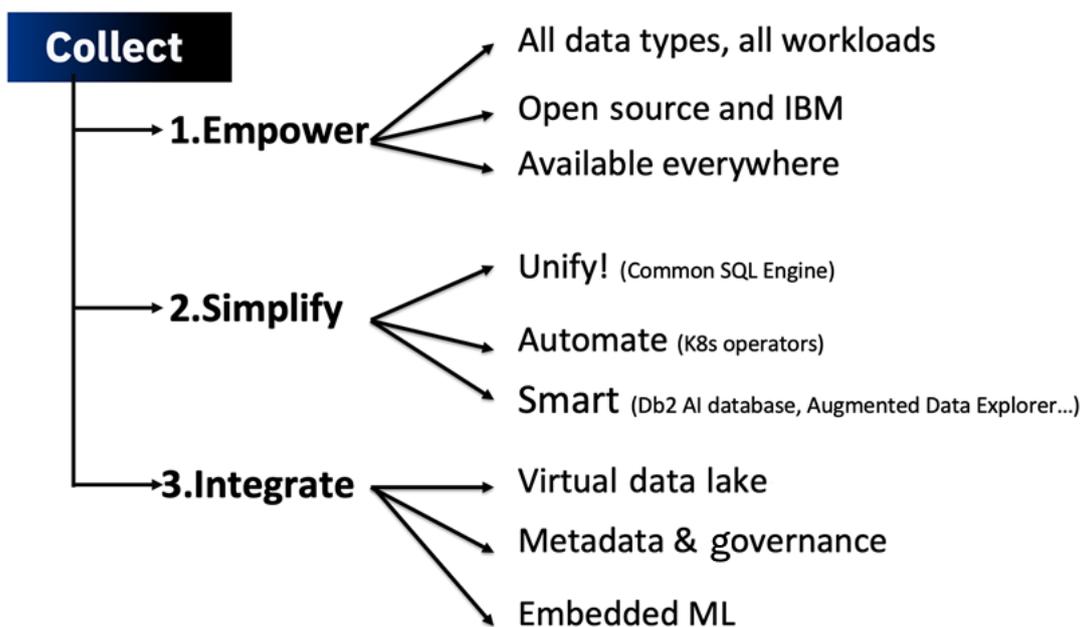
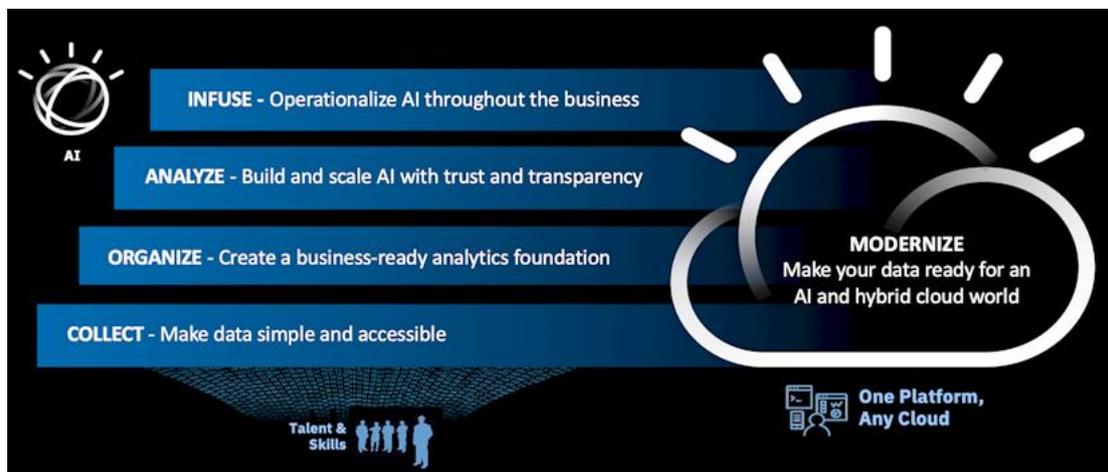
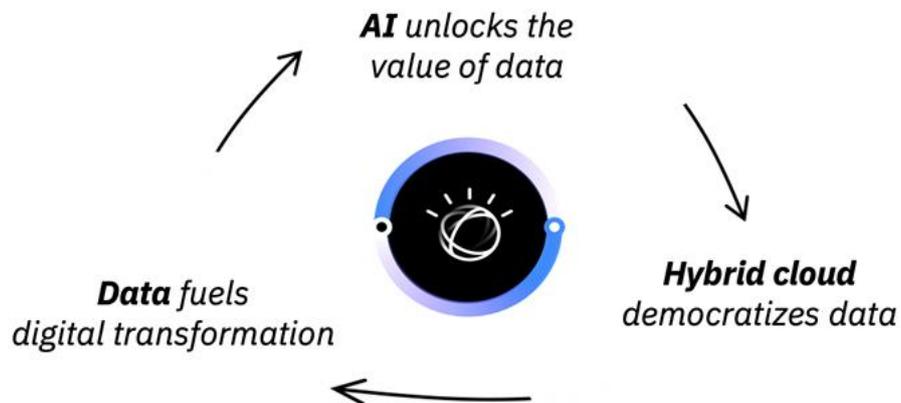
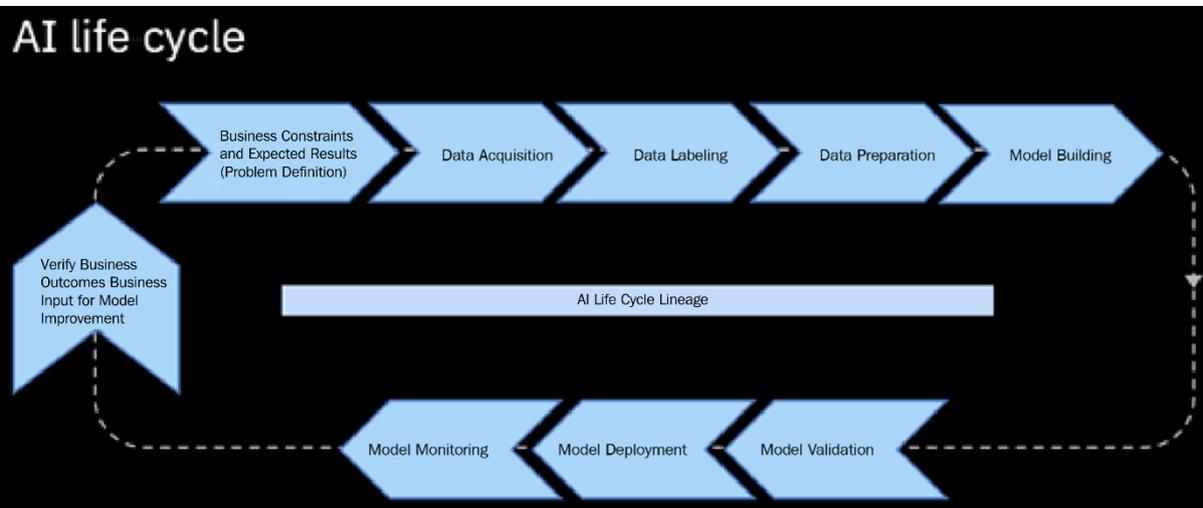
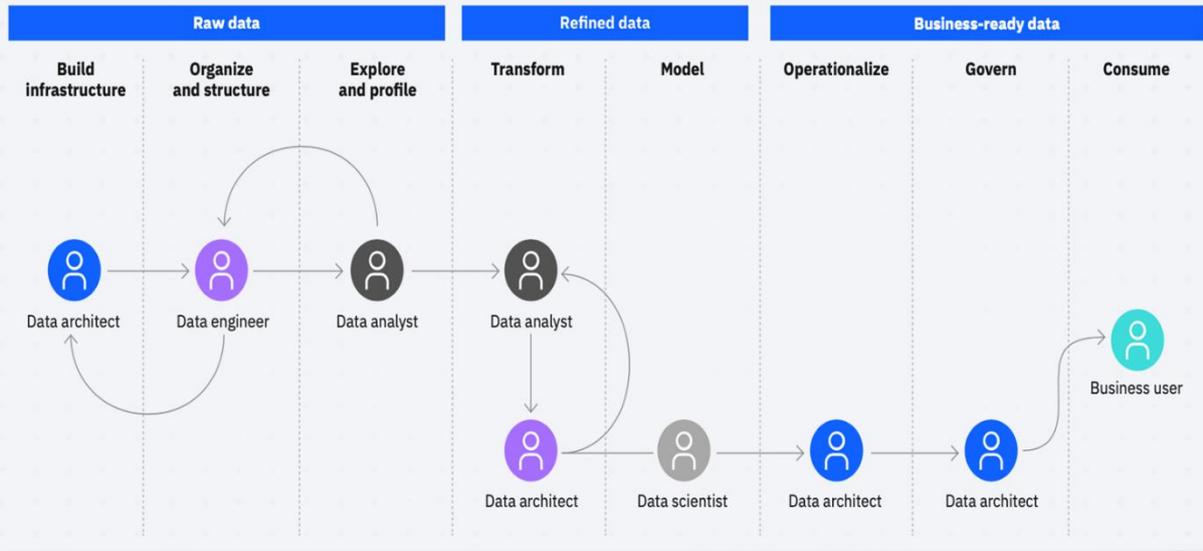
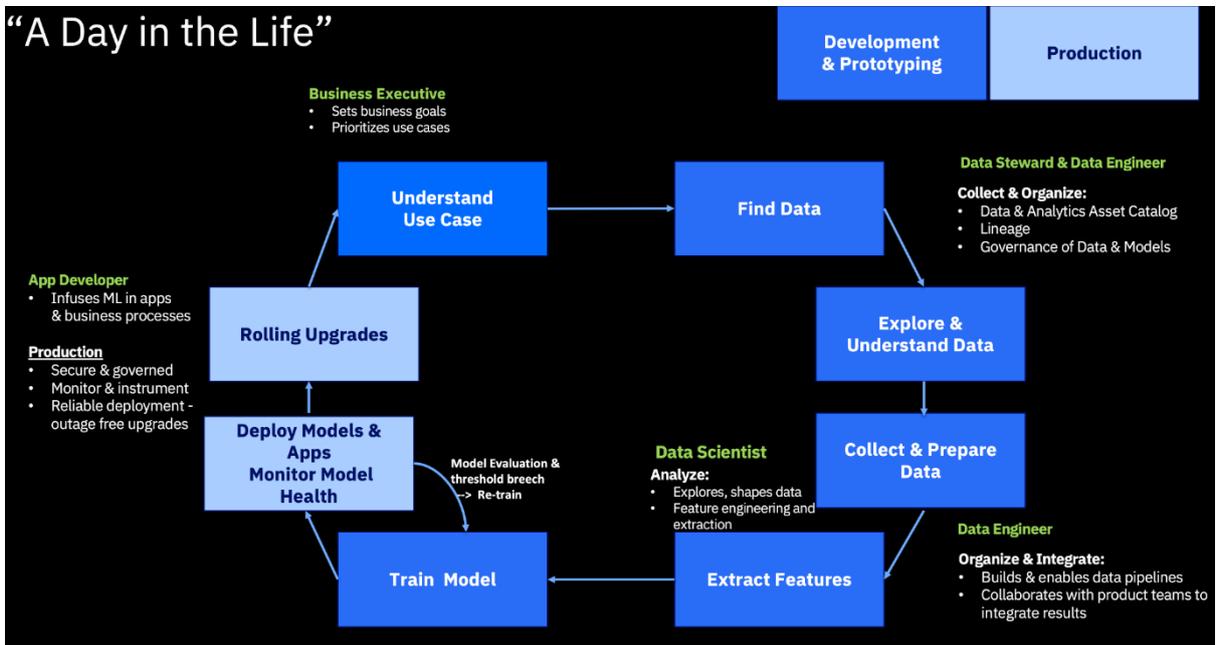
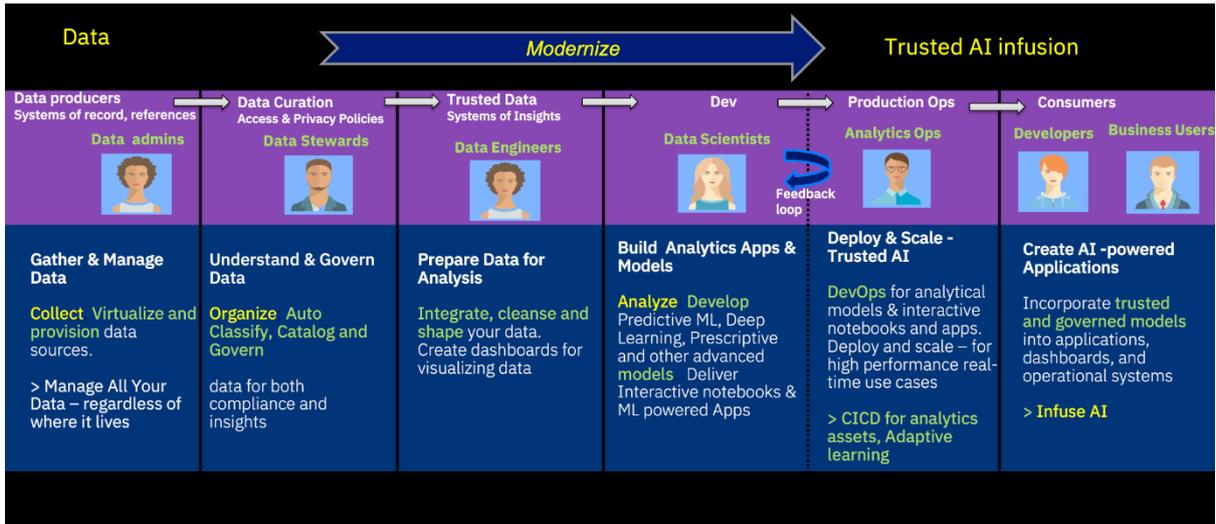


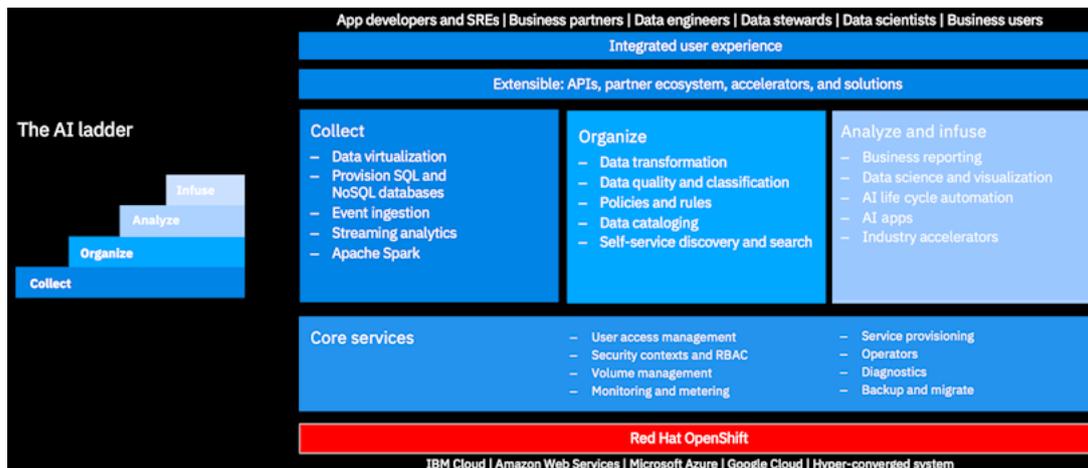
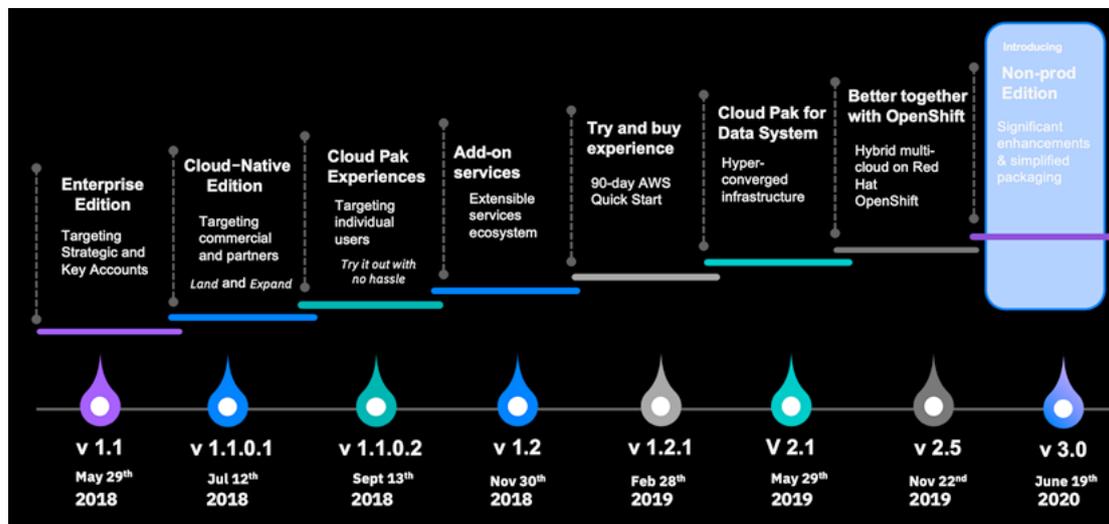
# Chapter 1: The AI Ladder: IBM's Prescriptive Approach





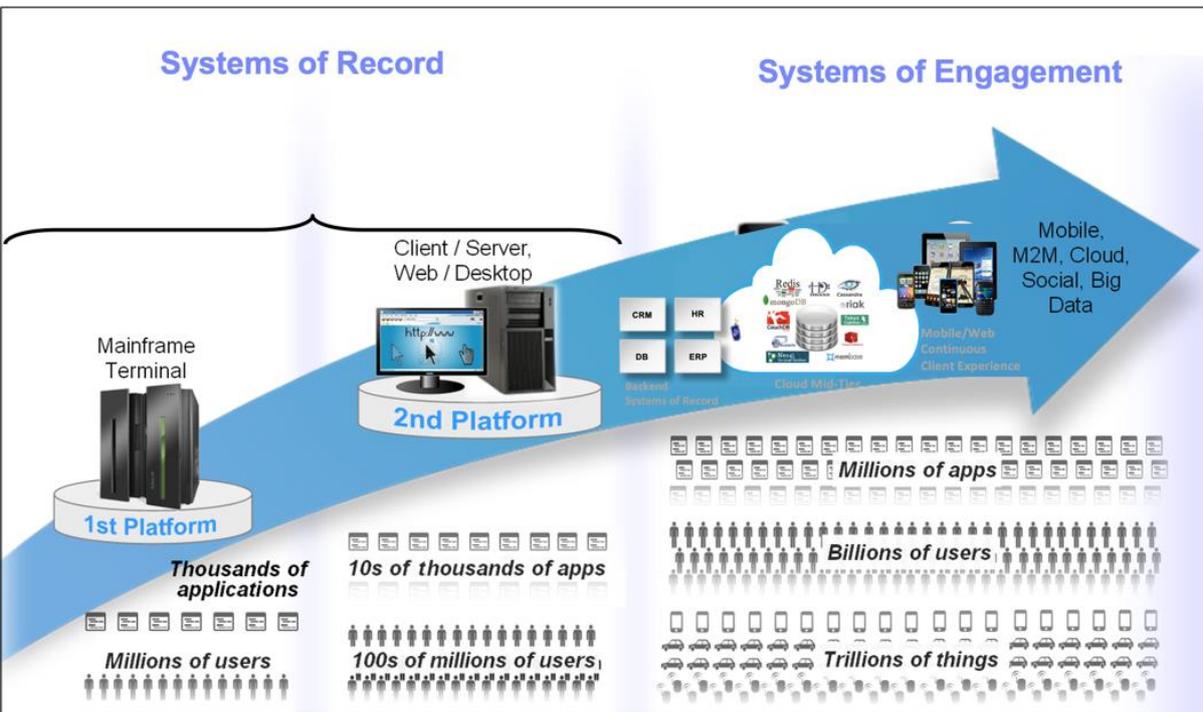
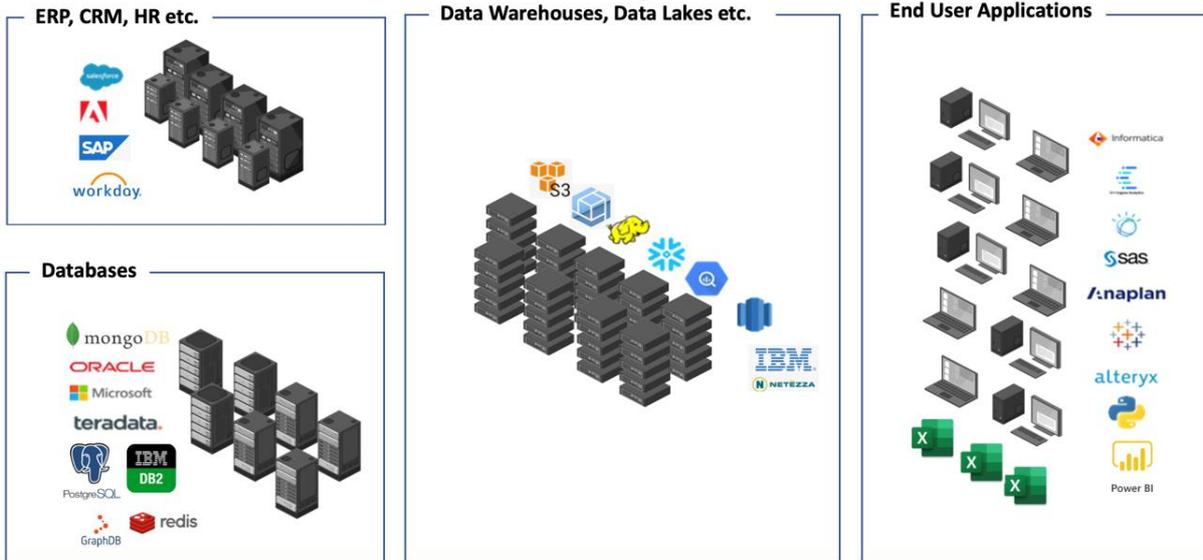


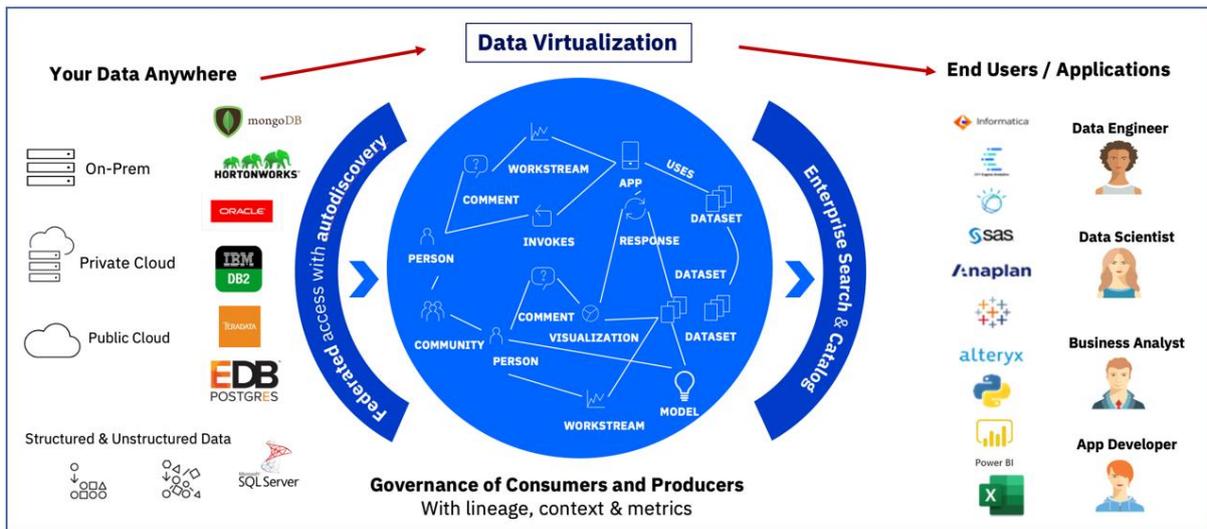
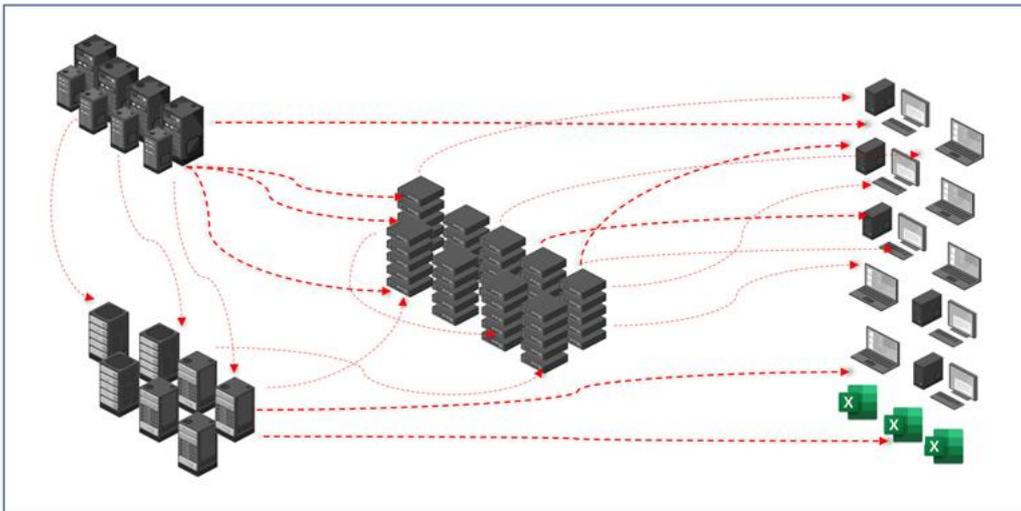
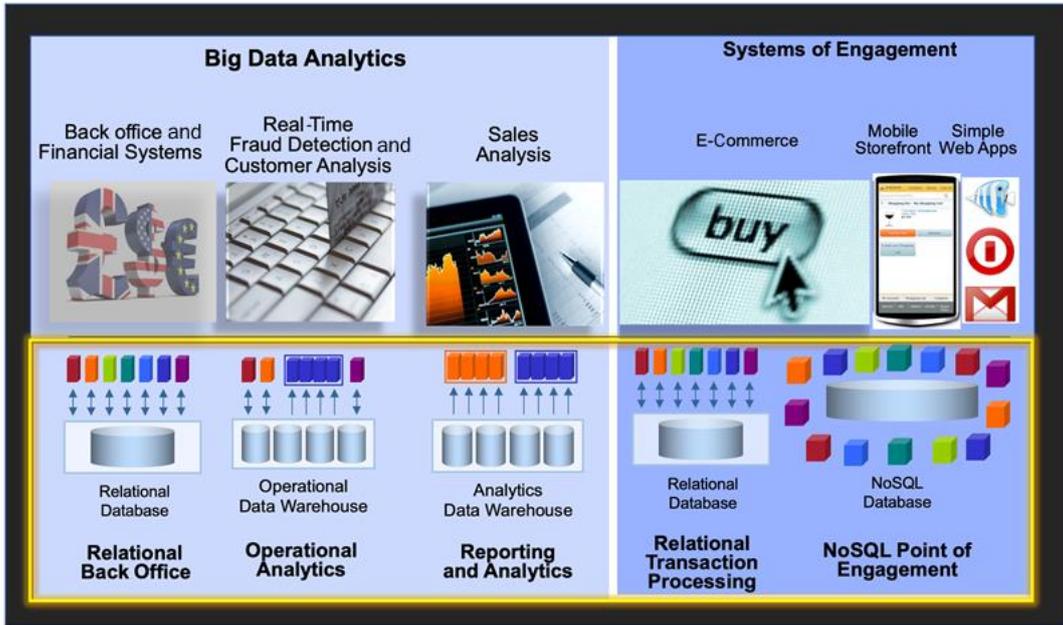
## Chapter 2: Cloud Pak for Data: A Brief Introduction

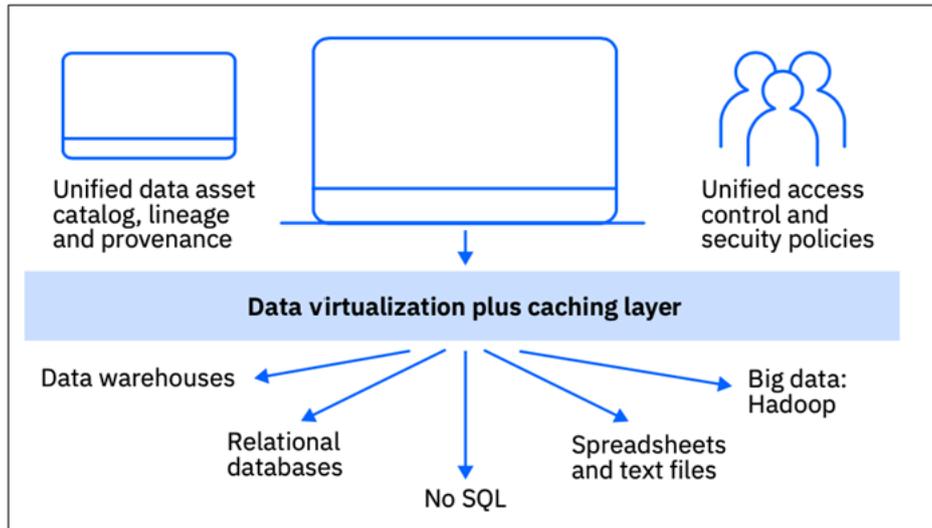


	Kubernetes	OKD	RedHat OpenShift
<b>Platform</b>			
Automated node configuration and cluster updates		X	X
Multi-host container scheduling	X	X	X
Self-Service provisioning	X	X	X
Service discovery	X	X	X
Enterprise operating system			X
Image Registry		X	X
Validated storage plugins		X	X
Validated networking plugins		X	X
Monitoring		X	X
Log aggregation		X	X
Multi-tenancy		X	X
Metering and chargeback			X
<b>Developer Exterience</b>			
Cloud service broker		X	X
Automated image builds		X	X
CI/CD and DevOps workflows		X	X
Validated third-party Kubernetes operators.			X
Certified databases			X
Certified middleware			X
Certified ISV Solutions			X
Serverless applications (Knative)		X	X
<b>Operations</b>			
Built-on operational management			X
Zero downtime patching and upgrades			X
Enterprise 24/7 support			X
9-year support life cycle			X
Security response team			X

# Chapter 3: Collect - Making Data Simple and Accessible







IBM Cloud Pak for Data All

### New connection

Select the type of data source that you want to connect to.

**IBM**

Analytics Engine HDFS	Data Virtualization Manager for z/OS	Db2 Hosted	Informix
Cloud Object Storage	Databases for PostgreSQL	Db2 on Cloud	Netezza (PureData System for Analytics)
Cloud Object Storage (infrastructure)	Db2	Db2 Warehouse	Planning Analytics
Cloudant	Db2 Big SQL	HDFS via Execution Engine for Hadoop	Storage volume
Cognos Analytics	Db2 Event Store	Hive via Execution Engine for Hadoop	
Compose for MySQL	Db2 for i	IBM SPSS Analytic Server	
Data Virtualization	Db2 for z/OS	Impala via Execution Engine for Hadoop	

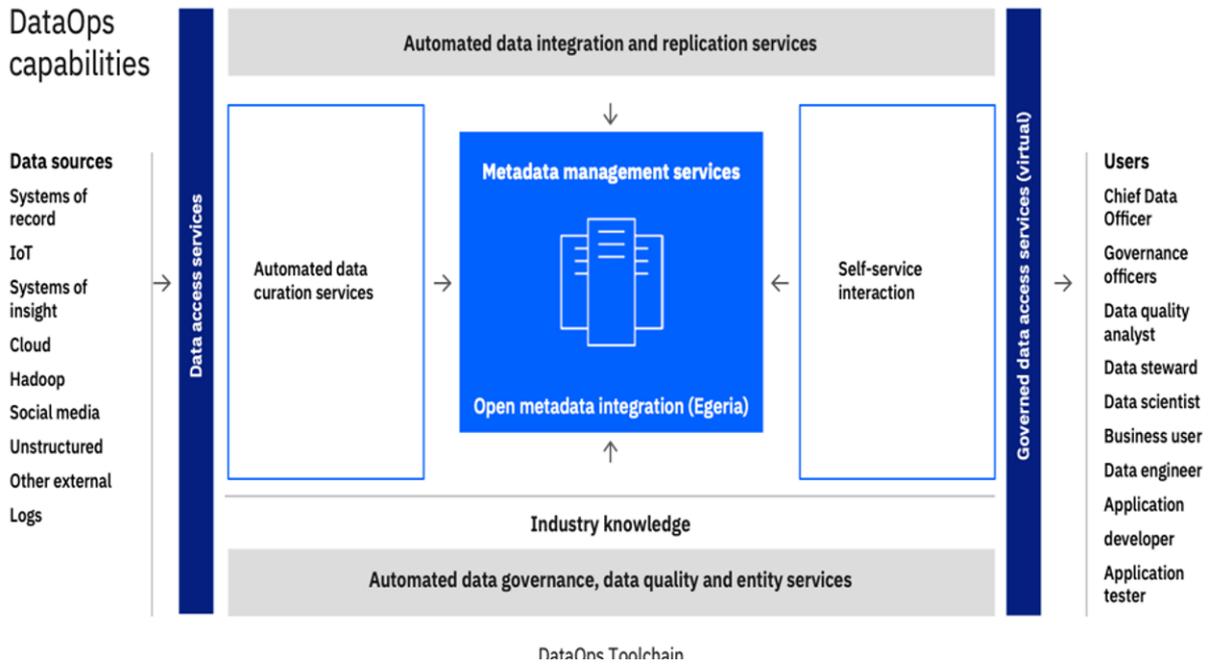
**Third-party**

Amazon RDS for MySQL	Dropbox	Microsoft Azure Data Lake Store	Salesforce.com
Amazon RDS for PostgreSQL	Elasticsearch	Microsoft Azure SQL Database	SAP HANA
Amazon Redshift	FTP	Microsoft SQL Server	SAP OData
Amazon S3	Google BigQuery	MinIO	Snowflake
Apache Cassandra	Google Cloud Storage	MongoDB	Sybase
Apache Derby	HTTP	MySQL	Sybase IQ
Apache HDFS	Looker	OData	Tableau
Apache Hive	MariaDB	Oracle	Teradata
Box	Microsoft Azure Blob Storage	Pivotal Greenplum	
Cloudera Impala	Microsoft Azure Cosmos DB	PostgreSQL	

**User-defined** ⓘ

Generic JDBC
--------------

# Chapter 4: Organize - Creating a Trusted Analytics Foundation



IBM Cloud Pak for Data | All | Search | [Notifications] | [Profile]

Business terms / Abandonment Rate

## Abandonment Rate Published

Overview | Related content

### General

**Description** ✎  
The number of applications begun on a company's recruitment platform, but not finished and submitted for consideration.

**Primary category** 📄 ✎  
Retail  
Categories / GoSales

**Secondary categories** 📄 ⊕  
No secondary category added yet.

### Related terms

**Type relationships**

**Is a type of** Show hierarchy ⊕  
No business term added yet.

**Has a type of** ⊕  
No business term added yet.

**Part relationships**

**Is a part of** ⊕  
No business term added yet.

**Has a part of** ⊕  
No business term added yet.

### About this business term

**Stewards** ⊕  
datasteward  
datasteward@ibm.com

**Effective dates** ✎

**Start date**  
Jan 19, 2021 12:01 AM

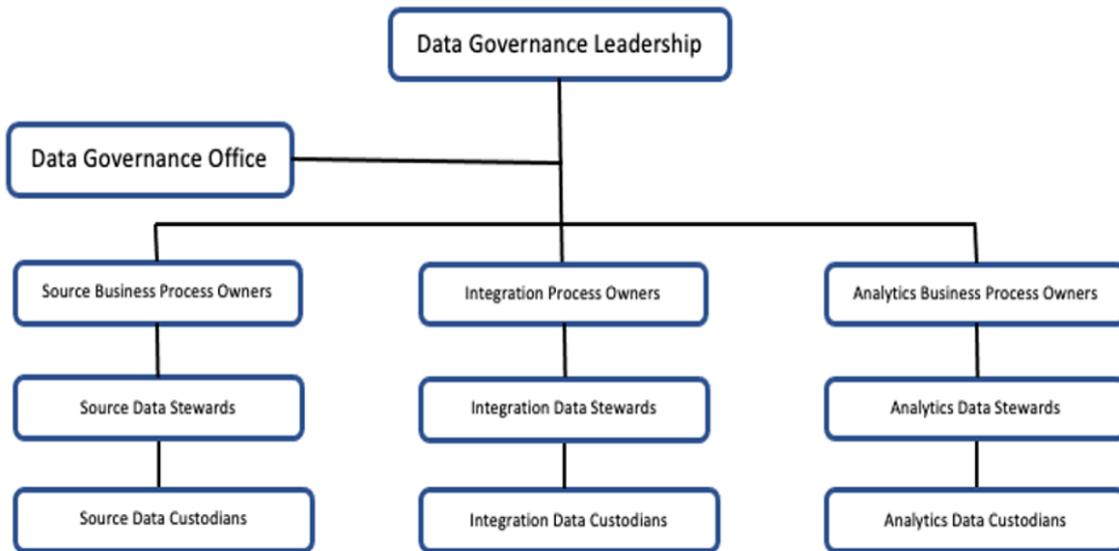
**End date**  
Jan 4, 2024 11:59 PM

**Tags** ⊕  
No tags added yet.

---

**Created by**  
datasteward  
Jan 18, 2021 11:15 PM

**Modified by**  
datasteward  
Jan 18, 2021 11:15 PM



IBM Cloud Pak for Data | All | Search

Catalogs / Platform assets catalog / Sales Fact Data | Add to Catalog

Data asset: Sales Fact Data

Remove | Download | Add to project

Overview | Asset | Access | Review | **Profile** | Activities



**Data profile not yet created**  
 You can create a profile from the first 5000 rows of the data set.  
[Select data classes](#)  
[Create profile](#)

IBM Cloud Pak for Data | All | Search

Catalogs / Platform assets catalog / Sales Fact Data | Add to Catalog

Data asset: Sales Fact Data

Overview | Asset | Access

**Data classes** [X]

Deselect all | Search

165 of 165 selected

**Available data classes**

- Health Insurance Claim Number
- US Social Security Number
- UK Province Code
- Computer Host Name
- US County
- Middle Name

Classifier: Computer Host N:  
 Description: Hostname is a las assigned to a device connectcomputer network and is used to identify.

Cancel | Apply

Data asset

### Sales Fact Data

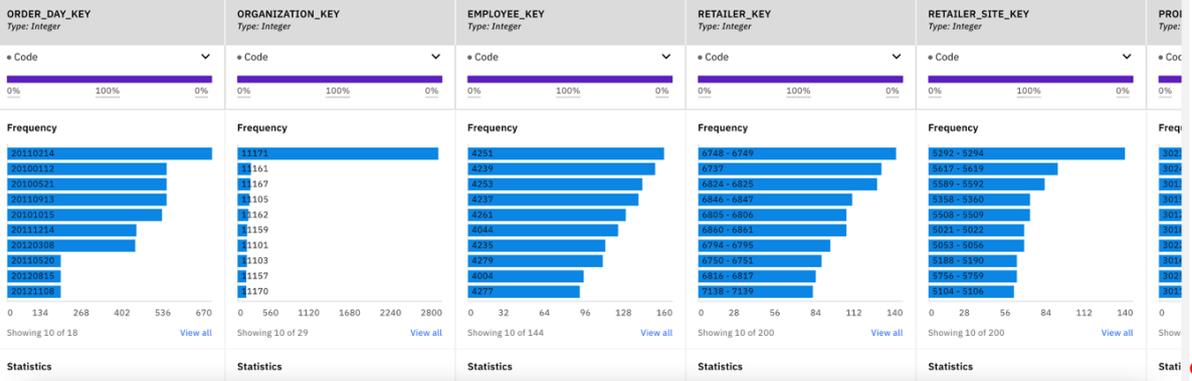
Remove Download Add to project

Overview Asset Access Review **Profile** Activities

Current profile Last profile Columns Rows ■ Matches ■ Mismatches ■ Missing

18 classifiers 18 Jan 2021 - 11:58 pm View log 18 5,000

Delete profile Update profile



## Restrict Confidential Data Access

Edit rule

Delete rule

### About this rule

**Business definition**

Corporate guidelines require that confidential data be limited to employees in the business unit.

**Type**

Access

**Created by**

admin  
Jan 19, 2021 11:10 PM

**Modified by**

admin  
Jan 19, 2021 11:10 PM

**Criteria**

**Condition 1**

If **Classification** contains any **Confidential**

**Action**

Then deny access to data

# Quick scan job

Cancel Discover

Connection ⓘ  
 GoSalesConnection

Discovery Root (Optional) ⓘ  
 schema[GSDDB|GOSALES] [Browse](#)

Project ⓘ  
 GoSalesDataQuality

**Discovery options**

- Analyze columns ⓘ
- Analyze data quality ⓘ
- Assign terms ⓘ
  - Use machine learning to assign terms ⓘ
- Use data sampling

The maximum number of records included in the data set sample:  
 1000

Data discovery / Job results

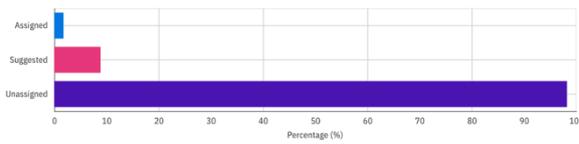
## Job results

Job: qs\_1611131080281

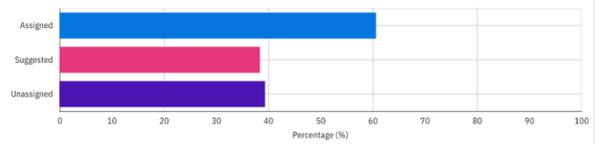
[Explore assets](#)

Discovery insights Data quality insig...

### Business term assignment



### Data class assignment



### Business terms and data classes

5 terms, 18 data classes

Name	Business terms			Data classes			
	Columns	Tables	Schemas	% assigned data class	% suggested data class	Assigned data class	Suggested data class
Branch Code	5	0	0	80%	0%	Indicator + 1 more	-
City	2	0	0	100%	100%	City	Province Name + 3 more
Address line 1	1	0	0	100%	100%	Address Line 1	Text + 1 more
Address line 2	1	0	0	0%	100%	-	US Street Name

### Assets

Job: qs\_1611131080281

**Asset type**

File

Schema

Table

Column

---

**Filters**

**Tables**

BRANCH (13)

INVENTORY\_LEVELS (10)

ORDER\_DETAILS (9)

ORDER\_HEADER (10)

ORDER\_METHOD (30)

[See 19 more](#)

**Schemas**

GOSALES (508)

**Status**

Ready for review (508)

**Quality score (0 - 100)**

Discovered columns (508)

Column name	Context	Quality	Assigned data class	Suggested data class	Assigned business term	Suggested business term	Actions
ADDITIONS	INVENTORY_LEVELS	100%	-	US Zip Code 2%	-	-	<a href="#">🗑️</a>
ADDRESS1	BRANCH	96.6%	Address Line 1 100%	US St... 31% + 1 more	Address line 1 82%	Email Address 50%	<a href="#">🗑️</a>
ADDRESS1_MB	BRANCH	93.3%	Text 100%	US Street Name 24%	-	Address line 1 67%	<a href="#">🗑️</a>
ADDRESS2	BRANCH	78.6%	-	US Street Name 22%	Address line 2 82%	Email Address 50%	<a href="#">🗑️</a>
ADDRESS2_MB	BRANCH	78.6%	-	US Street Name 22%	-	Address line 2 67%	<a href="#">🗑️</a>
ASSIGNED_TO	RETURNED_ITEM	100%	-	-	-	-	<a href="#">🗑️</a>
AVERAGE_UNIT_COST	INVENTORY_LEVELS	100%	Quantity 100%	-	-	-	<a href="#">🗑️</a>
BASE_PRODUCT_NU...	PRODUCT_FORECAST	100%	-	Boolean 4%	-	-	<a href="#">🗑️</a>
BASE_PRODUCT_NU...	PRODUCT	100%	-	-	-	-	<a href="#">🗑️</a>
BRANCH_CODE	BRANCH	100%	Identifier 100%	-	Branch Code 100%	-	<a href="#">🗑️</a>

### Assets

Job: qs\_1611131080281

**Asset type**

File

Schema

Table

Column

---

**Filters**

**Schemas**

GOSALES (24)

**Status**

Ready for review (24)

**Quality score (0 - 100)**

0  100

Min Max

**Published to**

No filters of this type

Discovered tables (24)

1 item selected Publish assets [🗑️](#) Audit assets [🗑️](#) Cancel

Table name	Context	Quality	Status	Published to	Connection name	Details
<input checked="" type="checkbox"/> BRANCH	GOSALES	87.8%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> CONVERSION_RATE	GOSALES	100%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> COUNTRY	GOSALES	97.5%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> CURRENCY_LOOKUP	GOSALES	98.9%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> EURO_CONVERSION	GOSALES	100%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> INVENTORY_LEVELS	GOSALES	98.7%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> ORDER_DETAILS	GOSALES	98.7%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> ORDER_HEADER	GOSALES	99.6%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> ORDER_METHOD	GOSALES	100%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>
<input type="checkbox"/> PRODUCT	GOSALES	98.6%	🗑️ Ready for review	-	GoSalesConnection	<a href="#">View details</a>

Items per page 10  1-10 of 24 items 1 of 3 pages [⏪](#) [⏩](#)



### Chief Risk Officer's mission statement

We need to have a full audit trail for sensitive personal information. We need to know who, what, when, and where.

Reach out to Nadine, our lead IBM Guardium admin, with questions.



### IBM Cloud Pak for Data admin

My application is connected to numerous data sources that contain sensitive personal information (SPI).



### IBM Guardium admin

I can give you information about the IBM Guardium appliances that we use to audit the data sources connected to your application.

Let me generate the client secret that IBM Cloud Pak for Data can use to connect to IBM Guardium.



### IBM Cloud Pak for Data admin

Thanks for your help. I've configured a connection to each appliance and added the groups we discussed.



### IBM Cloud Pak for Data admin

We're going to use IBM Guardium to audit assets that contain SPI. Please use the auditing feature in IBM Cloud Pak for Data.



### Data steward

OK. I'll make sure to add the assets to IBM Guardium groups when I review the discovery results. That way, we can audit the assets.

☰ IBM Cloud Pak for Data All ▼ 🔍 Search 🔔

## Automated discovery job

Cancel

Discover

Connection ⓘ

GoSalesConnection ▼

Discovery Root (Optional) ⓘ

table[GSDB[GOSALES|PRODUCT]];table[GSDB[GOSALES|ORDER\_METHOD]];table[GSDB[GOSA

[Browse](#)

Project ⓘ

GoSalesDataQuality ▼

[Go to project settings](#) →

### Discovery options

- Analyze columns ⓘ
- Analyze data quality ⓘ
- Assign terms ⓘ
- Publish results to catalog ⓘ
- Use data sampling

Set the maximum number of records that you want to include in your data set sample:

100

### Select the method that you want to use to create your sample:

- Use the first x number of rows (where x = maximum number of records allowed) ⓘ
- Use every Nth value (up to maximum number of records) ⓘ

**GoSalesDataQuality**  
Not governed

Description  
GoSales project to hold data quality results

Data assets: 33  
Reviewed data assets: 0/33

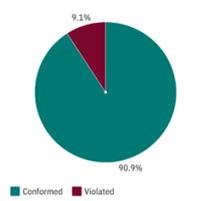
Connections: 1  
Created by: admin  
Last modified: Jan 20, 2021

Sampling enabled

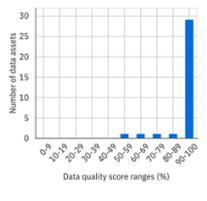
Project  
**GoSalesDataQuality**

Dashboard Data assets Data rules Relationships Settings

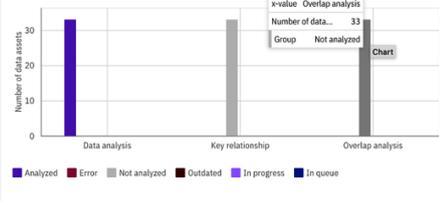
**Data quality threshold**



**Data quality score distribution**



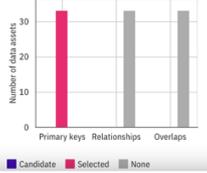
**Analysis**



**Rule run status**

Create and run a rule to see its status

**Relationships**



**Top 5 selected classes**



## Create a quality rule

Rule definition  Bindings  Summary and test

**Rule expression**

qtyValue >= qtyMinValue AND qtyValue <= qtyMaxValue

**Variables**

3 of 3 bindings complete

Bind

Variable name	Data type	Implemented bindings
<input type="radio"/> qtyminvalue	ANY	100
<input type="radio"/> qtyvalue	ANY	ORDER_DETAILS.QUANTITY
<input checked="" type="radio"/> qtymaxvalue	ANY	1000

**Available data sources**

Drag data to a variable, or select data and click Bind. You can bind only one data source to a variable.

Columns Global variables Literals

Data Type  
 String  Numeric  Date  Time

Literal value  
 1000

Cancel Back Next

Projects / GoSalesDataQuality / Data assets / ORDER\_DETAILS / Columns

### QUANTITY

Description: No description.

Data quality score: 61% (36% target)  
Rows: 1000 (900% target)

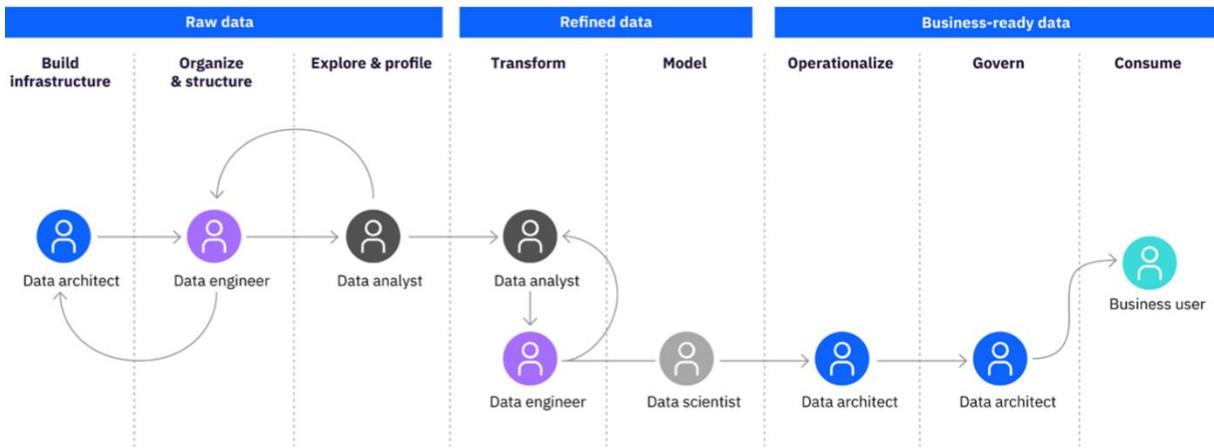
Analysis status: Completed  
Last column analysis: Jan 20, 2021  
Last DQ analysis: Jan 20, 2021

Name (p...)	Score	Delta
ORDER_DETAI	37%	62%
ORDER_NUM...	37%	62%
SHIP_DATE	100%	0%
PRODUCT_N...	93%	7%
PROMOTION...	97%	3%
<b>QUANTITY</b>	<b>61%</b>	<b>36%</b>
UNIT_COST	91%	3%
UNIT_PRICE	90%	2%
UNIT_SALE...	90%	3%

#### Data quality dimension results

Showing 11 dimensions

Dimension name	# of findings ↓	% of findings	Delta	Ignore
Rule violations: QtyInvalidRange	387	39	0 changes	Off
Suspect values	21	2	0 changes	Off
Values out of range	0	0	0 changes	Off
Suspect values in correlated columns	0	0	0 changes	Off
Missing values	0	0	0 changes	Off
Inconsistent representation of missing values	0	0	0 changes	Off
Inconsistent capitalization	0	0	0 changes	Off
Format violations	0	0	0 changes	Off
Duplicated values	0	0	0 changes	Off



Automation rules / Check order details for valid product quantity

Automation Rule details

## Check order details for valid product quantity

### Rule Logic

**Created By**  
admin

**Created On**  
20 January 2021, 4:29:47 pm

**Modified By**  
admin

**Modified On**  
20 January 2021, 4:29:47 pm

### Short description:

Check order details for valid product quantity

**Status:** Candidate

```

if [the asset] has the term GoSales/Retail/Product Quantity assigned
then use the data quality dimension Rule violations
    
```

IBM Cloud Pak for Data

Rules / Obfuscate Email Address Rule

### Obfuscate Email Address Rule

[Edit rule](#) [Delete rule](#)

**Criteria**

Condition 1

If Business term contains any Email Address

**Action**

Then Obfuscate data in columns containing Email Address

**Advanced masking options**

Custom options for this data class

**Reversibility**  
Irreversible masking

**Consistency**  
Repeatable output

**Input validation**  
Input validation

**About this rule**

Business definition  
Obfuscate Email Address Rule with Business Term criteria.

Type  
Access

Created by  
admin  
Jan 13, 2021 1:33 PM

Modified by  
admin  
Jan 14, 2021 11:09 AM

IBM Guardium

09:21 User Interface

Active Threat Analytics

Risk Spotter Outlier Mining Administration Apr 4, 2019 - Apr 5, 2019

Top Cases: High 1 Med 11 Low 31

Databases with open cases: High 1 Med 3 Low 2

DB Users with open cases: High 0 Med 4 Low 14

File Systems with open cases: High 0 Med 0 Low 0

File Users with open cases: High 0 Med 0 Low 0

DB User Behavioral Analytics

DB User: DB\_USER\_27  
Database: MYGDB  
Server IP: 9.42.17.165

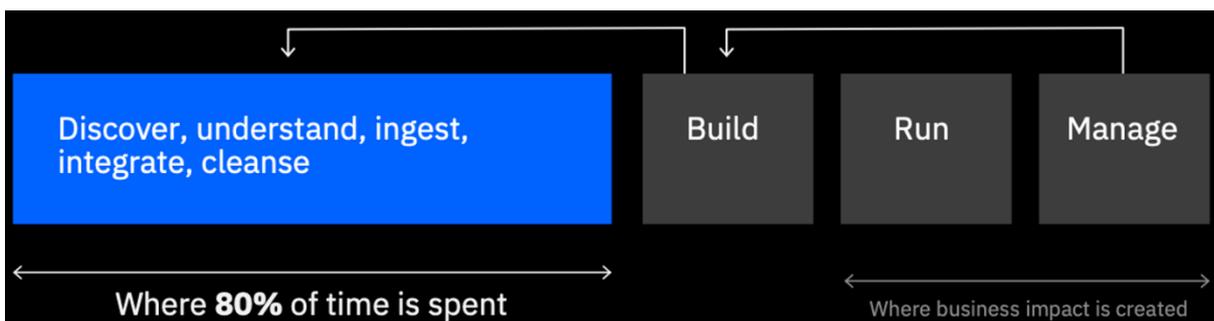
Tracked since: 2019-04-07 00:00  
Avg. activities per hour: 44929.95  
Max. activities in any one hour: 44953  
Avg. exceptions per hour: 59.09

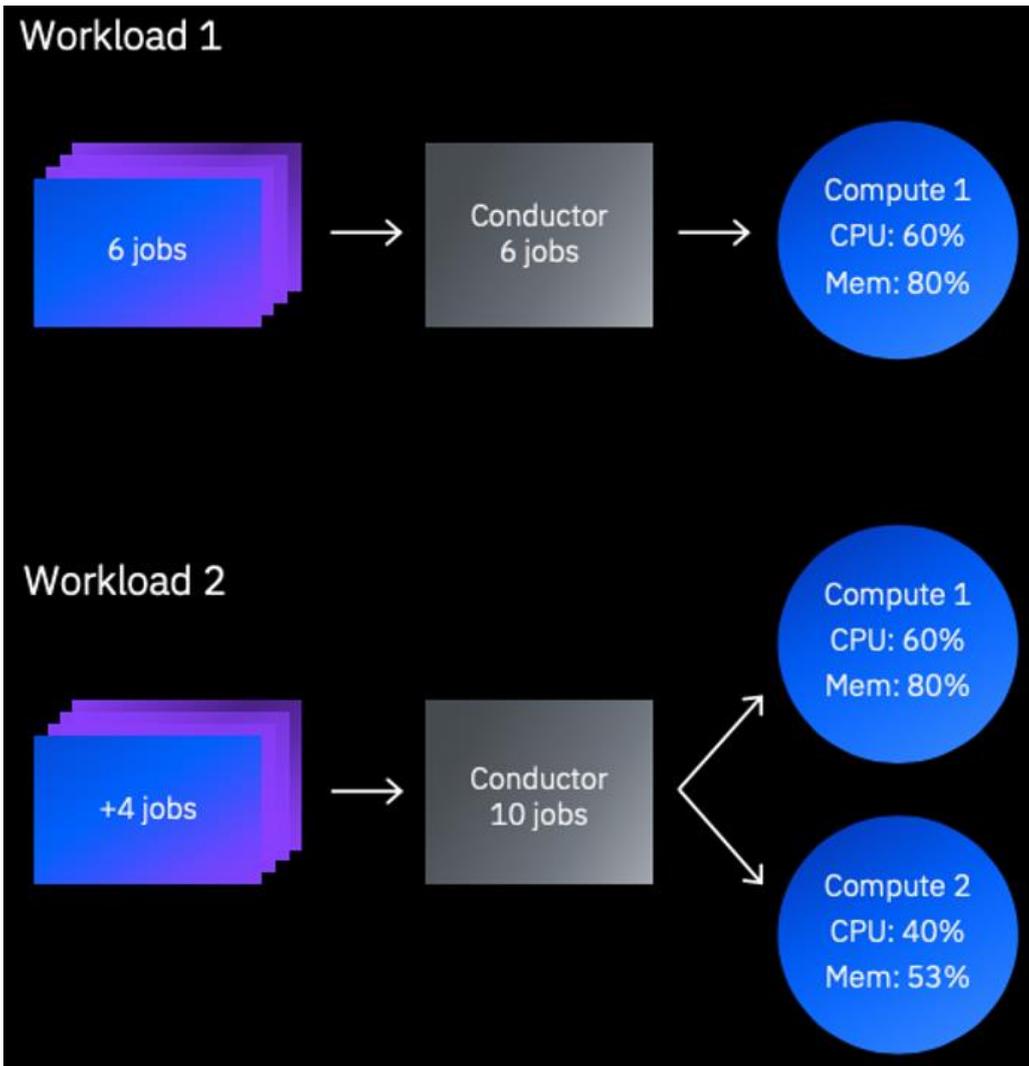
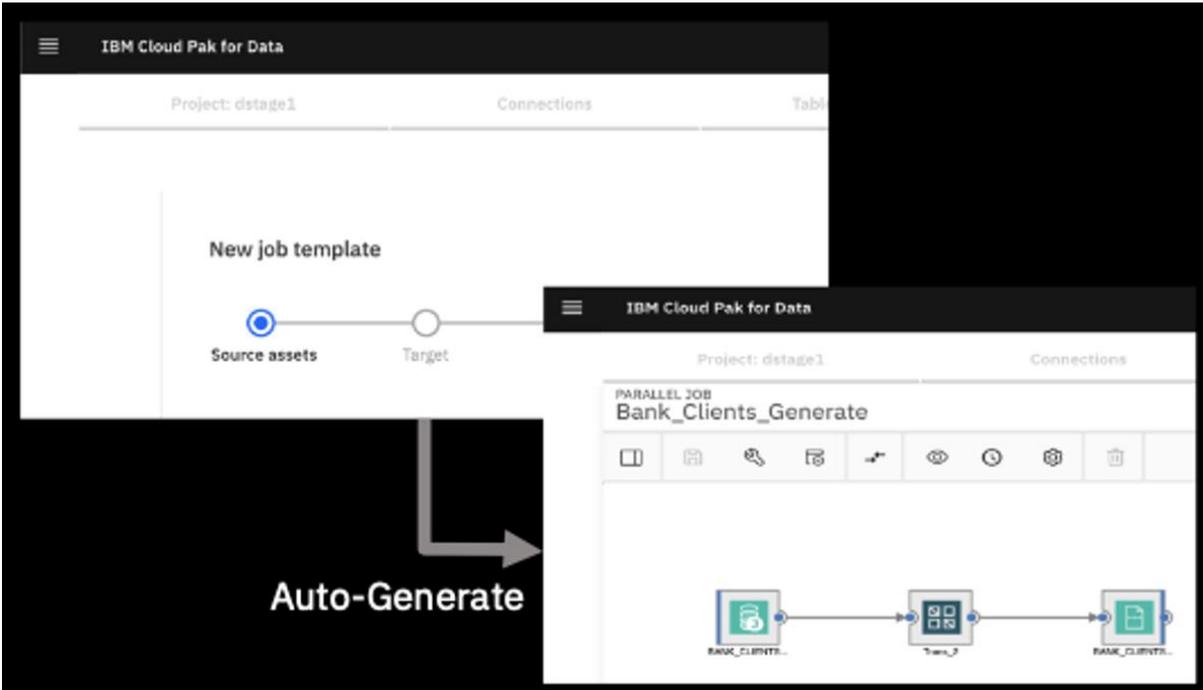
Distribution of Working Hours

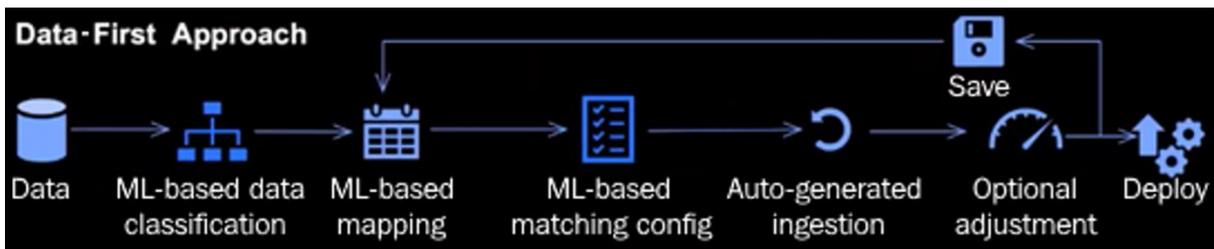
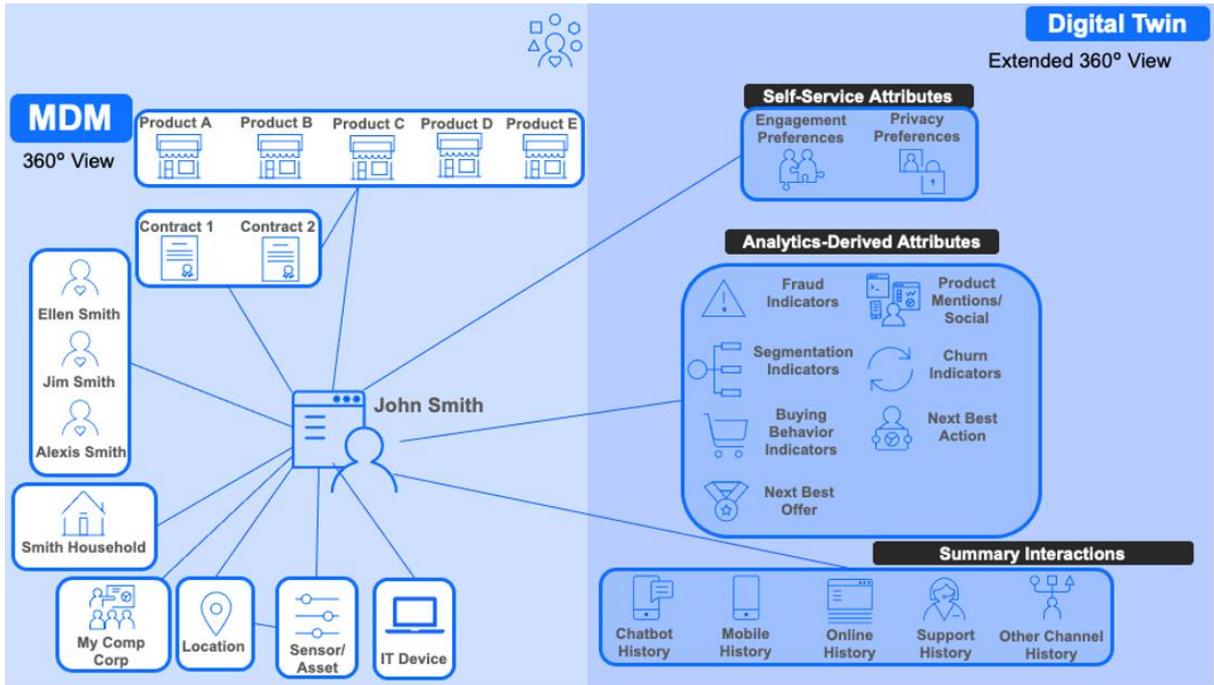
Distribution of Verbs

User Risk Indicators

Case #	Assign Case	Severity	Observations	Date	Assigned
20	Assign Case	Med	Exceptionally high volume of Exceptions	Apr 4, 2019 10:00 AM	No







IBM Cloud Pak for Data | All | Search | Upgrade | Michael Adams's Account

IBM Match 360 with Watson home / Master data configuration / Mapping

Data setup ▾

Assets | Mapping | Modeling

4-sources-data5.csv | Last mapping update: Apr 15, 2021, 1:45 PM | Auto map | Profile

Find an attribute or field | Filter by: All columns (11)

IN4	COLUMN5	COLUMN6	COLUMN7	COLUMN8	COLUMN9
sd to: Primary residence - s line 1 US Street Name teness : 99% e: string	Mapped Mapped to: Primary residence - City City Completeness : 100% Data type: string	Not Mapped Mapped to: US State Code Completeness : 100% Data type: string	Not Mapped Mapped to: US Zip Code Completeness : 99% Data type: int32	Not Mapped Mapped to: NoClassDetected Completeness : 99% Data type: string	Not Mapped US S Comp Data t
RICHELIEU ST.	PORT NORRIS	WI	8349	713-6906	788-
AST DELTA ST.	KEWAUNEE	WI	84216	850-0526	926-
I SIREN ST.	CHENEY	KS	67025	768-4917	129-
IORTHEAST CAMP ST.	PRAGUE	NE	68080	216-8165	806-
OUTH CALIFORNIA ST.	TUCSON	AZ	85708	267-3352	462-
CALIFORNIA ST.	TUCSON	AZ	85708	267-3352	462-

Mapping targets

COLUMN6 - Primary residence - State/Province value

- Primary residence - State/Province value (100%)
- Legal name - Prefix (7%)
- Mailing address - Street Name (5%)

# Master data explorer ▼

Export

Search    Explore    Settings

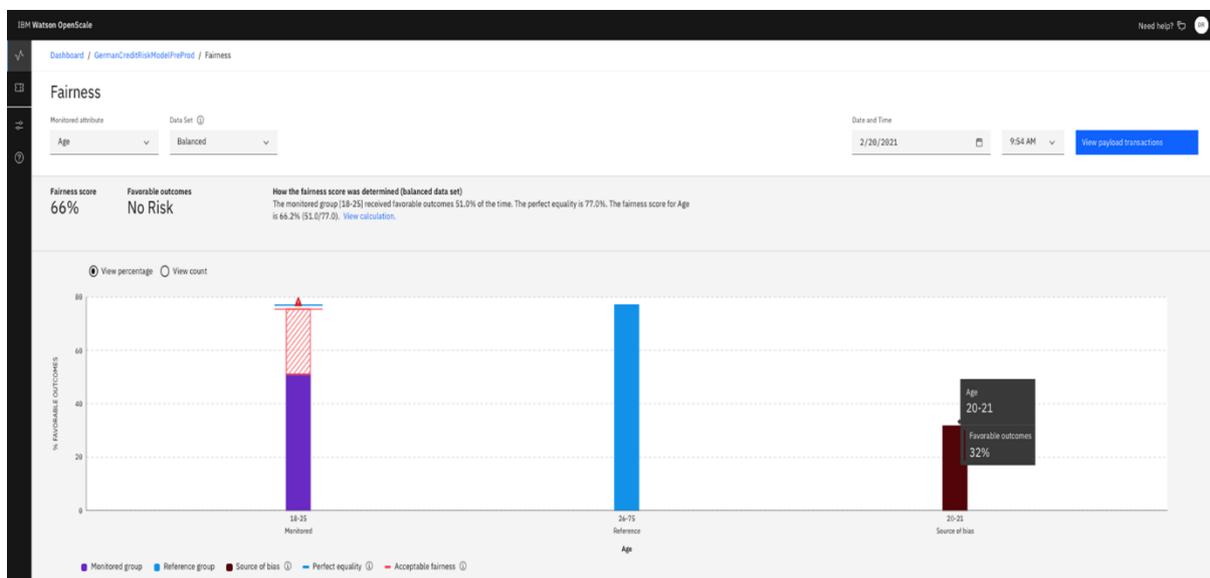
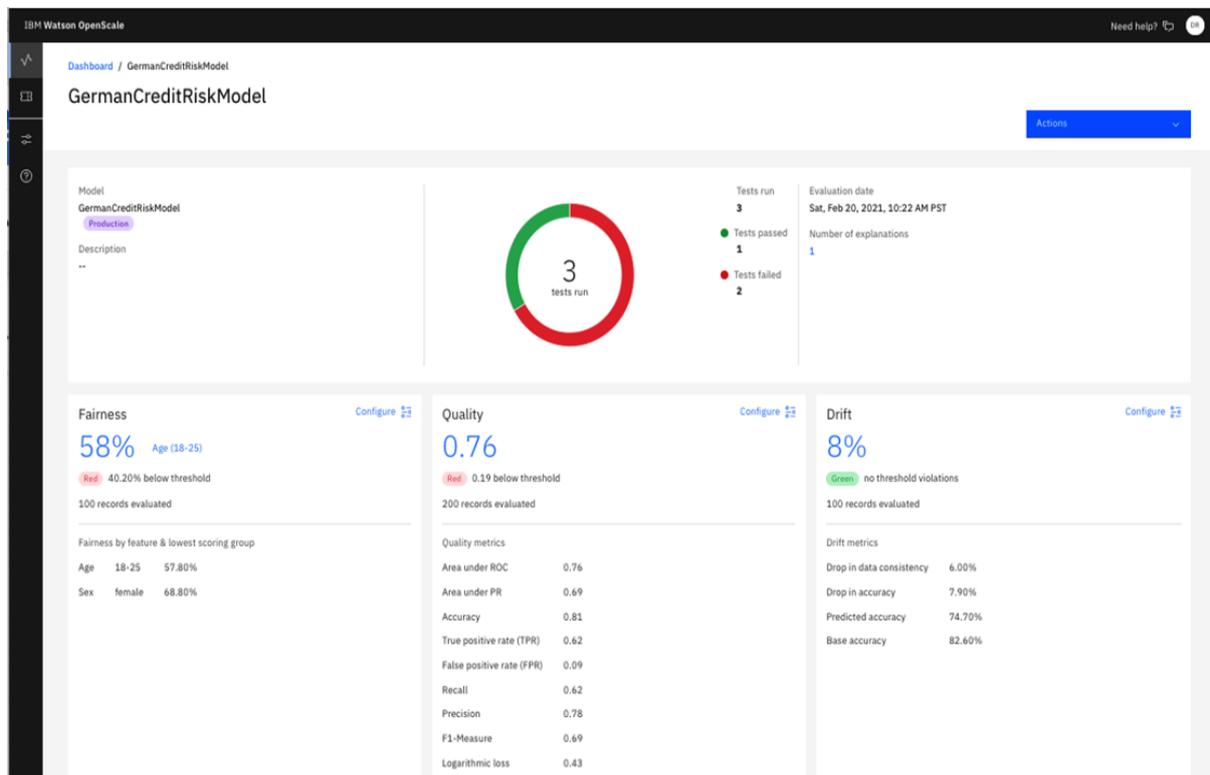
Enter search criteria here
Records ▼
Simple
Advanced

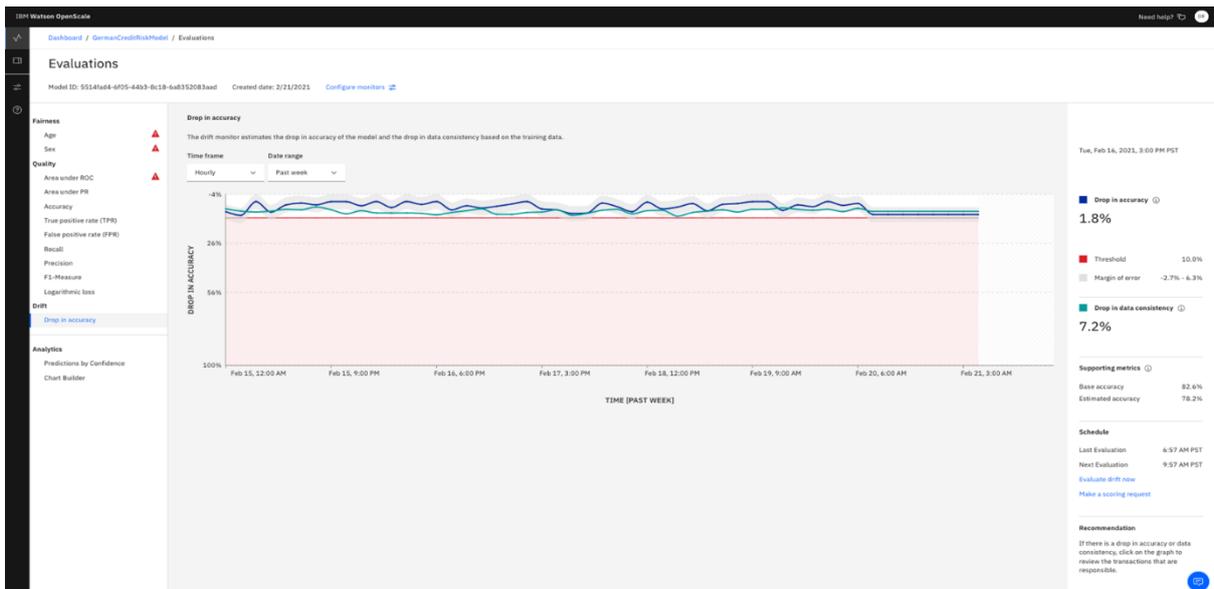
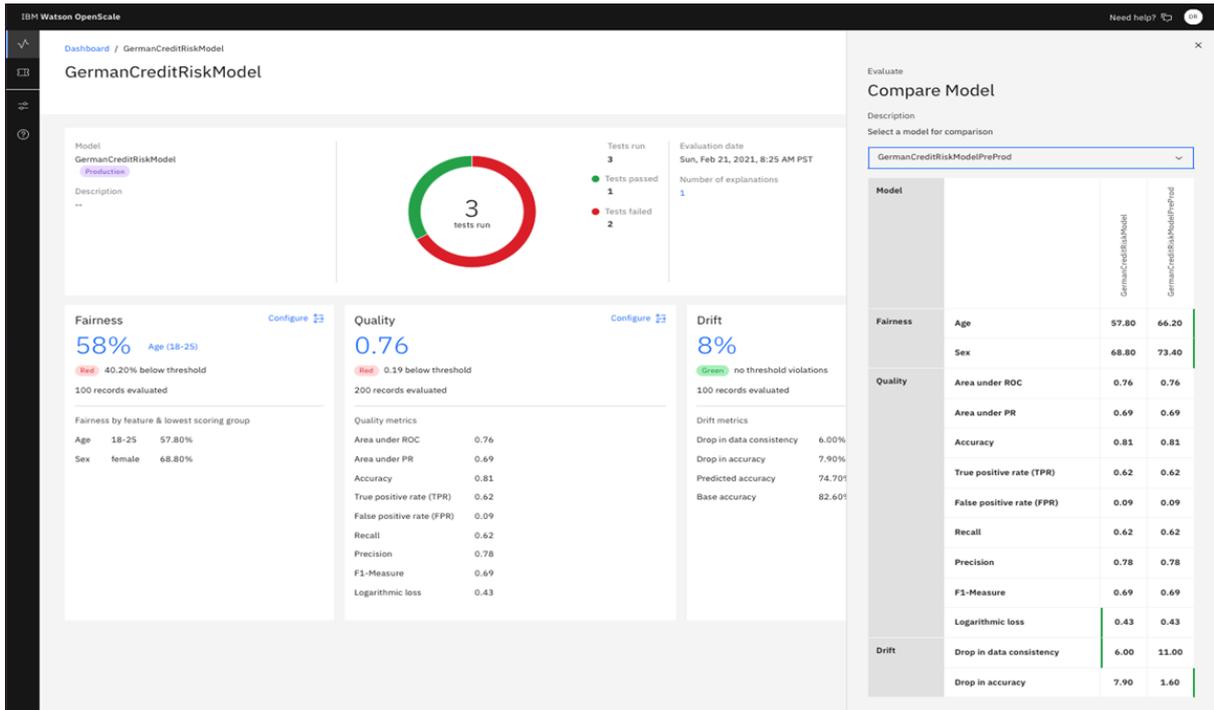
Brand (1)    **Organization (8)**    Person (1)    [Show graph](#)

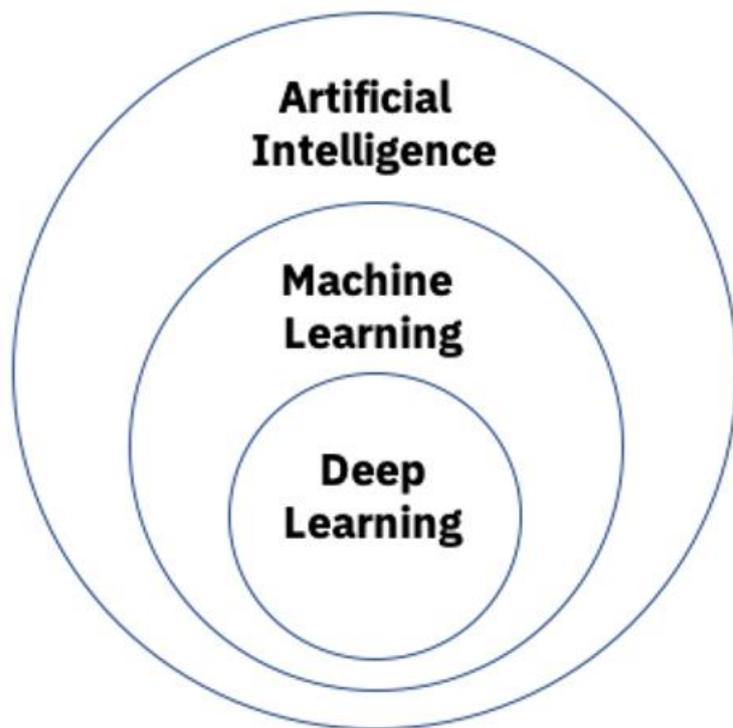
ID	Business address				Business name			Record display label	System record number
	Address line 1	City	Postal code	State/Province value	Name	Business tax identification	Business telephone		
0	7932 NORTHWEST JOHNSON ST.	OVERGAARD	85933	AZ	HEMODYNAMICS	59-8363961	9289188968	HEMODYNAMICS	
0	1240 NORTHWEST JOHNSON STREET	MELBOURNE	32935	FL	Ruegsegger Farms	52-2735232	3219242039	Ruegsegger Farms	
0	1240 NORTHWEST JOHNSON STREET	MELBOURNE	32935	FL	Ruegsegger Farms	52-2735232	3219242039	Ruegsegger Farms	
0	9823 E JOHNSON	WESTVIEW	40178	KY	DIAGNOSTIC LAB MEDICINE	14-8643014	2703714924	DIAGNOSTIC LAB MEDICINE	

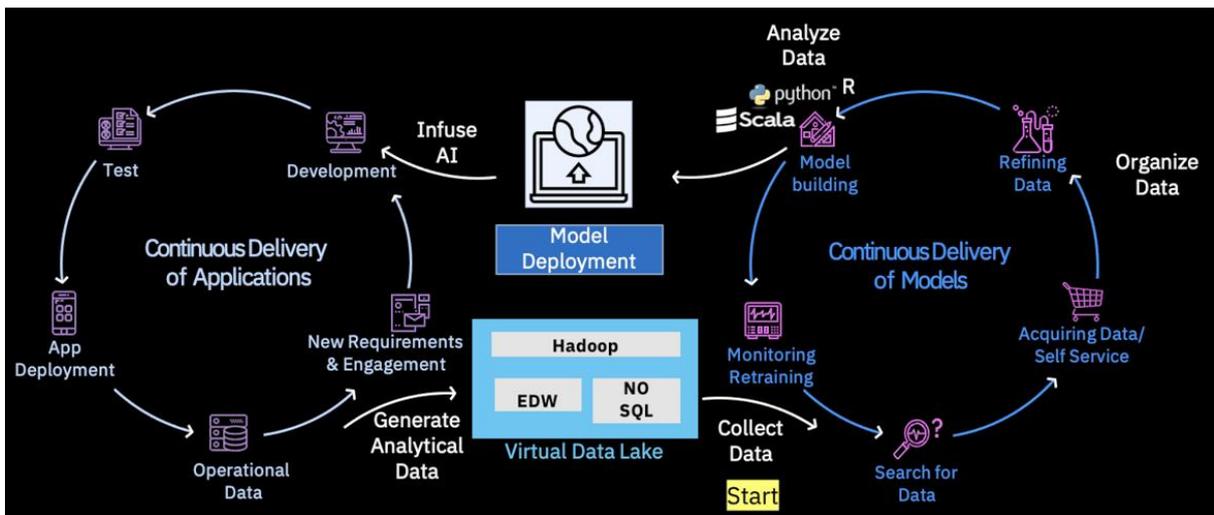
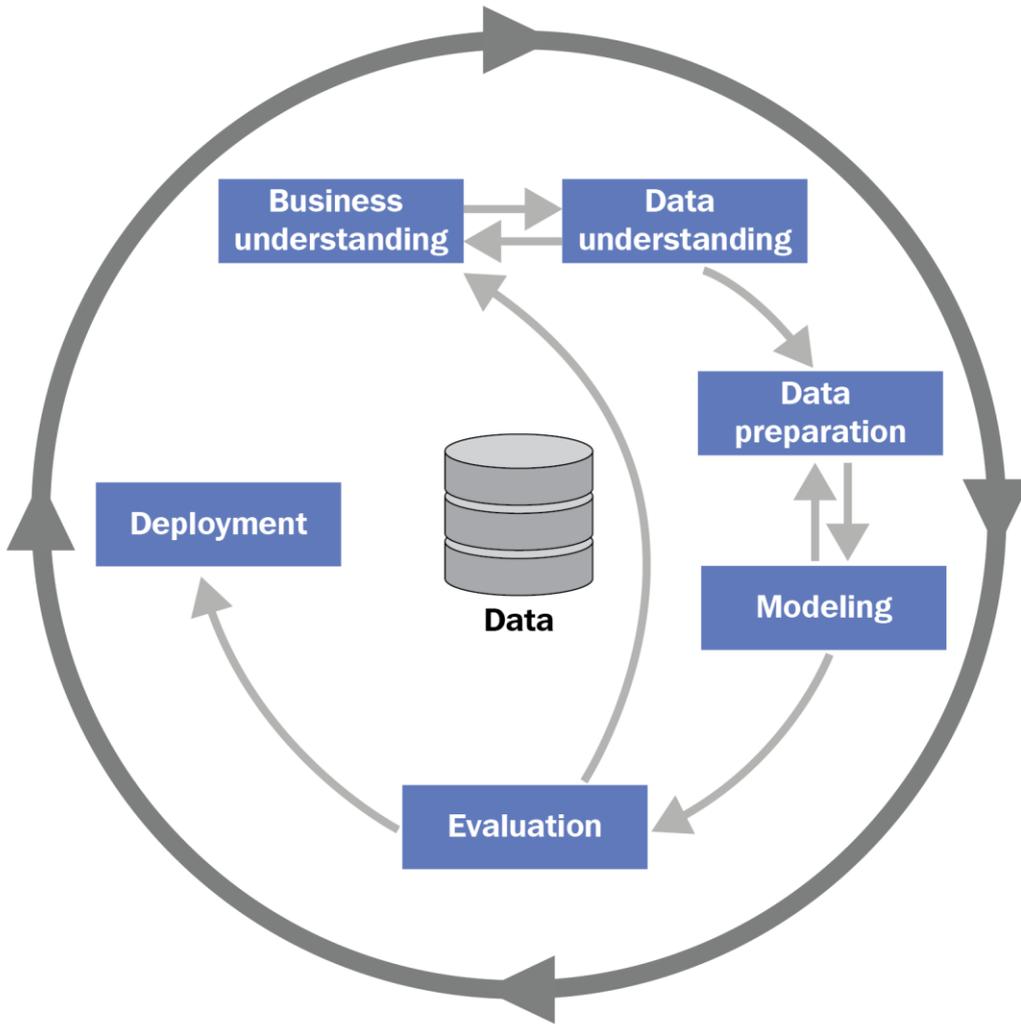
showing 1-10 of 186 record results    **1**    2    ...    19    >    Results per page: 10 ▼

# Chapter 5: Analyzing: Building, Deploying, and Scaling Models with Trust and Transparency









## Data Engineer



1. Run ETL jobs
2. Virtualize data
3. Add assets to the catalog

Configure connections to data sources typically used by the data science team

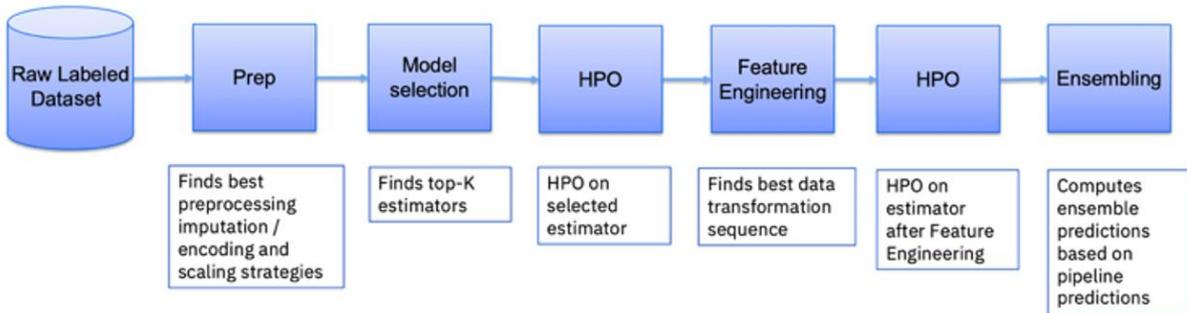
Security is provided by data sources  
Governance is provided by the catalog



1. Search for data in the catalog
2. If data is not found, submit a request

Use data sources, virtualized data, and assets from the catalog (data sources, tables, files, etc.)

## Data scientist



IBM Cloud Pak for Data

Projects / ModelOps\_Deployment

Overview **Assets** Environments Jobs Access Control Settings

Q What assets are you looking for?

▼ Data assets

0 assets selected.

Name	Type	Created by	Last modified
CSV <a href="#">customer-profile.csv</a>	Data Asset	DEEPAK RANGARAO	Feb 17, 2021, 7:29 AM
CSV <a href="#">churn.csv</a>	Data Asset	DEEPAK RANGARAO	Feb 17, 2021, 7:29 AM
CSV <a href="#">new_customers.csv</a>	Data Asset	DEEPAK RANGARAO	Feb 17, 2021, 7:29 AM

▼ Notebooks [New Notebook +](#)

Name	Shared	Scheduled	Status	Language	Last editor	Last modified
<a href="#">Predict_Customer_Churn</a>				Python 3.7	DEEPAK RANGARAO	Feb 17, 2021
<a href="#">Introduction to WML Client</a>				Python 3.7	DEEPAK RANGARAO	Feb 17, 2021
<a href="#">Customer_Data_Analysis</a>				Python 3.7	DEEPAK RANGARAO	Feb 17, 2021
<a href="#">Deployment_with_API</a>				Python 3.7	DEEPAK RANGARAO	Feb 17, 2021
<a href="#">Predict_Customer_Churn_SparkML</a>				Python 3.7 With Spark	DEEPAK RANGARAO	Feb 17, 2021

Data

Load Files Catalog

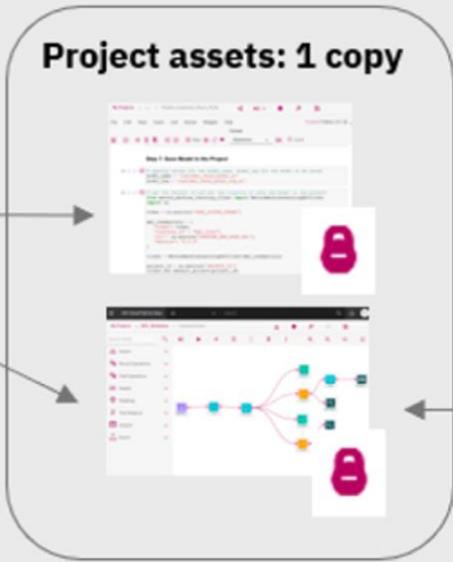
Drop files here or browse for files to upload.

### Collaborator 1



Edit

View only until unlocked by collaborator 2 or admin



### Collaborator 2

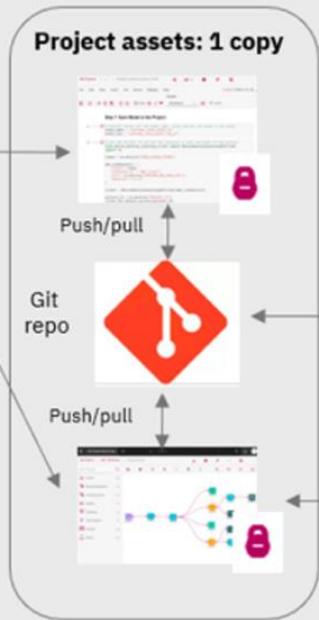


### Collaborator 1



Edit

View only until unlocked by collaborator 2 or admin



CI/CD

### CI/CD Tool



### Collaborator 2

IBM Cloud Pak for Data | All | Search | Upgrade | DEEPAK RANGARAO's Acc... | DR

Projects / ModelOps\_Deployment / customer-profile.csv / Refine data

Operation + Code an operation to cleanse and shape your data

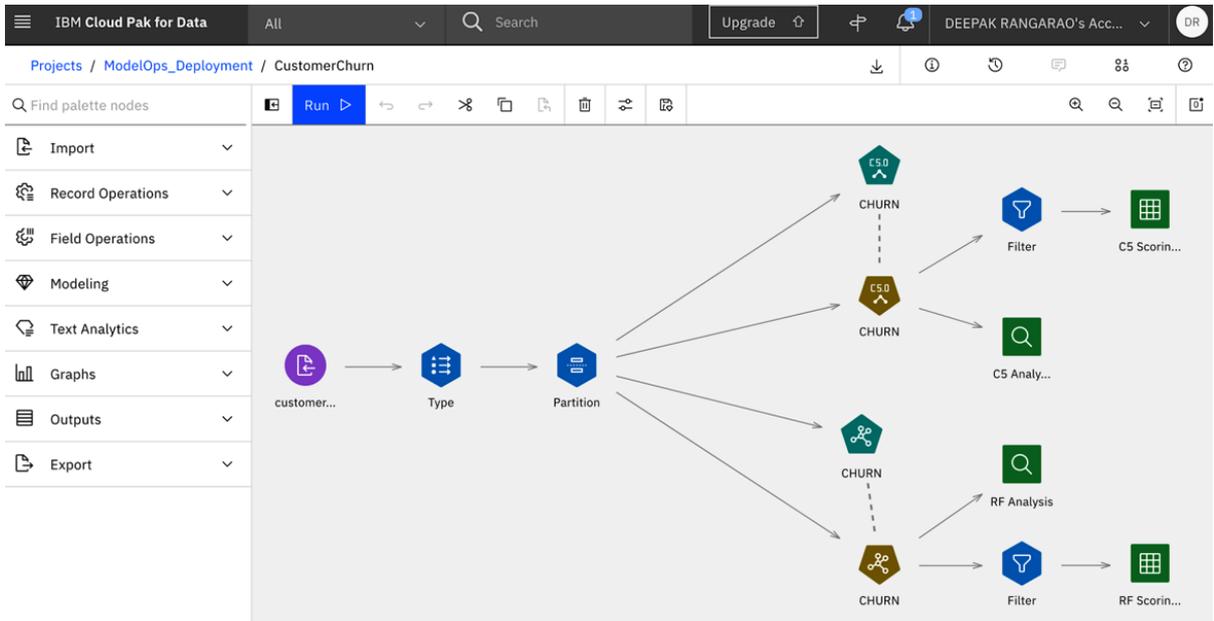
Data Profile Visualizations

ID	Gender	Status	Children	Est Income	Car Owner	Age	
Integer	String	String	Decimal	Decimal	String	Decimal	
1	843	M	M	1	120000	N	53.213333
2	1774	F	M	1	98680.4	Y	70
3	2146	F	M	1	98680.4	Y	42
4	2415	F	M	1	98680.4	Y	48
5	2659	F	M	1	98680.4	Y	43
6	2923	F	M	1	98680.4	Y	58.706667
7	3411	M	M	1	98680.4	Y	58.706667
8	3655	M	M	1	98680.4	Y	58.706667
9	1779	F	M	1	98453.6	Y	62
10	2664	F	M	1	98453.6	Y	52
11	2928	F	M	1	98453.6	Y	43
12	87	M	M	1	97180.4	N	50.6
13	87	M	M	1	97180.4	N	50.6
14	1623	M	M	1	97180.4	N	33
15	2661	F	M	1	95620.3	Y	46
16	3413	M	M	1	95620.3	Y	49.12
17	3657	M	M	1	95620.3	Y	49.12
18	3296	M	M	1	95448.6	Y	53.68

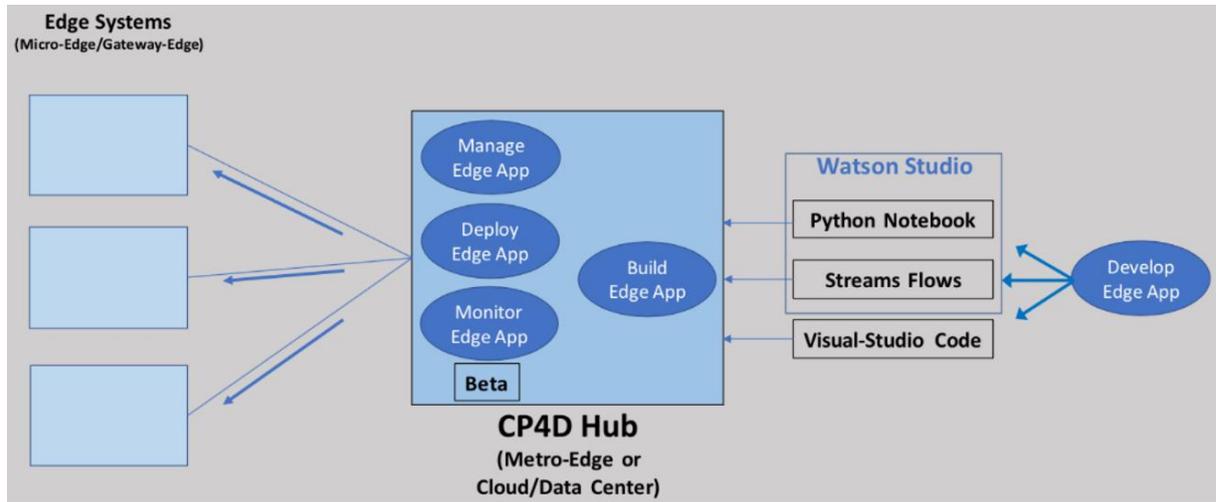
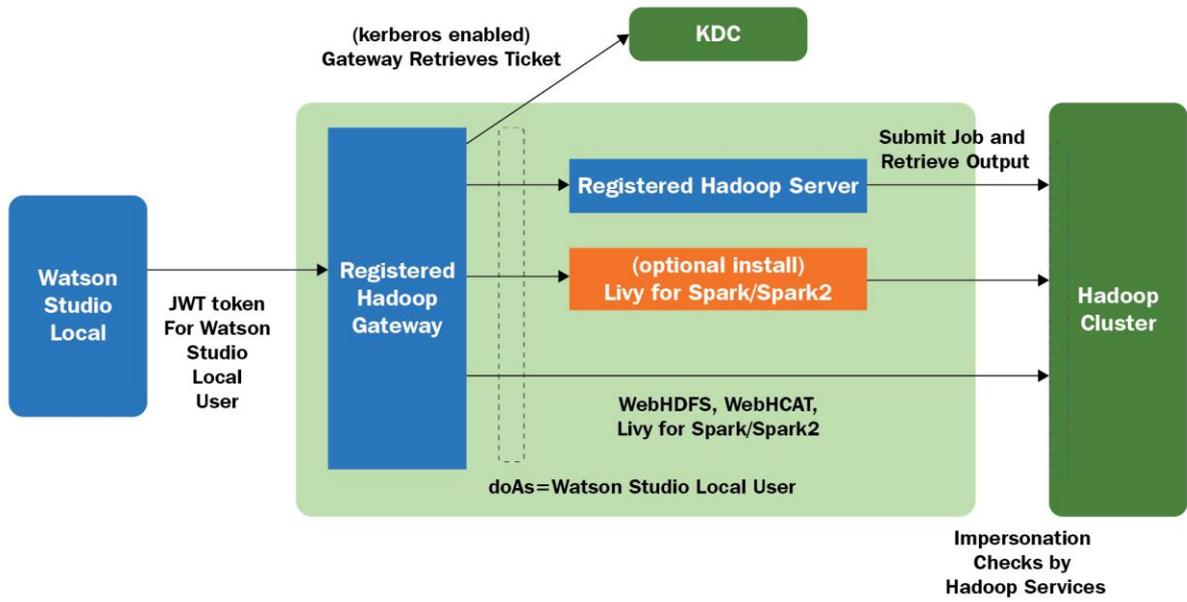
4 Steps

- Data Source: customer-profile.csv
- Convert column type: AUTOMATIC. Automatically converted one or more columns to inferred data types. Strings that are converted to decimal use a dot (.) for the decimal symbol.
- Sort descending: Sorted rows by Est Income
- Filter: Filtered by: Status where value contains M
- Filter: JUST ADDED. Filtered by: Children where value is equal to 1

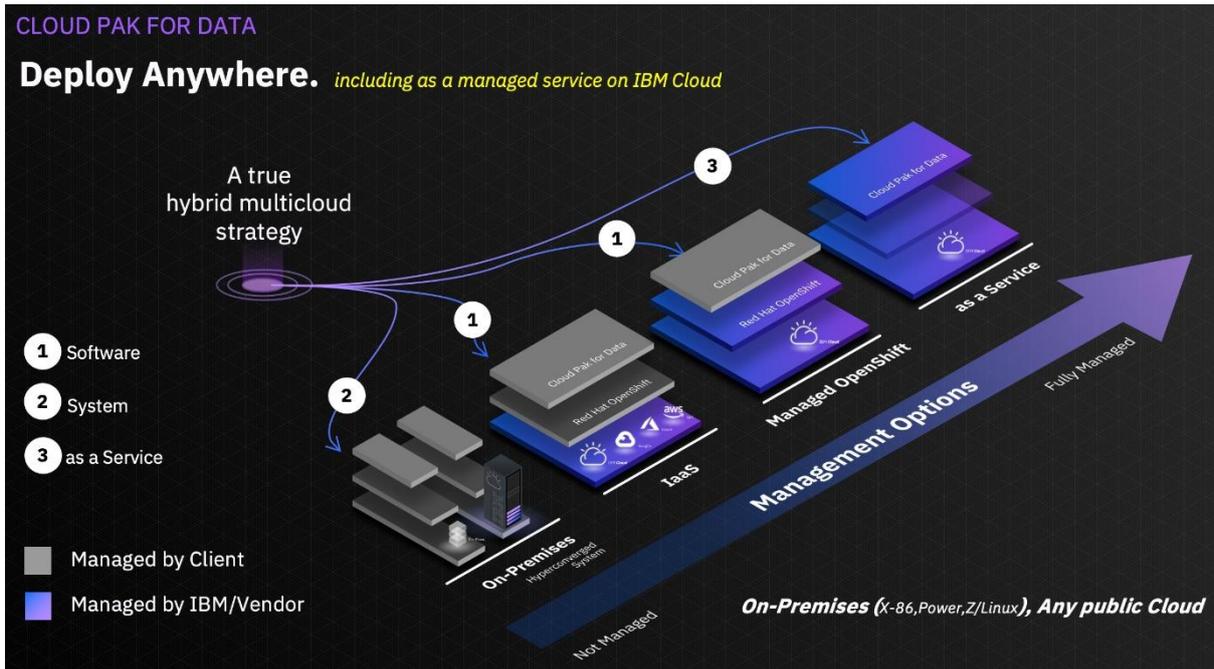
SOURCE FILE: customer-profile.csv | FULL DATA SET: 150 rows







# Chapter 6: Multi-Cloud Strategy and Cloud Satellite



## Cloud Pak for Data on AWS

Integration into AWS services

Elasticity

Availability

Agility

Global Footprint

Security

Native Services For Elasticity

Native Services For Availability

### Features:

- Automation for Infrastructure Provisioned Installer (IPI) to set up OpenShift and storage under 90 mins
- Supports parameterization and customization of the templates for deployments

### Support:

- ✓ Supported on AWS and Azure
- ✓ OpenShift 4.3
- ✓ Storage
  - AWS: Portworx and OpenShift Container Storage
  - Azure: NFS and Portworx



## Azure 60-day trial

# OpenShift Container Platform 60-day trial

## Quick Start and Marketplace

*Using Azure Resource Manager Templates*



Products > IBM Cloud Pak for Data



## IBM Cloud Pak for Data save for later

Tenant-IBMAIalliance-MPN-CORE-1560886

★★★★★ (0) [Write a review](#)

[Overview](#) [Plans](#) [Reviews](#)

Accelerate your journey to AI and transform your business with an extensible data and AI platform

IBM Cloud Pak for Data is an analytics platform that helps you prepare your data for AI. It enables data engineers, data stewards, data scientists, and business analysts to collaborate using an integrated multiple-cloud platform.

Cloud Pak for Data uses IBM's deep analytics portfolio to help organizations meet data and analytics challenges. The required building blocks (collect, organize, analyze, infuse) for information architecture are available using Cloud Pak for Data on Azure.

**Key Features:**

- Connect ALL data and eliminate data silos for self service analytics
- Automate and govern a unified data & AI lifecycle
- Analyze your data in smarter ways, with ML and Deep Learning capabilities
- Operationalize AI with trust and transparency

**Pricing information**  
Cost of deployed template components

**Categories**  
Analytics  
Databases  
IT & Management Tools

**Support**  
Support  
Help

**Legal**  
License Agreement  
Privacy Policy

IBM Watson OpenScale Need help? AA

OpenScaleSample

Dashboard / GermanCreditRiskModelICP

## GermanCreditRiskModelICP

Actions

Model: GermanCreditRiskModelICP (Production)

Tests run: 3  
 Tests passed: 1  
 Tests failed: 2

Evaluation date: Tue, Oct 13, 2020, 10:16 AM JST  
 Number of explanations: 0



### Fairness

80% Sex (female)

Red 18.00% below threshold

100 records evaluated

Fairness by feature & lowest scoring group		
Age	18-25	80.00%
Sex	female	80.00%

### Quality

0.62

Red 0.18 below threshold

100 records evaluated

Quality metrics	
Area under ROC	0.76
Area under PR	0.69
Accuracy	0.81

### Drift

0%

Red above threshold

100 records evaluated

Drift metrics	
Drop in data consistency	7.00%
Drop in accuracy	0.00%
Predicted accuracy	85.70%

IBM Cloud Pak for Data Upgrade IBM

My Projects / COVID-19 Tracking with IBM Dat... / COVID Tracking - Colleges in Pro...

Find palette nodes

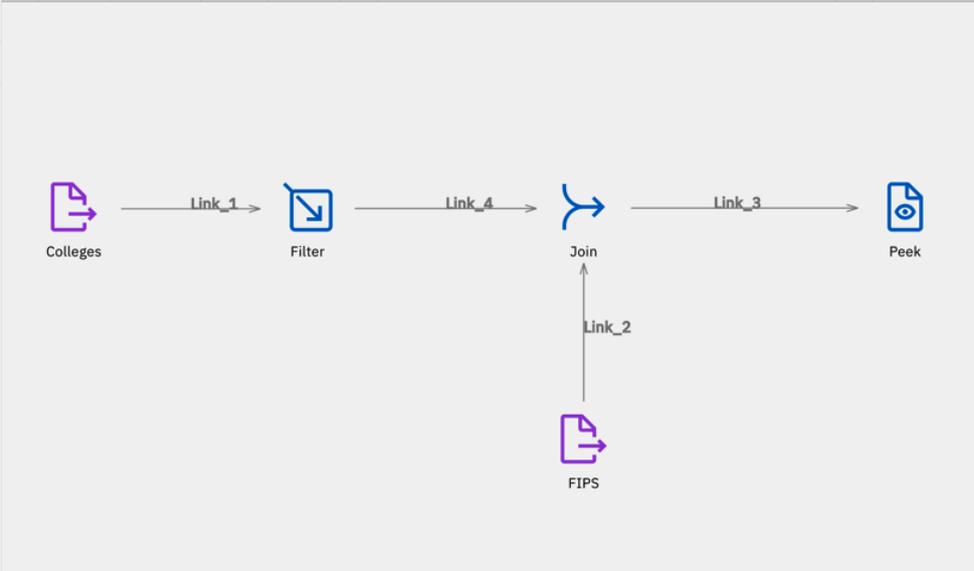
Save Compile Run

Loaded 1:01 PM

Connectors

Stages

- Aggregator
- Checksum
- Copy
- Filter
- Funnel
- Head
- Join
- Lookup
- Merge
- Modify
- Peek
- Remove Duplicates
- Row Generator
- Sort
- Tail



```

graph LR
    Colleges[Colleges] -- Link_1 --> Filter[Filter]
    Filter -- Link_4 --> Join[Join]
    FIPS[FIPS] -- Link_2 --> Join
    Join -- Link_3 --> Peek[Peek]
    
```

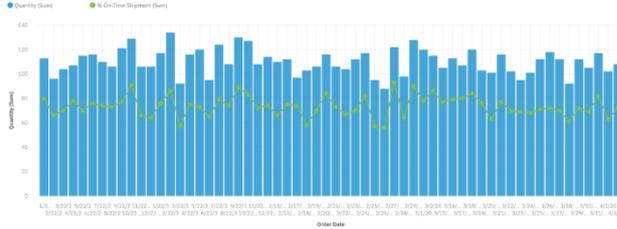
Web Orders

Product Mix

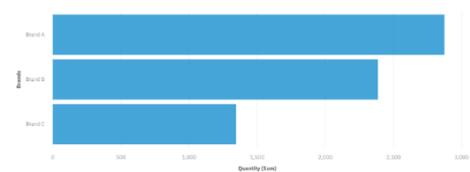


67%  
% On-Time Shipment

Daily Volume vs. On-Time Shipments



Orders by Brand



Resource list /

Analytics Engine-sv Active [Add tags](#)

[Details](#) [Actions...](#)

Getting started

Manage

[Launch Console](#)

Service credentials

Connections

Plan

Details

Compute nodes 1

Status  
**Active**

Software package  
AE 1.2 Spark and Hive

Components

- Apache Spark 2.3.2
- Apache Livy 0.4.1
- Ambari 2.7.3
- Hive 3.1.0
- Hadoop 3.1.1
- Knox 1.0.0
- Anaconda-Py 3.7.1
- Jupyter Enterprise Gateway 0.8.0

Price estimate

USD **1.40/hour**

United States

Credentials

User ID  
clsadmin

All clusters must be password-protected. Set or reset the password for your cluster. See [Credentials in IBM Analytics Engine](#).

[Generate password](#)

Nodes

Node type	Node name	Commission Time
master-management	██████████.us-south.ae.appdomain.cloud	2020-06-10T09:17:04Z

FEEDBACK

Check my balance

[Send us feedback](#)

Auto-save ⓘ

[Close](#)

Saved ⌚

Action started by

- Customer messages

1 <span data-send-event="Balance">I can help you with that. One moment while I retrieve a list of accounts...

Checking Savings +1

Continue to next step

2 1 = Savings Plus

Great! It looks like you have \$3,960.00 in Savings Plus. Would you like to see your Checking Account...

Yes No

Continue to next step

3 1 = Savings

Great! It looks like you have \$2,370.80 in Savings. Would you like to see your Checking Account balance?

Yes No

Continue to next step

4 3 = Yes 2 = Yes

Okay. You have \$3173.30 in your Checking Account.

Continue to next step

5 1 = Checking

Great! It looks like you have \$3173.30 in your Checking Account. Would you like to see your Saving...

Yes No

Continue to next step

New step +

Step 2 is taken with conditions

Conditions

1 condition

1. All of this is true:

If 1. <span data-send-event="Balance">I c... is Savings Plus

and Add condition +

2. New group +

Assistant says

B I @

Great! It looks like you have \$3,960.00 in Savings Plus.

Would you like to see your Checking Account balance?

Yes No

[Edit response](#) [Edit validation](#)

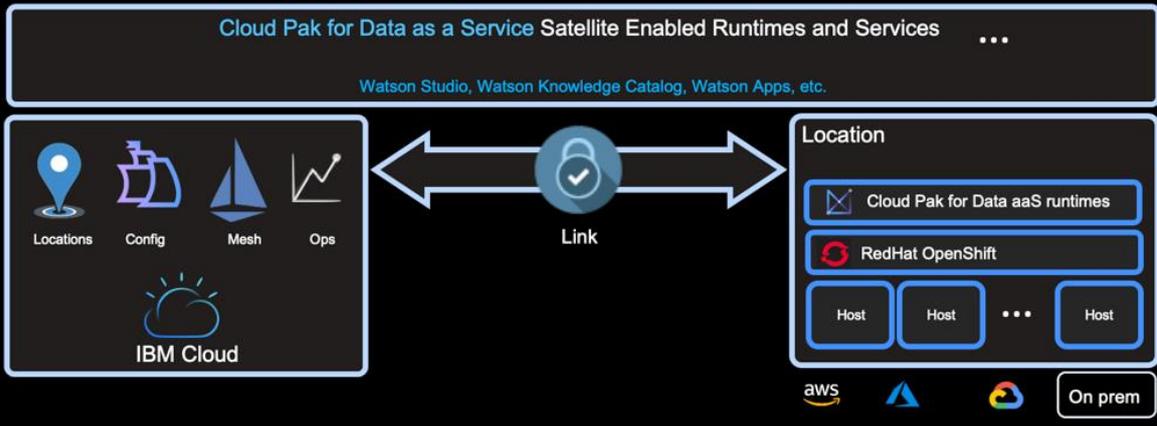
And then

Continue to next step

Preview ▶

# Run managed instances of **Data and AI workloads** on Satellite Locations

All workloads in Locations run on Red Hat OpenShift leveraging IBM's enterprise cloud managed OpenShift service



## Data Fabric : Manage your **Data and AI** across hybrid and multi-cloud stacks

- 1 **Data Orchestration & Virtualization**
  - Provision & manage databases across multiple clouds
  - Access data irrespective of where it resides (Data Virtualization)
  - Clone databases, create sample dev-test instances with masking across clouds
- 2 **Data Governance (Distributed & Synchronized)**
  - Auto-catalog all enterprise data & analytics assets, define policies & rules
  - Distributed data governance & meta data synchronization through OMRS
  - Location sensitive data governance : track data lineage, define & enforce policies etc.
- 3 **Model Management & Deployment**
  - Develop, train , deploy and manage ML & AI models for multi-cloud
  - Ensure trust and governance of models across enterprise
- 4 **Data Movement (Backup/DR/Operational)**
  - Ad-hoc data movement/copies across different clouds
  - Backup/disaster recovery of data: lift and shift
- 5 **Data Privacy & Security**
  - Policy enforcement and data masking to ensure data privacy across Cloud
  - Data access and security across clouds



IBM Cloud Pak for Data

My data: Databases / Informix-1 / Details

## Details: Informix-1 ✓

[Open](#)

About this database		Storage	
Database name	db1	Storage class (Storage)	rook-ceph-cephfs-internal
Database type	informix	Size (Storage)	1 GB
Database software version	14.10.0-x86_64	<b>Access information</b>	
Processor	x86-64	Username	user999
Deployment id	informix-1594806857820	Password	██████████
Created on	15 Jul 2020 10:54	JDBC Connection URL	jdbc:db2://10.16.7.155:30037/db1
Status	Available		

**Nodes**

HOSTNAME	CPU	MEMORY

DASHBOARD Filter by Appliances Organizations America/New\_York admin Last Updated: 10:10:19

### Applications\*

Managed(6) Unmanaged(81) Unscheduled(0)

**81**  
Unmanaged

Click on labels to view details

### Job Status

Current Jobs: Running(1) Queued(0)

Past: Last 24 Hours

Succeeded(108) Failed(0) Retried(0) Not Run(0)

**Humming along nicely**

Click on labels to view details

### Events (Last 24 Hours)\*

2 Errors 0 Warnings

### Managed Data

**1,1 TB**

2020-02-07 00:00:00

### Appliances Health

Disk Usage

Snapshot Pool	1	0	0
Primary Pool	1	0	0
Dedup Pool	1	0	0

Hardware

VDisk Utilization	1	0
-------------------	---	---

Software

Local Snapshot	1	0
Local Dedup	1	0
Remote Manage	0	1

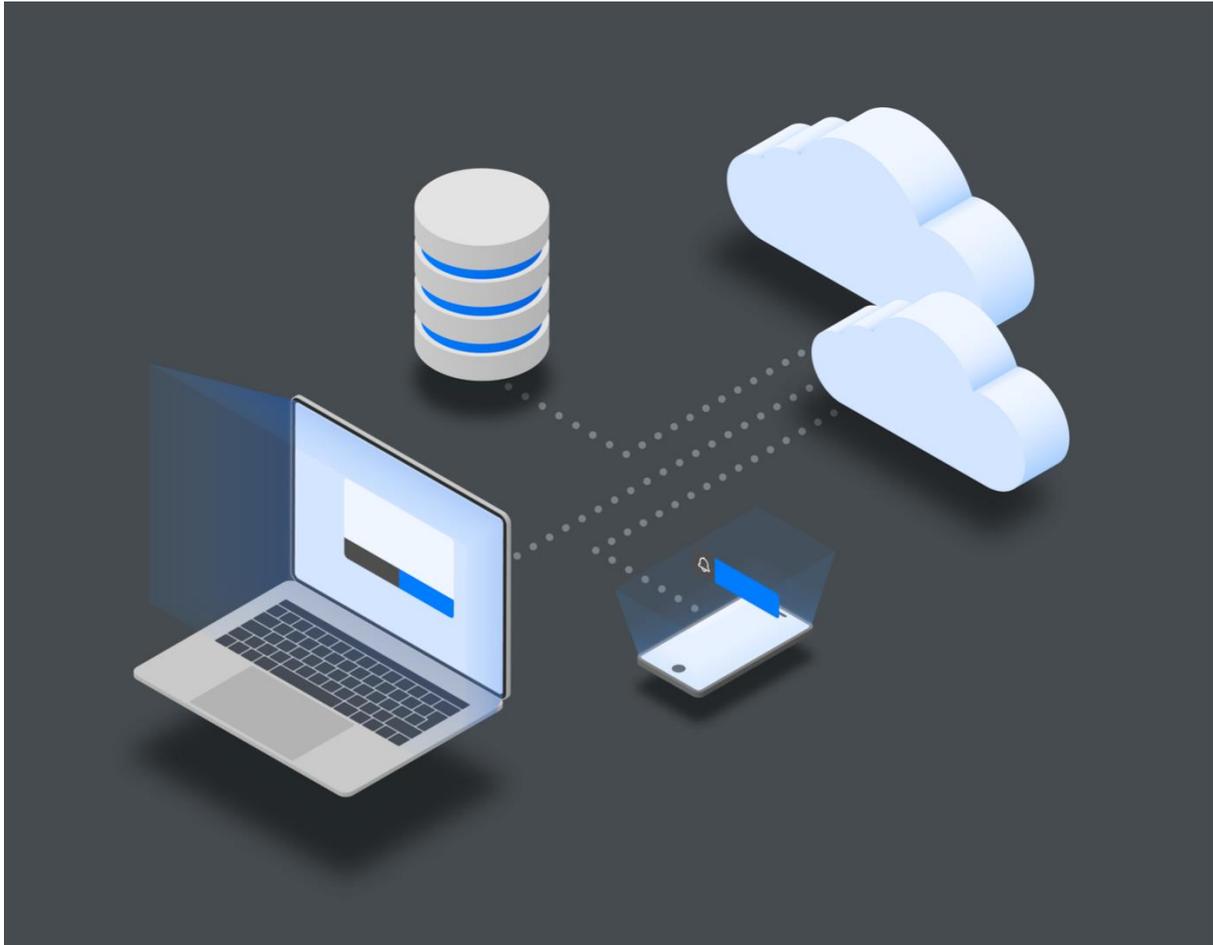
IVGM Connections

Appliances	1	0
------------	---	---

Click on numbers to view details

### Dedup

<b>30.6 TB</b> Logical Copy Data	<b>535.2 GB</b> Dedup Used	<b>58.6 : 1</b> Dedup Efficiency
-------------------------------------	-------------------------------	-------------------------------------



mongoDB Ops Manager 4.0.17 OK All Clusters Please set your time zone Admin firsttest -

CONTEXT **Mongo** Add New ...

PROJECT **ACTIVE**

**Deployment**

Processes Servers Agents Security More ▾

VIEW: TOPOLOGY LIST Search...

**mongodb-1586226366247-replica-set**  
Version 4.0.6-ent

DATA METRICS MODIFY ...

<b>TYPE</b>	<b>SSL</b>
Replica Set	Disabled
<b>MONGODB COUNT</b>	<b>AUTH</b>
3 Mongods	Disabled
<b>DATABASE SIZE</b>	
277 B	
<b>BI CONNECTOR CT</b>	
0 Connectors	

Operations  
R: 2.1 W: 0.02

Last 6 Hours

Connections  
22

Last 6 Hours

Last Login: 10:254.0.1 Version: 4.0.1750606.20200204T0832Z  
 Since: 2020-04-07T02:37:04Z Central URL: http://mongodb-1586226366247-ops-manager-svc:8080  
 ©2020 MongoDB, Inc.

IBM Cloud Pak for Data

Project: dstage1    Connections    Table definitions    Parameter sets    Jobs    Job\_1\*

PARALLEL JOB Job\_1\*

Subrecord    Vector

Peek    Pivot    PS    Remove    Row    Slowly Changing Dimension  
Enterprise    Subrecord    Duplicates (33%)    Generator

Sample    Split    Split    Surrogate    Switch    Tail  
Subrecord    Subrecord    Vector    Key    Generator

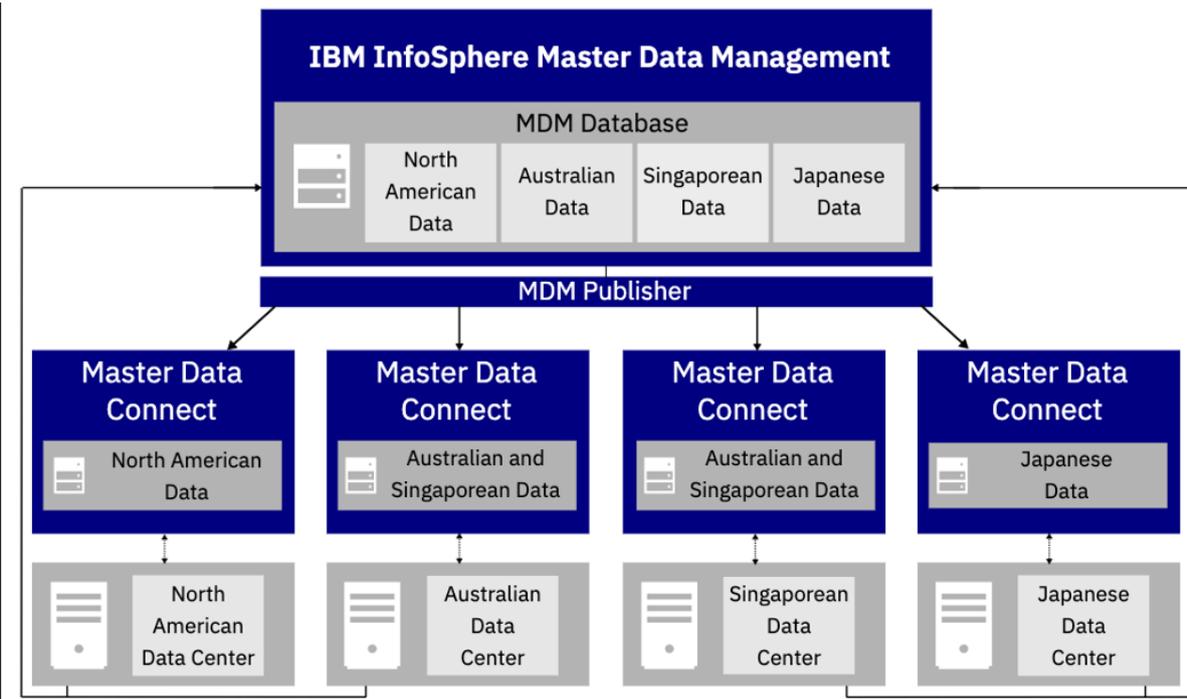
Wave    Write  
Generator    Range Map

GENERAL

Container    Shared Container

SMART STAGE SUGGESTIONS

Filter (33%)    Sort (33%)    Remove Duplicates (33%)





**Business Users**  
Consume and customize apps through digital web experiences

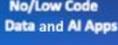


**Data Scientists**  
Create models based on trusted data and business PoVs



**App Developers**  
Rapidly build AI-infused business applications





No/Low Code  
Data and AI Apps



Data-Driven Decisions



Business Ontologies

## IBM Cloud Pak for Data



Collect Data



Organize Data



Analyze Data



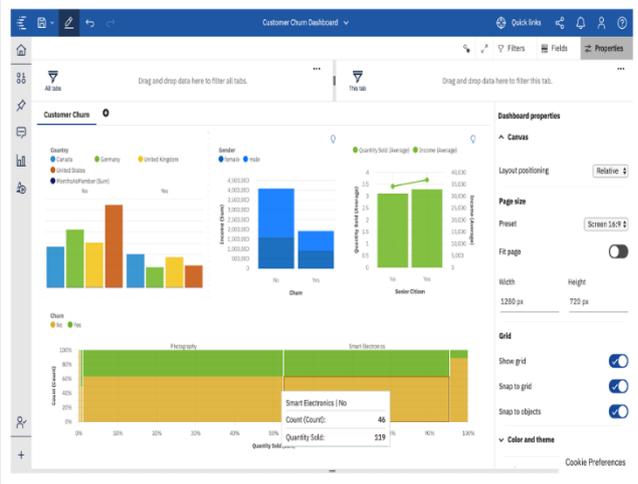
Infuse AI



Hybrid Cloud Platform



All Your Data, On Any Cloud



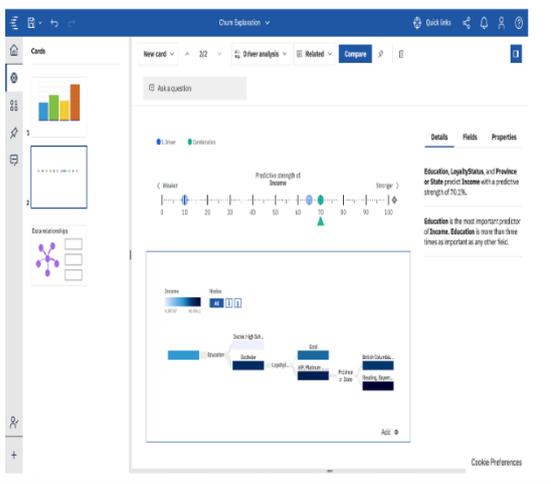
**Customer Churn Dashboard**

Customer Churn

Quantity Sold (Average)

Smart Electronics (No)

Quantity Sold: 119



**Churn Exploration**

Predictive strength of feature

Education, LoyaltyStatus, and Province or State predict Churn with a predictive strength of 76.2%.

Education is the most important predictor of Churn. Education is more than three times as important as any other factor.

Main
Corp Commentary

## Operating Expenses - Budget

➔

Months  
Total Year

**Select Organization:**

- Total Company
- Digital
- Website 01
- Website 02
- Kiosk
- East
- Boston, MA
- NY, NY
- Atlanta, GA
- West
- Seattle, WA
- San Francisco, CA
- Portland, OR
- Central
- Chicago, IL
- Denver, CO
- Dallas, TX

**Income Statement Variance Analysis:**

Account	Amount				Comments
	2018	Budget	Var\$	Var%	
4999 Revenue	34,124	76,026	41,902	122.8%	Tech release during year
5999 Cost of Sales	20,209	57,751	(37,542)	(185.8%)	majority related to increased sales
<b>Gross Margin</b>	<b>13,915</b>	<b>18,275</b>	<b>4,361</b>	<b>31.3%</b>	
Payroll	2,719	5,795	(3,076)	(113.1%)	Larger store needing more HC
Office Expenses	1,122	1,872	(750)	(66.8%)	Ins Liab Coverage increase and New Tele System
Travel Expenses	1,049	2,529	(1,480)	(141.0%)	Large event planning session at Corp
Occupancy	1,881	4,164	(2,283)	(121.3%)	More SQ FT to service and maint in Q3
Marketing	2,166	6,610	(4,444)	(205.2%)	Social Media/Innovation Campaign, Seasonal Events
Depreciation	1,401	2,212	(811)	(57.9%)	
<b>Total OPEX</b>	<b>10,339</b>	<b>23,182</b>	<b>(12,844)</b>	<b>(124.2%)</b>	
<b>Net Profit</b>	<b>3,576</b>	<b>(4,907)</b>	<b>(8,483)</b>	<b>(237.2%)</b>	

Select Account: Accounts Office Expenses

**\$ Variance and % Variance by Month:**

**Select Item Below:**

Office Expense

Occupancy

**Select Marketing Items:**

Sales Promotions

Advertising

**Ad Hoc Analysis:**

Ad Hoc

**Excel Integration:**

Excel Designed Report

Created during Demo

IBM Watson Assistant Plus
Learning center

Check my balance

[Send us feedback](#)
✔ Auto-save
Close
Saved

**Action started by**

- Customer messages

---

1. Checking Savings + 1

↓ Continue to next step

---

1. Savings Plus

Great! It looks like you have \$3,960.00 in Savings Plus. Would you like to see your Checking Account balance?

2. Yes No

↓ Continue to next step

---

1. Savings

Great! It looks like you have \$2,370.80 in Savings. Would you like to see your Checking Account balance?

3. Yes No

↓ Continue to next step

---

3. Yes    2. Yes

4. Okay. You have \$3173.30 in your Checking Account.

↓ Continue to next step

---

1. Checking

Great! It looks like you have \$3173.30 in your Checking Account. Would you like to see your Saving...

5. Yes No

↓ Continue to next step

New step +

**Step 2 is taken** with conditions

Conditions: 1 condition

- All of this is true:
  - If 1. <span data-send-event="Balance">I c... is Savings Plus

and Add condition +

---

2. New group +

---

**Assistant says**

Great! It looks like you have \$3,960.00 in Savings Plus.

Would you like to see your Checking Account balance?

Yes No

Edit response Edit validation

---

**And then**

↓ Continue to next step

Preview ▶

IBM Watson Discovery

Cookie Preferences Instance: Discovery-yu

Agent Assist [Configure data](#)

Overview Errors and warnings (1) Search settings

1 document

0 documents failed [View details](#)

Created on 5/28/2019 12:15:51 pm EDT  
Last updated 5/28/2019 12:15:51 pm EDT

[Upload documents](#)

Identified 1 field from your data

text

Need to identify more fields? [Add fields](#)

Added 4 enrichments to your data

Entity Extraction  
Administrator (1) | Alfa Romeo (1) | Aston Martin (1) | Audi (1) | BMW (1)

Sentiment Analysis  
100% positive 0% neutral 0% negative

Concept Tagging  
Car rental (1) | Damage waiver (1) | Insurance (1) | Vehicle insurance (1)

Category Classification  
finance → financial news

5 enrichments available. [Add enrichments](#)

Now you're ready to query!

Entities of type Company which have negative sentiment [Run](#)

Top people related to /finance/financial news [Run](#)

Entities of type Company which have positive sentiment [Run](#)

[Build your own query](#)

IBM Watson Knowledge Studio

reggie.raye@ibm.com

Back to Workspaces Back to Task Open document list

Assets Documents Entity Types Relation Types Dictionaries Rule-based Model Rules Versions Machine Learning Model Pre-annotation Annotation Tasks Performance Versions Settings Help

View Details Replace Concordance Attribute View View Guidelines In Progress Save

race.txt

1 Freddy Fraye, captain of the Harvard University squad, hit the tarmac with 10K to go in the Boston Beanpot Road Race.

2 Though his Cannondale's derailleur was damaged, he managed to regain the wheel of the Westpoint riders who'd been his breakaway companions.

3 Penn State, which had come with a stacked team, took up the chase in the peloton.

4 The gap to the break fell to 3 seconds with 100m to go, but the Penn boys ran out of road, allowing Freddy to cross the line first.

Entity		Mention
Type	Subtype	Role
-	bicycle	
-	crash	
-	distance	
-	race	
-	rider	
-	team	
-	win	



**IBM Financial Crimes Insight with Watson** Admin In-Basket

#5 - Possible Day Corp. Insights Profile Transactions Six-Month Variance Narrative

**Overall Risk: Exonerating**  
14% Watson Insight Score

**Transactions**

- Geographic Location
- Observed Behavior Pattern

**Customer**

- Prior Alerts
- Anomaly
- Peer Group Risk**
- Complexity Risk

**Risk Key**

- Highly Aggravating
- Exonerating

**Geographic Location**  
In the review period prior to the alert, the customer conducted 1 transactions with high risk countries, including (China).

**Customer**

- Prior Alerts**  
This party has no history of prior alerts.
- Anomaly**  
This party does not have unusual activity relative to its peer group number {cluster}.

**Transactions**

- Observed Behavior Pattern**  
This customer's behavior is similar to the behavior of other customers whose activity was noted as not suspicious. The following shows the alerted customer's transaction patterns as compared to characteristic actions of similar customers.

Indicator	Customer's value	Range of similar customer values
Maximum total quarterly sum of transactions	155.05	value <= 67,389.00
Variance of the total quarterly sum of transactions	0.00	value >= 0.00
Minimum number of days between transactions	0.00	value >= 0.00

Set status

**IBM Surveillance Insight** In-Basket

75% - Party Behavior - 13 Overview Network Involved Employees

**Reasoning Graph**

Start date: Jun 30, 2019 End date: Jun 30, 2019

00:01 Jun 30, 2019 00:01

Legend: NO SCORE (grey), LOW (yellow), MEDIUM (orange), HIGH (red)

**Evidence**

Weight	Type	Date	Description
85%	E-mail	Jun 30, 2019 09:10:00	Mass mails sent by Kimberly Johnst
85%	E-mail	Jun 30, 2019 09:10:00	Kimberly Johnst communicated in unusual timings
85%	E-mail	Jun 30, 2019 05:11:43	John Collins communicated in unusual timings
85%	E-mail	Jun 30, 2019 05:11:43	John Collins included unusual number of attachments
73%	Sentiment	Jun 30, 2019 05:11:43	John Collins was excessively negative

Reasoning Graph nodes: Comm Vol Anomaly (LOW), Behavior Anomaly (MEDIUM), Comm Anomaly (HIGH), Sensitive Content (LOW), Risk (MEDIUM)

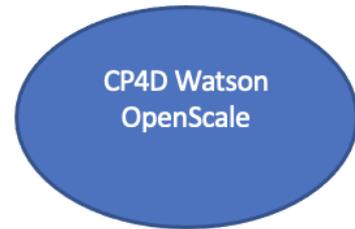
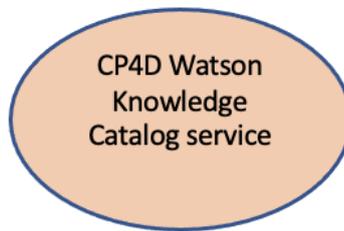
Set status

# Chapter 8: Customer Use Cases

Collect

Organize

Analyze



Data steward



Data scientist

Collect

Organize

Analyze

Infuse

CP4D Watson Knowledge Catalog service On-premise

CP4D Watson Studio Local service

CP4D Watson Discovery service

CP4D Watson Assistant service

CP4D Watson API kit



Business analyst

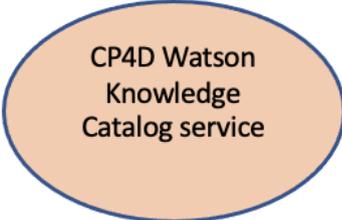
Data scientist

App developer

Collect

Organize

Analyze



Data steward

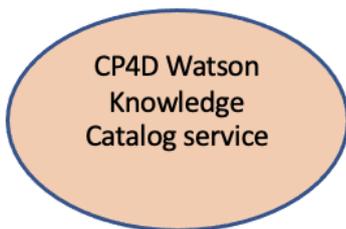


Data scientist

Collect

Organize

Analyze



Data steward



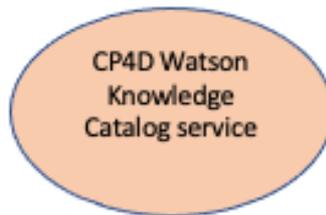
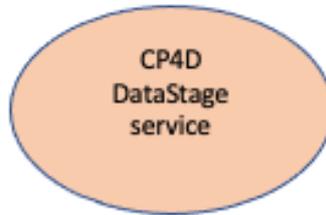
Data scientist

## Collect



Data engineer

## Organize



Data steward



Data scientist

## Analyze

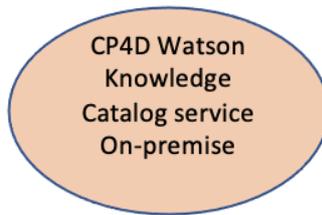


Business analyst

## Collect



## Organize



## Analyze



Data engineer



Business analyst

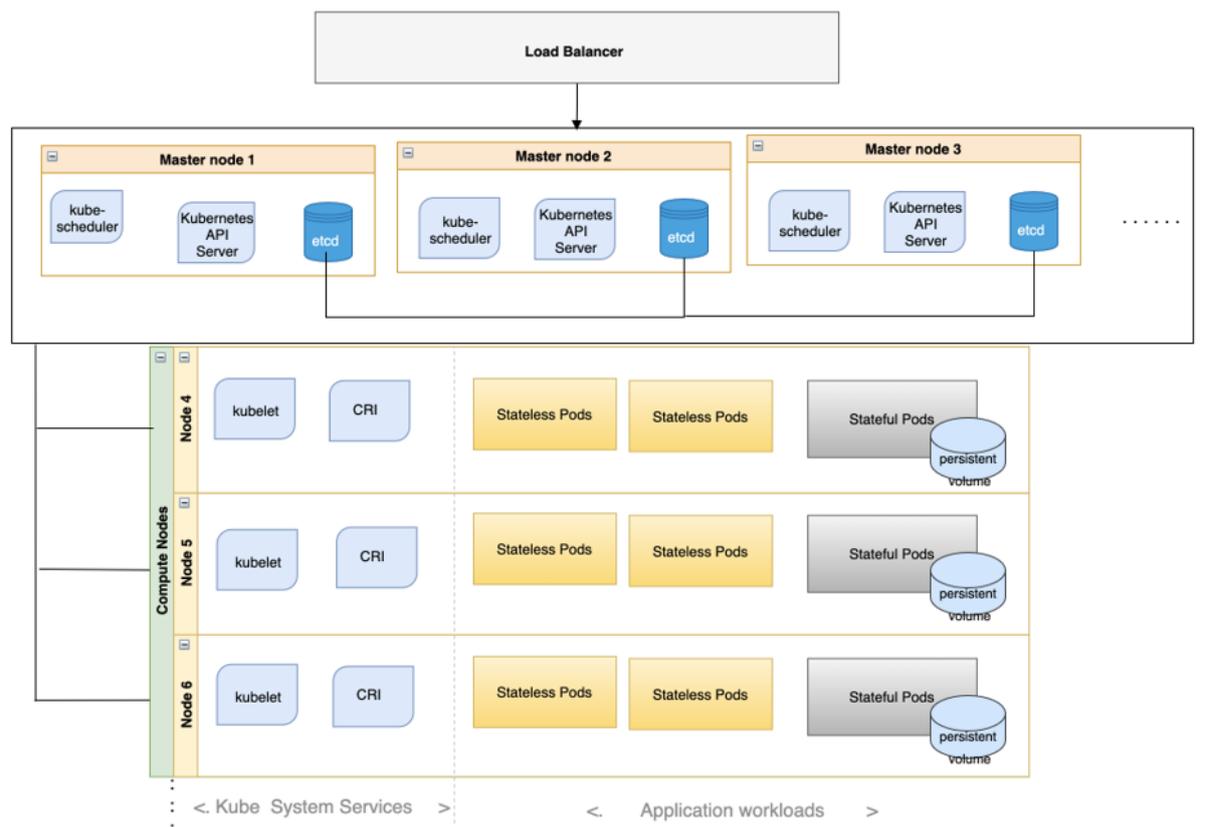
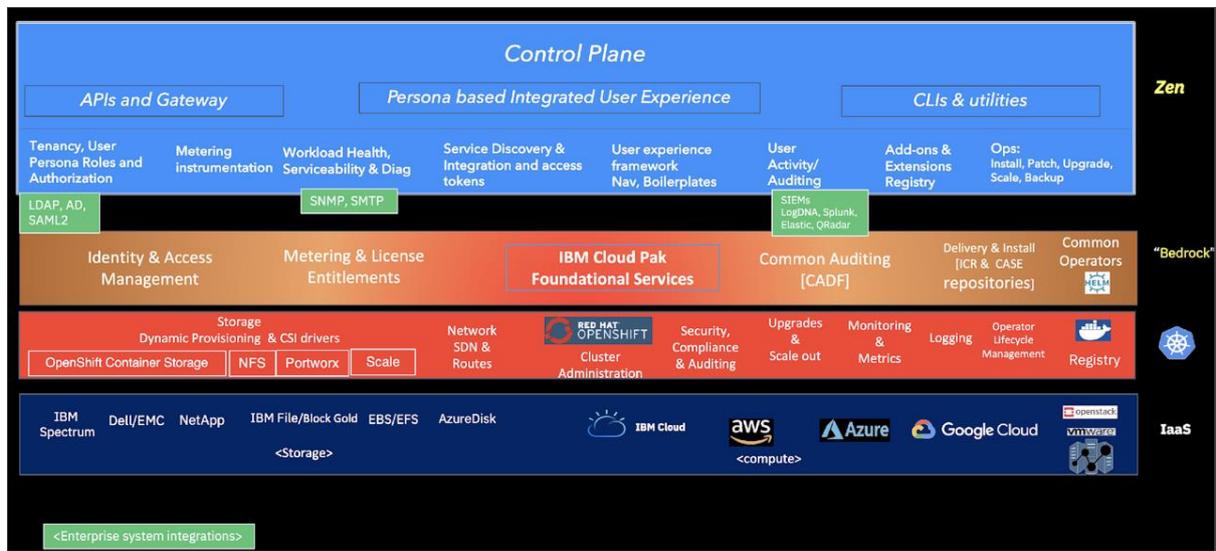


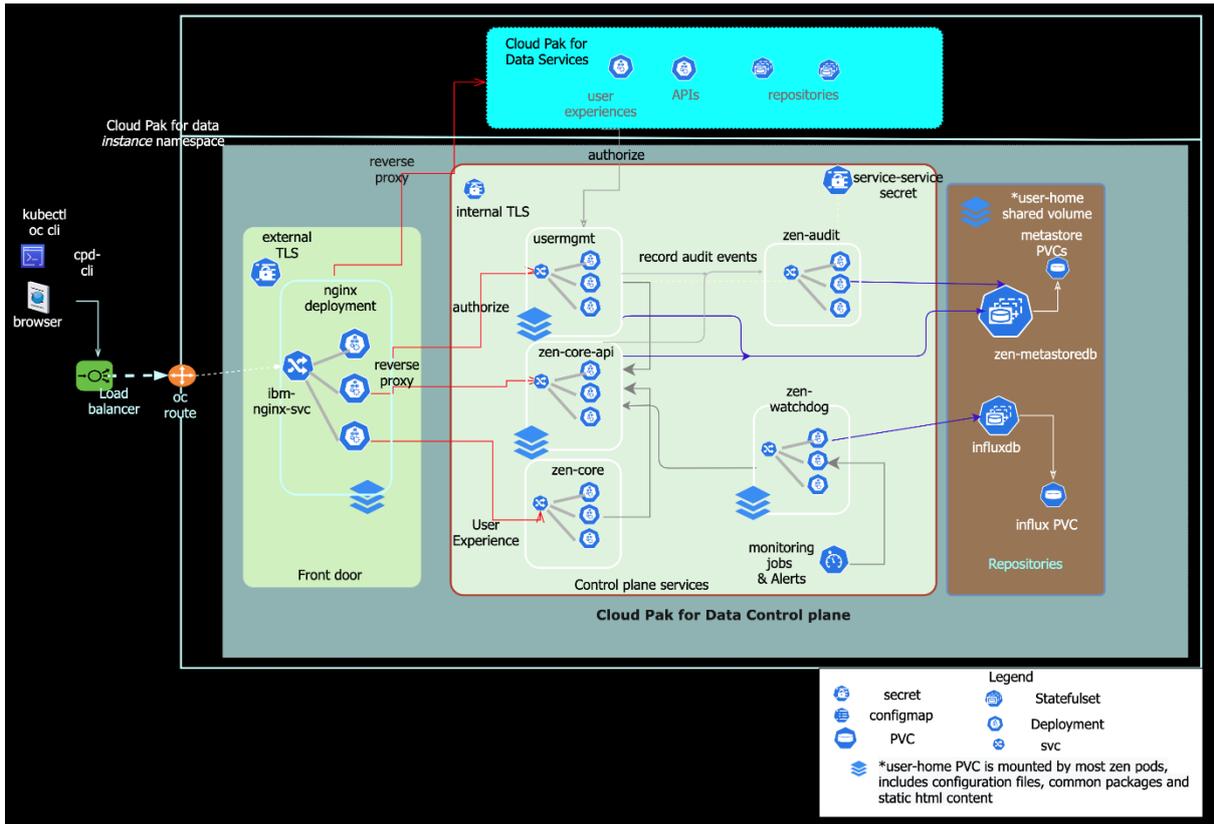
Data analyst



Data steward

# Chapter 9: Technical Overview, Management, and Administration





**Events** →

0 alerts 0 critical events

0 alerts 0 warning events

---

**Status summary**

**Services** 5 total

**Service instances** 0 total

**Environments** 0 total

**Pods** 155 total

3

**Platform resource overview** →

Platform vCPU (cores)	31.52	126.65	Platform memory (GiB)	78.95	285.51
3.30	Requests	Limits	60.54	Requests	Limits
Currently in use			Currently in use		

**Platform resource use - last 12 hrs** →

**Platform quotas** →

**vCPU (Cores)**

- No vCPU quota set

**Memory (GiB)**

- No memory quota set

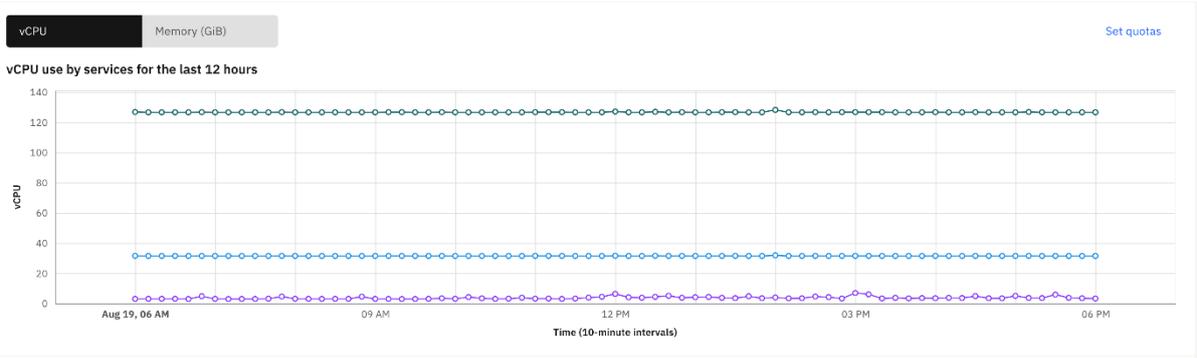
**Service quotas** →

**vCPU (Cores)**

- 0 Quota reached
- 0 Over alert threshold
- 0 Below alert threshold
- 5 No quota set

**Memory (GiB)**

- 0 Quota reached
- 0 Over alert threshold
- 0 Below alert threshold
- 5 No quota set



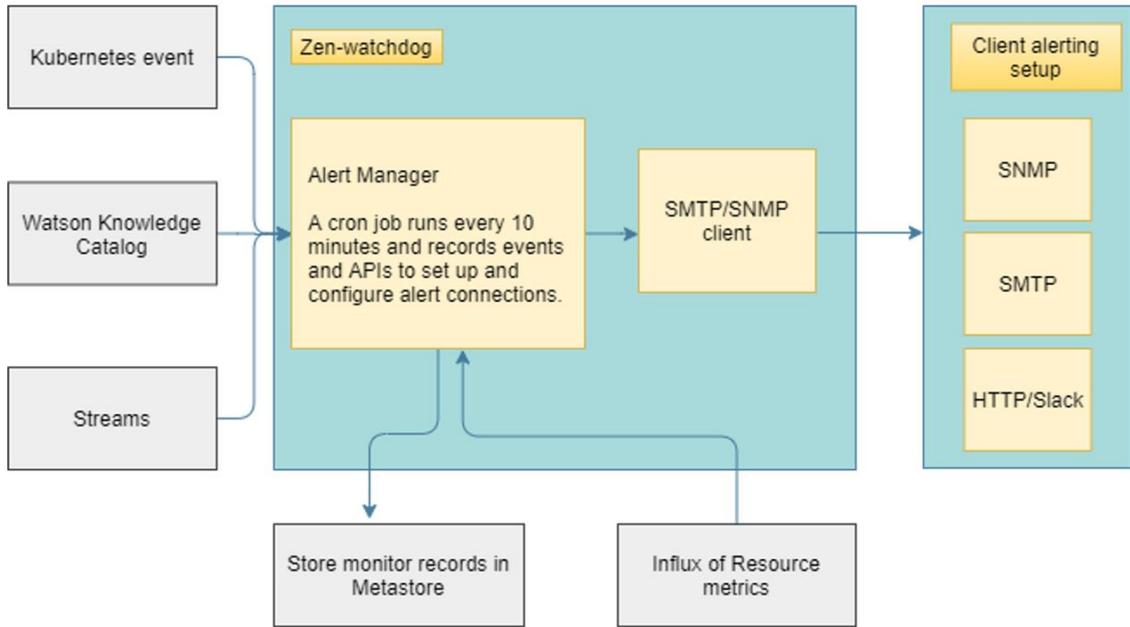
Detailed breakdown by service for: 8/19/2021 6:06 PM (Latest) ⓘ

Find services Filter by: Status vCPU quota status Memory quota status Clear all

Service name	Status	Service instances	Environments	Jobs	vCPU status	Memory status
Total						
Common Core Services	Normal	0	0	0	No quota set	No quota set
Common Database Services	Normal	0	0	0	No quota set	No quota set
Data Refinery	Normal	0	0	0	No quota set	No quota set
Watson Knowledge Catalog	Normal	0	0	0	No quota set	No quota set
IBM Cloud Pak for Data Control Plane	Normal	0	0	0	No quota set	No quota set

Find pods Filter by: Pod status Service Namespace Type Clear all Refresh

Pod name	Service	Current vCPU use	vCPU requests	Current memory use (GiB)	Memory requests (GiB)
Total		3.79	31.52	60.44	78.95
shop4info-mappers-service-0	Watson Knowledg...	0.01	0.20	0.29	0.50
shop4info-rest-0	Watson Knowledg...	0.00	0.20	0.48	1.00
shop4info-scheduler-0	Watson Knowledg...	0.00	0.10	0.33	1.00
shop4info-type-registry-servic...	Watson Knowledg...	0.00	0.10	0.32	0.50
solr-0	Watson Knowledg...	0.01	0.15	0.84	1.00
spaces-7556bf4568-5jz7w	Common Core Ser...	0.01	0.50	0.16	0.50
spawner-api-54bfc796d5-hx6j2	Common Core Ser...	0.00	0.01	0.03	0.06
usermgmt-5b79c55dd9-5j49b	IBM Cloud Pak for...	0.01	0.20	0.09	0.25
usermgmt-5b79c55dd9-kr7j6	IBM Cloud Pak for...	0.00	0.20	0.09	0.25
wdp-activities-85...	Watson Knowledg...	0.06	0.10	0.51	0.78
wdp-connect-con	Common Core Ser...	0.02	0.15	1.40	0.63
wdp-connect-con	Common Core Ser...	0.02	0.20	1.40	0.63
wdp-connect-flig	Common Core Ser...	0.01	0.20	1.50	0.63
wdp-couchdb-0	Common Core Ser...	0.11	1.00	1.80	0.50
wdp-couchdb-1	Common Core Ser...	0.11	1.00	1.90	0.50
wdp-couchdb-2	Common Core Ser...	0.11	1.00	1.90	0.50



### Create a database

- Type
- Configure
- Storage
- Finalize

### Configure

Number of nodes ⓘ

1 ▾

Cores per node

4 ●  128 4

Memory (GB) per node

16 ●  2048 16

Deploy database on dedicated nodes

Value for node label ⓘ

database-db2wh

Namespace ⓘ

Use an existing namespace

zen

#### Storage structure

Select a storage option ⓘ

- Single location for all data
- Separate locations for all data

### Add users ×

Specify which users can access Db2 Warehouse-1, and define their role.

🔍

<input type="checkbox"/>	Name	Username	Role
<input type="checkbox"/>	alex1e	alex1e	User ▾
<input type="checkbox"/>	alexe	alexe	User ▲
<input type="checkbox"/>	alexs	alexs	Admin
<input type="checkbox"/>	alexu	alexu	User
<input type="checkbox"/>	DV User 1 Data Engineer	dv_user1	User ▾
<input type="checkbox"/>	DV user 2 USER	dv_user2	User ▾

Cancel Add

### Storage volumes

Manage the storage volumes that users can access from the platform.

🔍 Find volumes New volume +

Name	Type	Created by	vCPU (Cores)	Memory (GB)	Users	Status	Created on
<a href="#">sample-vol</a>	volumes	admin	0.10	0.10 Gi	1		Nov 9, 2020

### Grant access to users ×

Specify the users who can access Control Plane Volumes and the role of each user.

🔍 Find users

Users Filter by: All ▾

<input type="checkbox"/>	user2 administrator user2, jerry@ibm.com	Choose a role ▲
<input type="checkbox"/>	user1 user user1, carson@ibm.com	Editor ▾
<input type="checkbox"/>	user3 can_provision user3, alex@ibm.com	Viewer ▾
		Admin ▾

## Storage volume details

Volume type ⓘ

Select type ^

- External NFS
- External SMB
- Existing PVC
- New PVC

## Storage volume details

Volume type ⓘ

New PVC v

Storage class ⓘ

ocs-storagecluster-cephfs v

Size in GB ⓘ

1 ⇅

Mount path ⓘ

/mnts/ test-vol

Mount path: /mnts/test-vol

# Storage volumes

Manage the storage volumes that users can access from the platform.

You can enable access to a volume on an external NFS server or to a persistent volume claim (PVC).

## Storage volume overview

Name

sample-vo1

Description (optional)

Description for the volume

## Storage volume details

Volume type ⓘ

External NFS

NFS server ⓘ

zen-dev-bart-master.fyre.ibm.com

Exported path ⓘ

/sample-data

Mount path ⓘ

/mnts/ data

Mount path: /mnts/data

# Watson Knowledge Catalog quotas

×

Monitor platform resource use against your target use

## Current requests and limits ⓘ

vCPU requests 18.51

vCPU limits 57.35

Memory requests (GiB) 70.90

Memory limits (GiB) 154.27

## What do you want to set the quotas on?

Requests ⓘ  Limits ⓘ

## Quotas

vCPU quota

24

vCPU alert threshold

90

% of quota

Memory quota (GiB)

100

Memory alert threshold

95

% of quota

# Chapter 10: Security and Compliance

## LDAP configuration

### LDAP server information

Work with your LDAP administrator to ensure that you have the required information.

LDAP protocol	LDAP hostname	LDAP port	Domain search user (optional) ⓘ
ldaps://	sample.com	636	A user that can perform lookups in the LDAP server
User search base ⓘ	Domain search password (optional)		
For example, dc=sample,dc=com	The password for domain search user		👁
User search field ⓘ			
For example, cn, uid, or sAMAccountName			
<input type="checkbox"/> Use LDAP group ⓘ			
Group search base ⓘ	Group search field ⓘ		

### LDAP attribute mapping

If you want to use LDAP to manage access to the platform, provide the LDAP attributes that map to the following values. (All attributes are required if you plan to use LDAP groups.)

First name (optional)	Group membership
For example, givenName	
Last name (optional)	Group member field
For example, sn	
Email (optional)	
For example, mail	

Users      User groups      **Roles**

Filter by: Enabled permissions ▾

Q Find roles New role +

Name	Description	Modified on	Enabled permissions
<a href="#">Administrator</a>	Administrator role	Aug 4, 2021 3:36 AM	Access governance artifacts + <a href="#">10 more</a> 
<a href="#">Business Analyst</a>	Business analyst role	Aug 4, 2021 3:36 AM	Access information assets view, View data quality  
<a href="#">Data Engineer</a>	Data engineer role	Aug 4, 2021 3:36 AM	Access catalogs + <a href="#">8 more</a>  
<a href="#">Data Quality Analyst</a>	Data quality analyst role	Aug 4, 2021 3:36 AM	Access catalogs + <a href="#">6 more</a>  
<a href="#">Data Scientist</a>	Data scientist role	Aug 4, 2021 3:36 AM	Access catalogs  
<a href="#">Data Steward</a>	Data steward role	Aug 4, 2021 3:36 AM	Access catalogs + <a href="#">6 more</a>  
<a href="#">Developer</a>	Developer role	Aug 4, 2021 3:36 AM	Access catalogs, Create service instances  
<a href="#">User</a>	Default role for users with limited permissions on the platform	Aug 4, 2021 3:36 AM	Access assigned services  

Permission	Category	Services that contribute the permission
Access governance artifacts	Governance artifacts	Watson™ Knowledge Catalog
Administer platform	Platform administration	Cloud Pak for Data control plane
Create deployment spaces	Deployments	– Watson Knowledge Catalog – Watson Studio
Create projects	Projects	– Watson Knowledge Catalog – Watson Studio
Create service instances	Service instances	Cloud Pak for Data control plane
Manage asset discovery	Data curation	Watson Knowledge Catalog
Manage catalogs	Catalogs	– Watson Knowledge Catalog – Watson Studio
Manage data protection rules	Governance artifacts	Watson Knowledge Catalog
Manage data quality	Data curation	Watson Knowledge Catalog
Manage data transformation	Data integration	DataStage®
Manage governance categories	Governance artifacts	Watson Knowledge Catalog
Manage information assets	Catalogs	Watson Knowledge Catalog
Manage metadata	Data curation	Watson Knowledge Catalog
Manage workflows	Workflows	Watson Knowledge Catalog

Permission	Category	Services that contribute the permission
Access catalogs	Catalogs	<ul style="list-style-type: none"> <li>– Watson Knowledge Catalog</li> <li>– Watson Studio</li> </ul>
Create deployment spaces	Deployments	<ul style="list-style-type: none"> <li>– Watson Knowledge Catalog</li> <li>– Watson Studio</li> </ul>
Create projects	Projects	<ul style="list-style-type: none"> <li>– Watson Knowledge Catalog</li> <li>– Watson Studio</li> </ul>

Permission	Category	Services that contribute the permission
Create deployment spaces	Deployments	<ul style="list-style-type: none"> <li>– Watson Knowledge Catalog</li> <li>– Watson Studio</li> </ul>
Create projects	Projects	<ul style="list-style-type: none"> <li>– Watson Knowledge Catalog</li> <li>– Watson Studio</li> </ul>

## Permissions

Add permissions needed for this new role. Click labels to view details. Click checkboxes to add.

🔍 Find permissions

- Catalogs ▼
- Data curation ▼
- Deployments ▼
- Governance artifacts ▼
- Platform administration ▲
- Administer platform**
- Manage configurations
- Manage platform health
- View platform health
- Projects ▼
- Service instances ▲
- Create service instances
- Manage service instances
- User administration ▲
- Manage platform roles
- Manage user groups
- Manage users

### Administer platform

#### Description

This permission offers the most comprehensive set of actions for managing and monitoring the platform.

Users with this permission have elevated privileges and can grant or revoke all permissions, including other administrative permissions.

#### 38 actions

See the actions listed in the following permissions:

- Manage configurations
- Manage platform health
- Manage platform roles
- Manage users
- Manage user groups
- Manage service instances

Users | User groups | Roles

Filter by: Roles ▾

Find users New user +

Name	User ID	Username	Created on	Roles
 admin	1000330999	admin	May 31, 2021 3:43 AM	Administrator + 6 more

Administrator  
 Business Analyst  
 Data Engineer  
 Data Quality Analyst  
 Data Scientist  
 Data Steward  
 Developer

## Add user

Authorize user access to the platform.

- Profile information
- Platform access
- Roles
- Summary

### Roles

Assign roles to this user or [create a new role](#).

Search: administrator

<input checked="" type="checkbox"/> Administrator	<h4>Administrator</h4> <p><b>Description</b> Administrator role</p> <p><b>Modified on</b> Jul 20, 2021 4:34 AM</p> <p><b>24 permissions, 110 actions</b> <a href="#">Expand all</a></p> <hr/> <p>Access advanced governance</p> <hr/> <p>Access advanced mapping</p> <hr/> <p>Access catalogs</p> <hr/> <p>Access data quality</p> <hr/> <p>Access governance artifacts</p> <hr/> <p>Administer platform</p>
---	--

## New user group

Form groups of users to widely administer permissions.

- Details
- Users**
- Roles
- Summary

### Users

Select users to be added to the user group. You can also add users or groups of users directly from your connected identity providers.

Existing users | Identity provider users | Identity provider groups

Selected 2 existing users

	Name	Username	Email	Roles
<input checked="" type="checkbox"/>	builder	builder	builder@hursley.igc.ibm.com	Data Steward
<input checked="" type="checkbox"/>	Developer	developer	developer@ibm.com	Developer

## New user group

Form groups of users to widely administer permissions.

- Details
- Users**
- Roles
- Summary

### Users

Select users to be added to the user group. You can also add users or groups of users directly from your connected identity providers.

Existing users | Identity provider users | **Identity provider groups**

Selected None

Search for the identity provider groups you want to add to this user group.

  
2 results returned  
CN=aaa group,CN=Groups,DC=hursely,DC=igc,DC=ibm,DC=com  
CN=Allowed RODC Password Replication Group,CN=Groups,DC=hursely,DC=igc,DC=ibm,DC=com

# New user group

Cancel

Back

Next

Details

Users

Roles

Summary

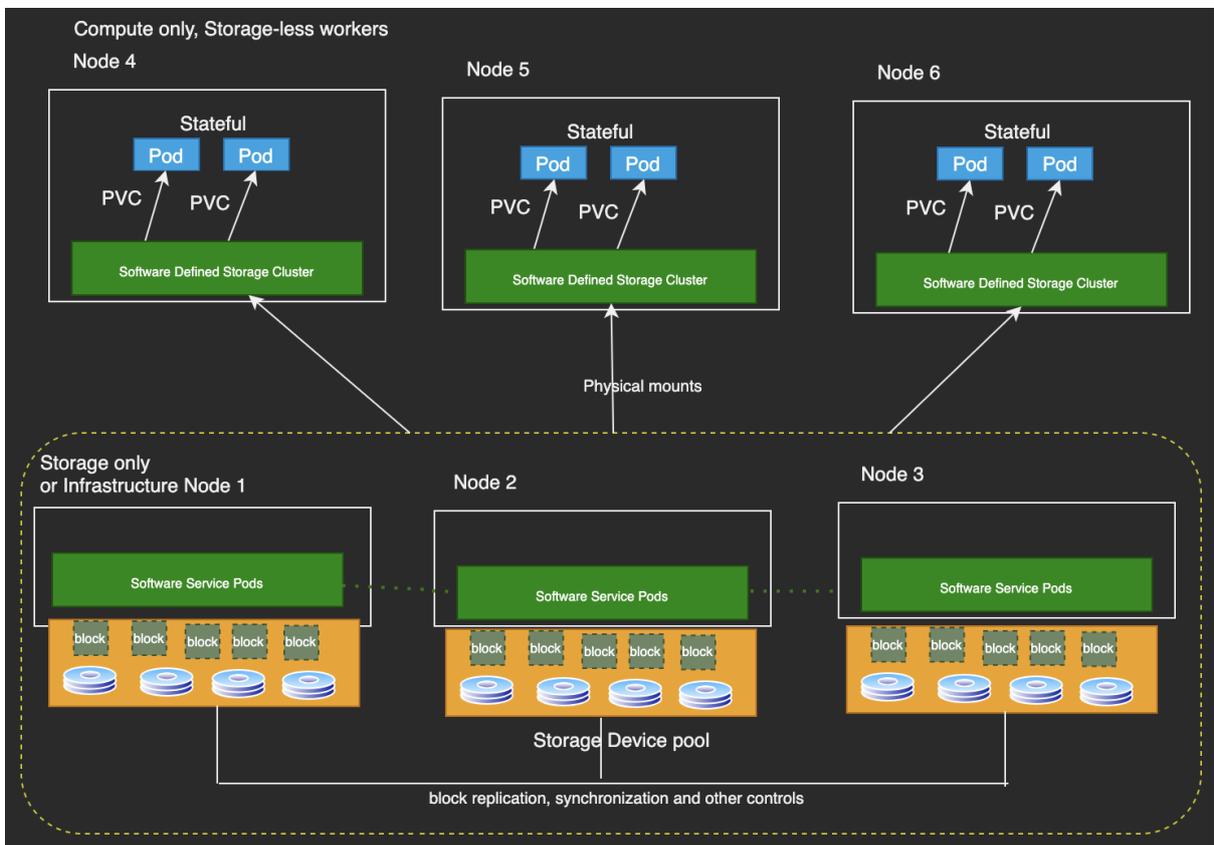
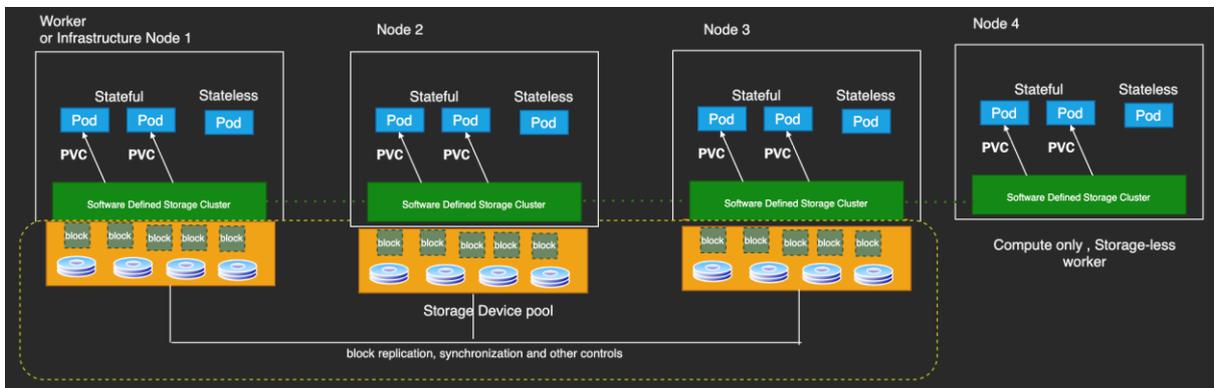
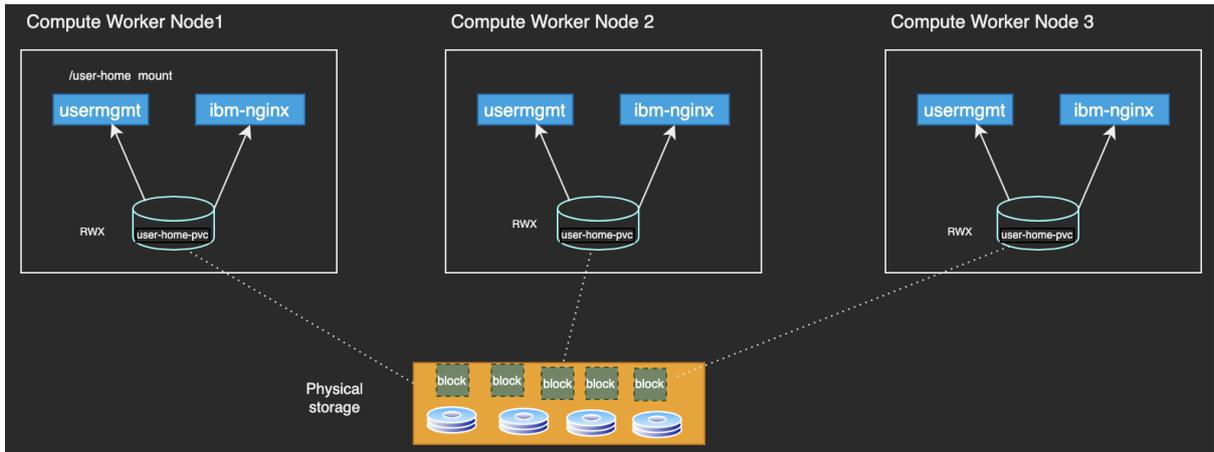
## Roles

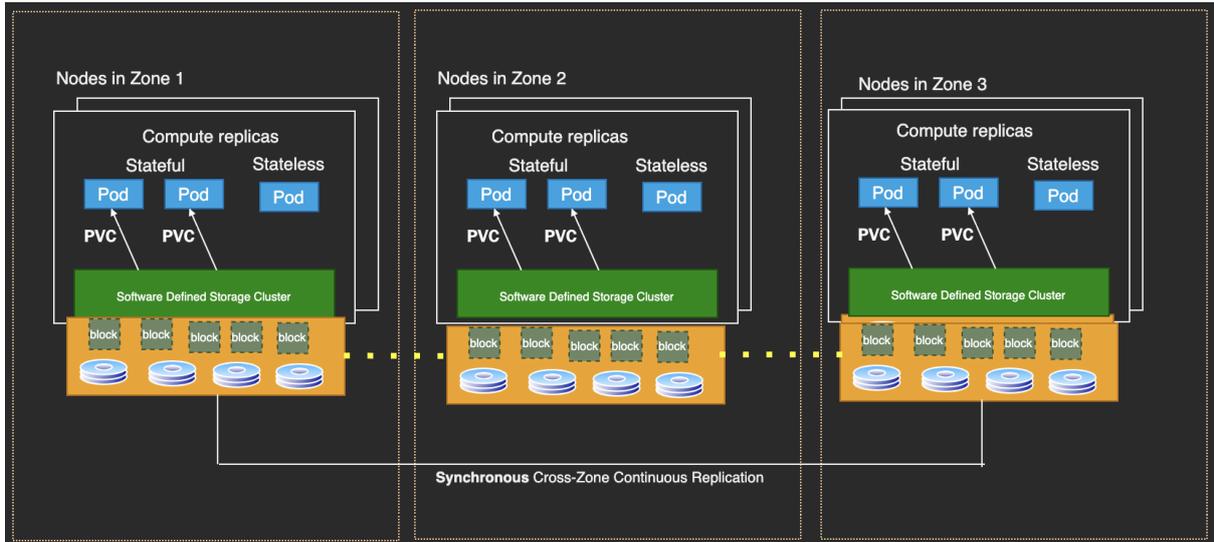
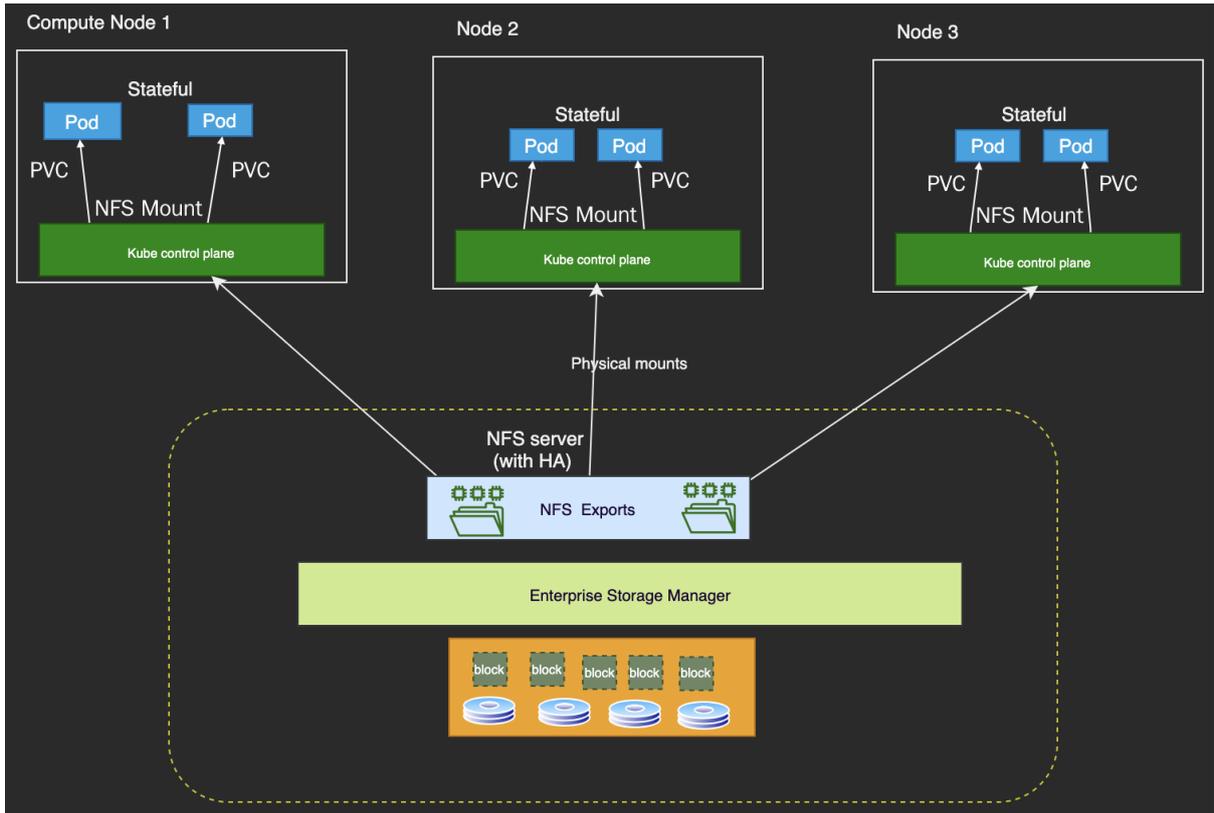
Select the roles to assign to this group. If you can't find the appropriate role in the list, you can [create a new role](#).

Find roles

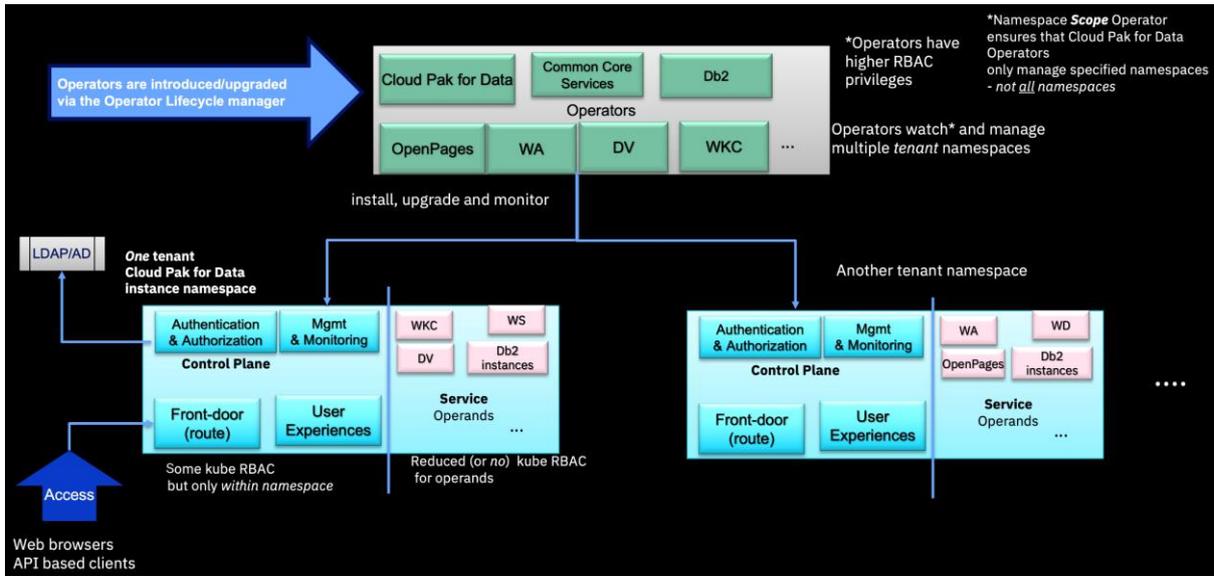
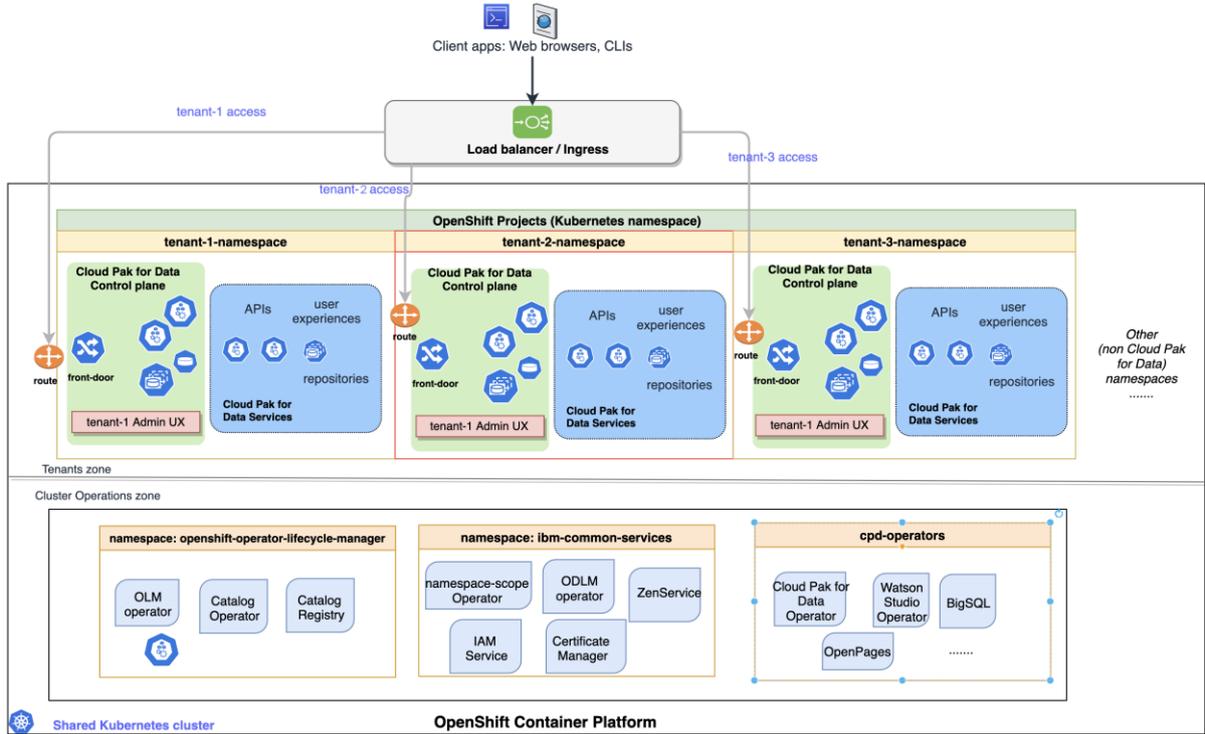
<input type="checkbox"/>	Name	Enabled permissions
<input type="checkbox"/>	Administrator	Access governance artifacts + <a href="#">10 more</a>
<input type="checkbox"/>	Business Analyst	Access information assets view, View data quality
<input type="checkbox"/>	Data Engineer	Access catalogs + <a href="#">8 more</a>
<input type="checkbox"/>	Data Quality Analyst	Access catalogs + <a href="#">6 more</a>
<input type="checkbox"/>	Data Scientist	Access catalogs
<input checked="" type="checkbox"/>	Data Steward	Access catalogs + <a href="#">6 more</a>
<input type="checkbox"/>	Developer	Access catalogs, Create service instances
<input type="checkbox"/>	User	Access assigned services

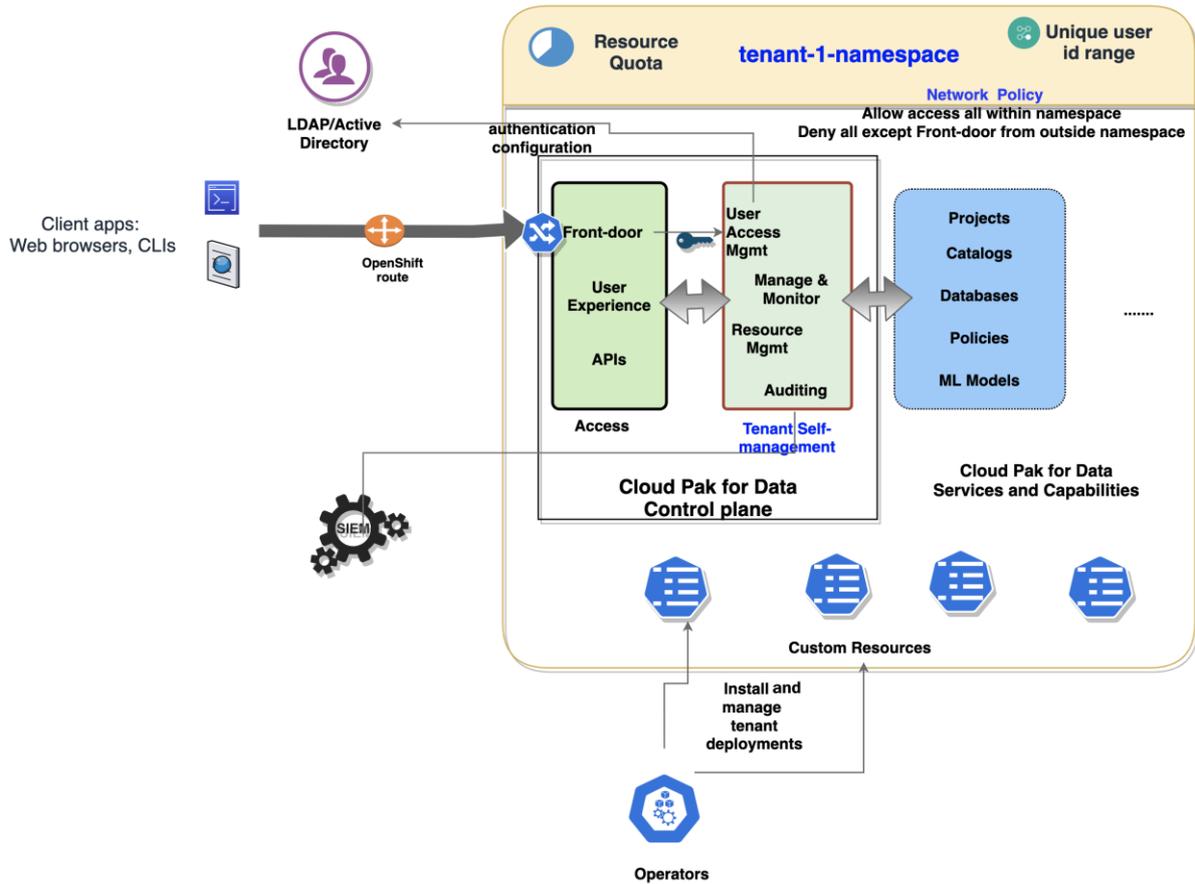
# Chapter 11: Storage



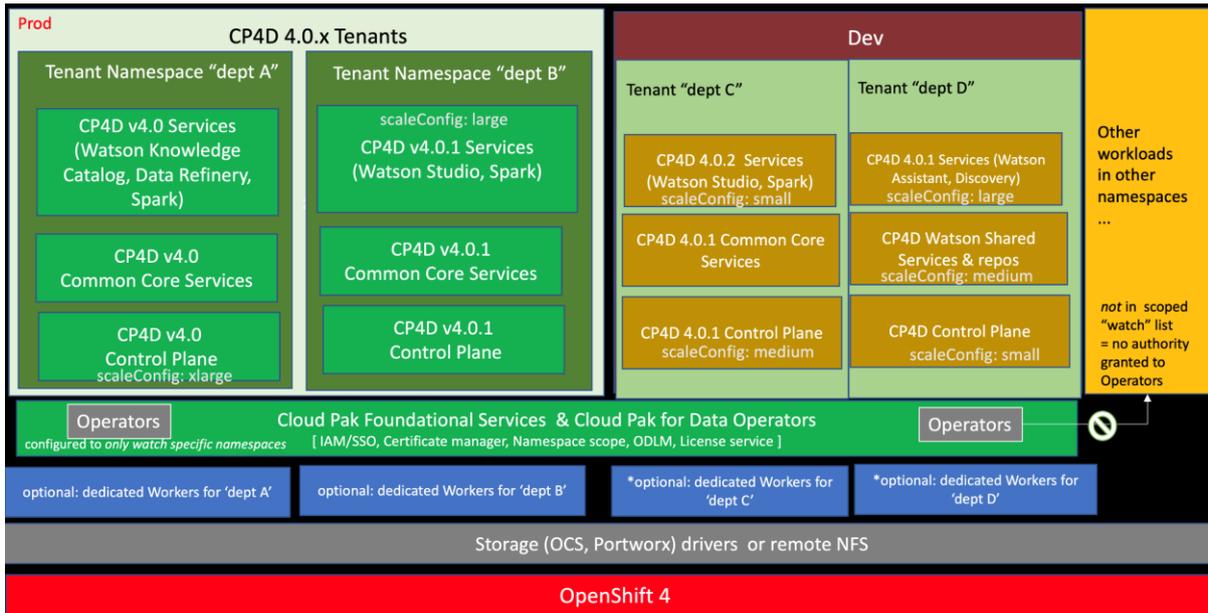


# Chapter 12: Multi-Tenancy





Tenancy Aspects	Mechanism
<b>Isolation</b> 1. Users – Authentication, Roles and Authorization 2. Compute Resources 3. Access to Instance 4. Storage Usage	<i>Individual Kube Namespace quotas/OpenShift Projects for each Cloud Pak for Data instance</i> <ul style="list-style-type: none"> <li>✓ If needed, can support <i>different</i> LDAP/AD or SAML configurations <i>per</i> instance (each instance has its own unique user management service)</li> <li>✓ Namespace quotas – set by the cluster admin.</li> <li>✓ Unique DNS-based routes for each instance.</li> <li>✓ Completely separated PVCs (in different namespaces) with different PVs – no sharing <i>across</i> instances.</li> </ul>
<b>Security</b> 1. Restricting Kube Cluster-Level Privileges 2. Network Access 3. Auditing	<i>RBAC privileges &amp; Service Accounts scoped only to that namespace</i> <ul style="list-style-type: none"> <li>✓ Service accounts and role binding – scoped to Project namespaces.</li> <li>✓ Kube network policies for hardening <i>like a firewall around an RHOS Project namespace</i>.</li> <li>✓ <i>Per-tenant</i> SIEM forwarding supported.</li> </ul>
<b>Management</b> 1. Monitoring/Metering 2. Ops – Backup/DR, Scale, Patch, Upgrade 3. Serviceability 4. Service Installations	<i>Administration privileges scoped to a namespace</i> <ul style="list-style-type: none"> <li>✓ Different tenant “Admin users for different instances. Tenants are managed independently.”</li> <li>✓ Ops privileges scoped to that OpenShift project’s namespace.</li> <li>✓ Serviceability utilities only capture diagnostics that are appropriate for that instance.</li> <li>✓ Different instances can have different sets of services at different versions scaled independently from other tenants.</li> </ul>



Tenancy Aspects	Mechanism
<b>Isolation</b>	<i>Shared instance in one Kubernetes namespace</i>
<ol style="list-style-type: none"> <li>Users – Authentication, Roles and Authorization</li> <li>Compute Resources</li> <li>Access to Instance</li> <li>Storage Usage</li> </ol>	<ul style="list-style-type: none"> <li>✓ Authorization controls with groups and roles with only one LDAP/AD/SAML configuration possible. <ul style="list-style-type: none"> <li>○ Note: Some services are still shared, and some admin roles cannot be <i>exclusively</i> granted to one sub-tenant.</li> </ul> </li> <li>• No quotas for individual sub-tenants.</li> <li>• A single OpenShift route for all users; harder to support sub-tenant-specific DNS domain names.</li> <li>• No separation of persistent volumes.</li> </ul>
<b>Security</b>	<i>RBAC privileges and service accounts scoped to that namespace</i>
<ol style="list-style-type: none"> <li>Restricting Kube Cluster-Level Privileges</li> <li>Network Access</li> <li>Auditing</li> </ol>	<ul style="list-style-type: none"> <li>✓ No cluster-level privileges.</li> <li>• No network isolation between sub-tenant user workloads.</li> <li>✓ Auditing covers <i>all</i> sub-tenant activities and individual users that have been uniquely identified, with only one SIEM configuration possible.</li> </ul>
<b>Management</b>	<i>Administration and Ops privileges cannot be granted to individual users</i>
<ol style="list-style-type: none"> <li>Monitoring/Metering</li> <li>Ops – Backup/DR, Scale, Patch, Upgrade</li> <li>Serviceability</li> <li>Service Installations</li> </ol>	<ul style="list-style-type: none"> <li>✓ A primary tenant "admin" assignment can ensure security. "Sub-tenants" should not be granted an admin role.</li> <li>• Day 2 operations cannot be at a sub-tenant scope.</li> <li>• Serviceability and diagnostics cannot be done at the sub-tenant level. The primary tenant admin can view all logs from all sub-tenants.</li> <li>• The responsibility of the tenant admin (or cluster admin) - cannot be at a sub-tenant scope. Different service versions for different sub-tenants cannot be supported.</li> </ul>