

Data Mesh

Was ist ein Datenprodukt?



JOCHEN CHRIST
@JOCHEN_CHRIST

Hi,
I am Jochen

Jochen Christ

Senior Consultant at INNOQ



Java



Software Architecture



Data-driven Development



O'REILLY®

Deutsche
Ausgabe

Data Mesh

Eine dezentrale Datenarchitektur entwerfen



Zhamak Dehghani

Vorwort von Martin Fowler

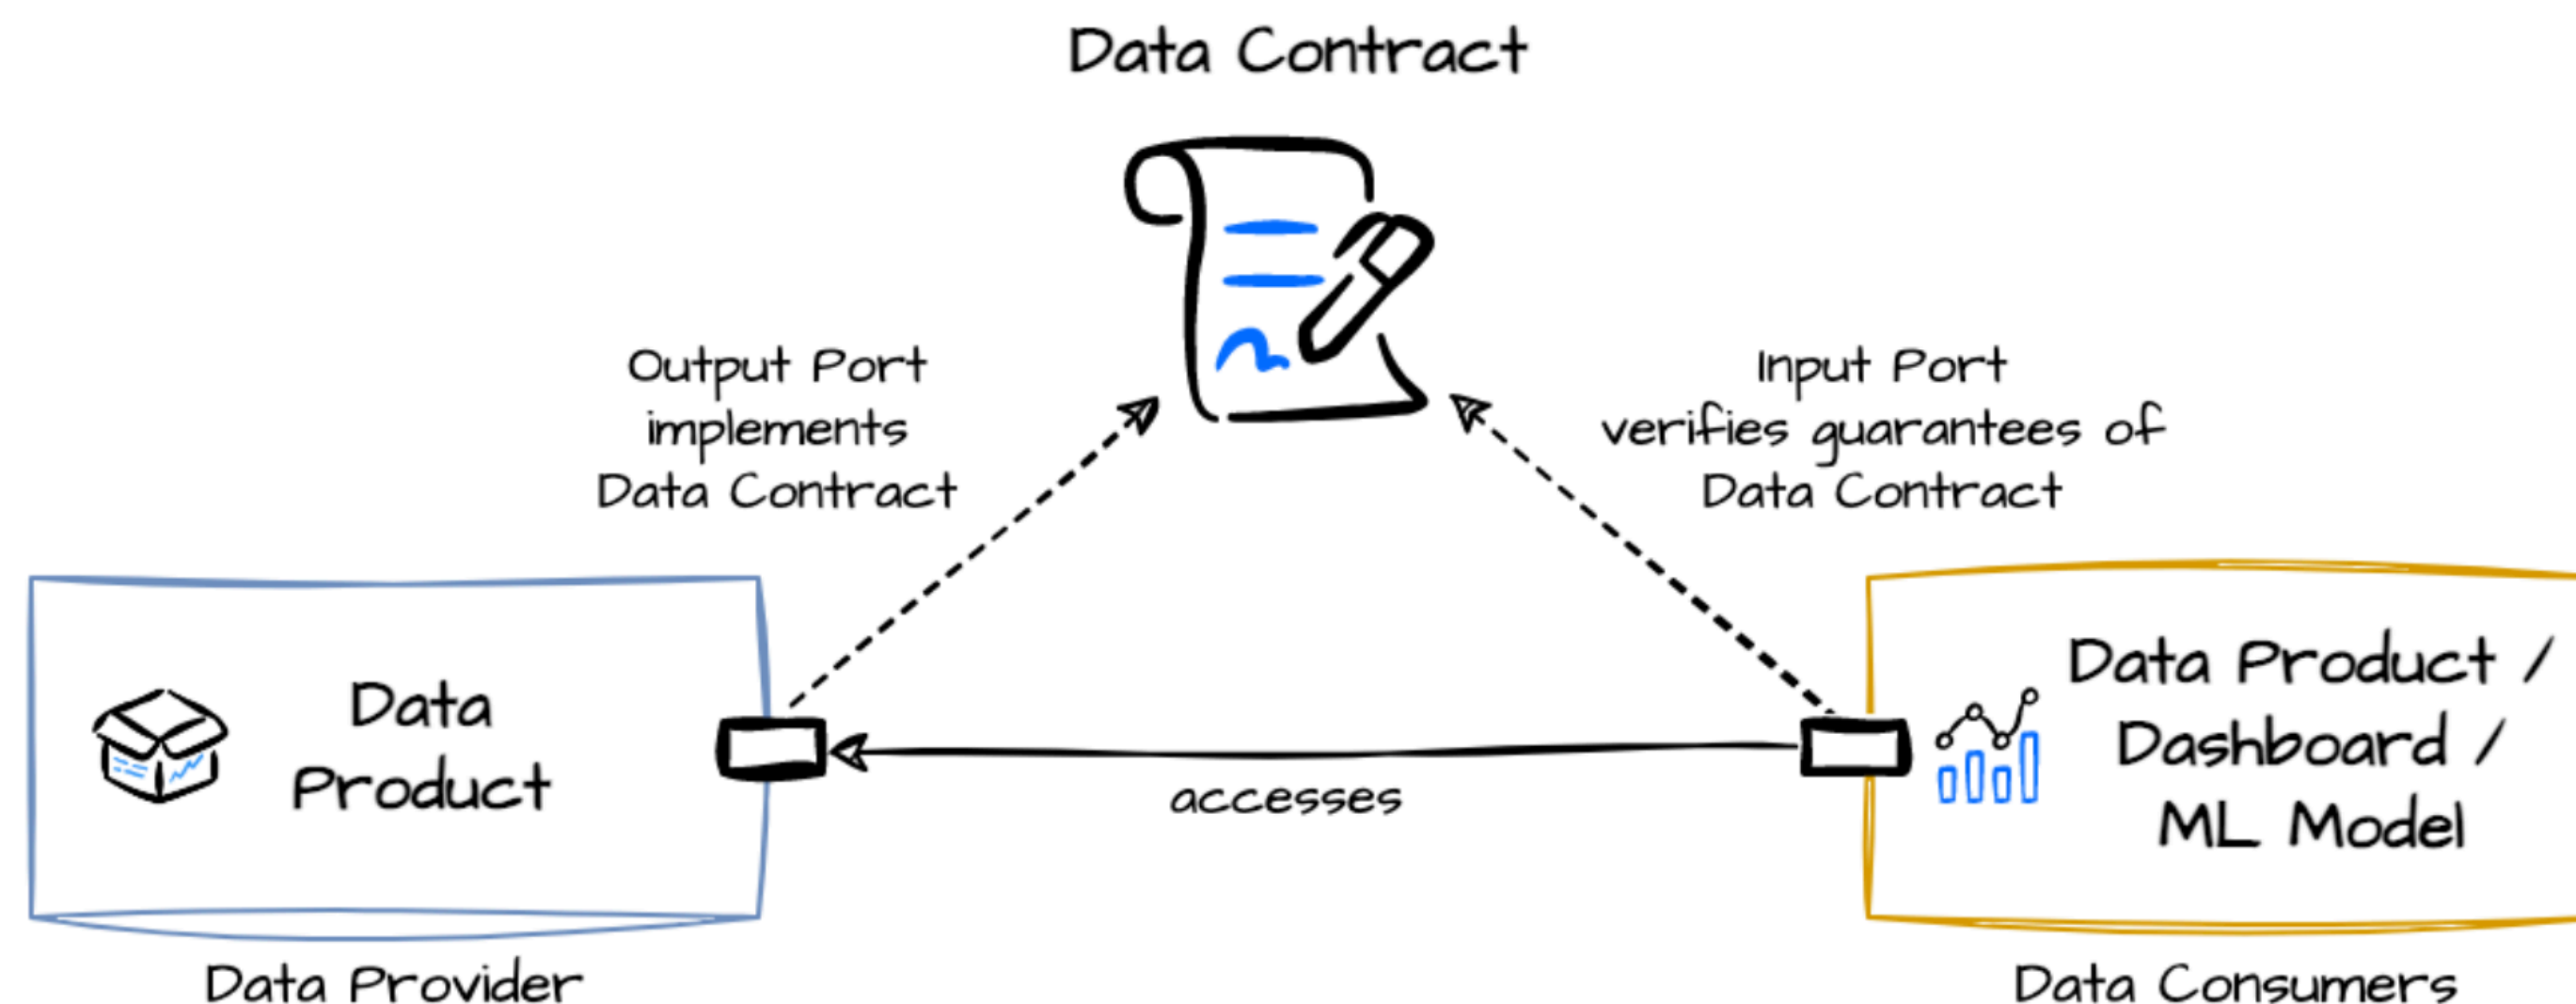
Übersetzung von Jochen Christ und Simon Harrer



datacontract.com



Data Contract Specification



Data contracts bring data providers and data consumers together.

A *data contract* is a document that defines the structure, format, semantics, quality, and terms of use for exchanging data between a data provider and their consumers. A data contract is implemented by a data product's output port or other data technologies. Data contracts can also be used for the input port to specify the expectations of data dependencies and verify given guarantees.

The *data contract specification* defines a YAML format to describe attributes of provided data sets. It is data platform neutral, yet supports well-known formats to express schemas (e.g., dbt models, JSON Schema, Protobuf, SQL DDL) and quality tests (e.g., SodaCL, SQL queries) to avoid unnecessary abstractions. The data contract specification is an open initiative to define a common data contract format. Think of an [OpenAPI specification](#), but for data sets.

datacontract.com



Home > Data Products > Search Queries All

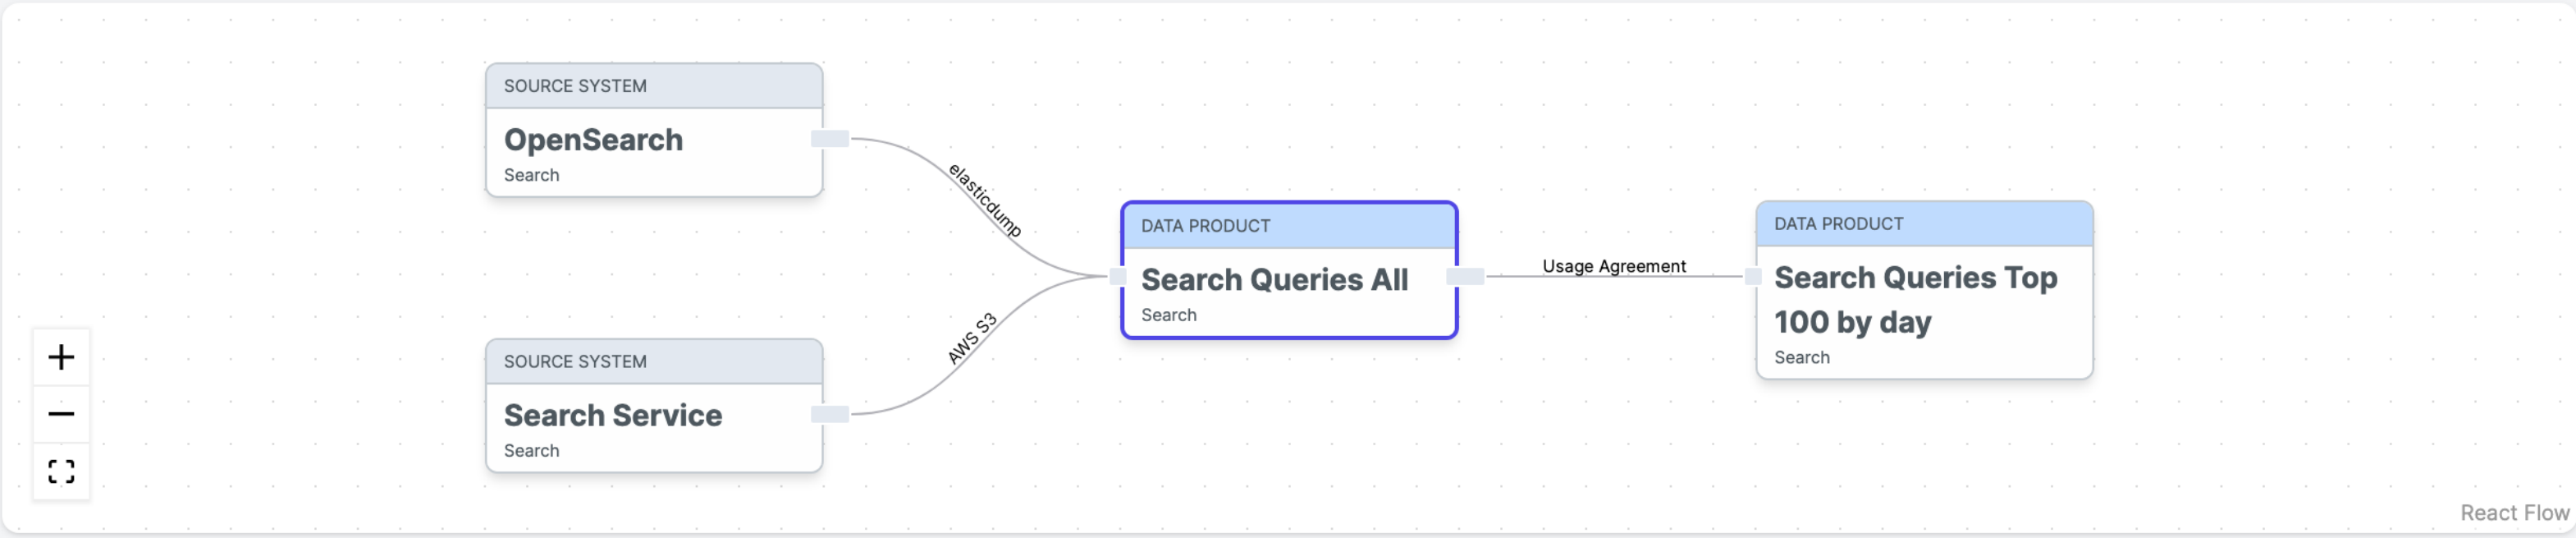
Search Queries All

Show specification

Edit

Request Access

Search source-aligned active managed demo



Info

Information about the data product

Data Catalog

Repository

Documentation

Name: Search Queries All
ID: urn:dataprodut:search:search-queries-all

Description: All search queries and results since 2020.

Metrics

Monitor business value, costs, and compliance

Consumers
1

Costs
\$6,560.00

Compliance
0 / 1 policies

Input Ports

The source of the data, source systems or other data products.

OpenSearch
acme.search.clicks

Output Ports

Technology, dataset, and version of provided data.

search_queries_all_v1
SEARCH_DB.SEARCH_QUERIES_ALL_NPII_V1

Data Contract

1 usage

Recap

Data Mesh

Decentralized Data Architecture

What?

**A decentralized data architecture gives
ownership and competence
for (analytical) data
to the teams that
understand the business context.**

-- Jochen

Decentralized Data Architecture

Why?



Make qualified data-driven decisions in your domain

Use data to better understand your users and system behavior. Derive features from insights, qualify value, and fast iterations. Also qualified rejection of unnecessary tasks.

Do the right things, purpose, motivation



Build innovative services in your domain

Enhance your customer experience with data technologies, such as LLMs, visualizations, classifications, and ML models for predictions and recommendations.

Customer value through innovation



Provide data as business value for other domains

Domain data is valuable for other business units as reference data and to aggregate. Needs managed, explained, high-quality and easy accessible data as products.

Company success

What Is Data Mesh?

Strategic
Domain-driven
Design

Socio-technical
Perspective

Technology

Domain
Ownership

Domain
Bounded Context

Domain Teams

Operational &
Analytical Data

Data as a
Product

Product Thinking

Data Product by
Domain Team

Interoperability
Interfaces

Self-serve
Data Platform

Domain-agnostic

Data Platform
Team

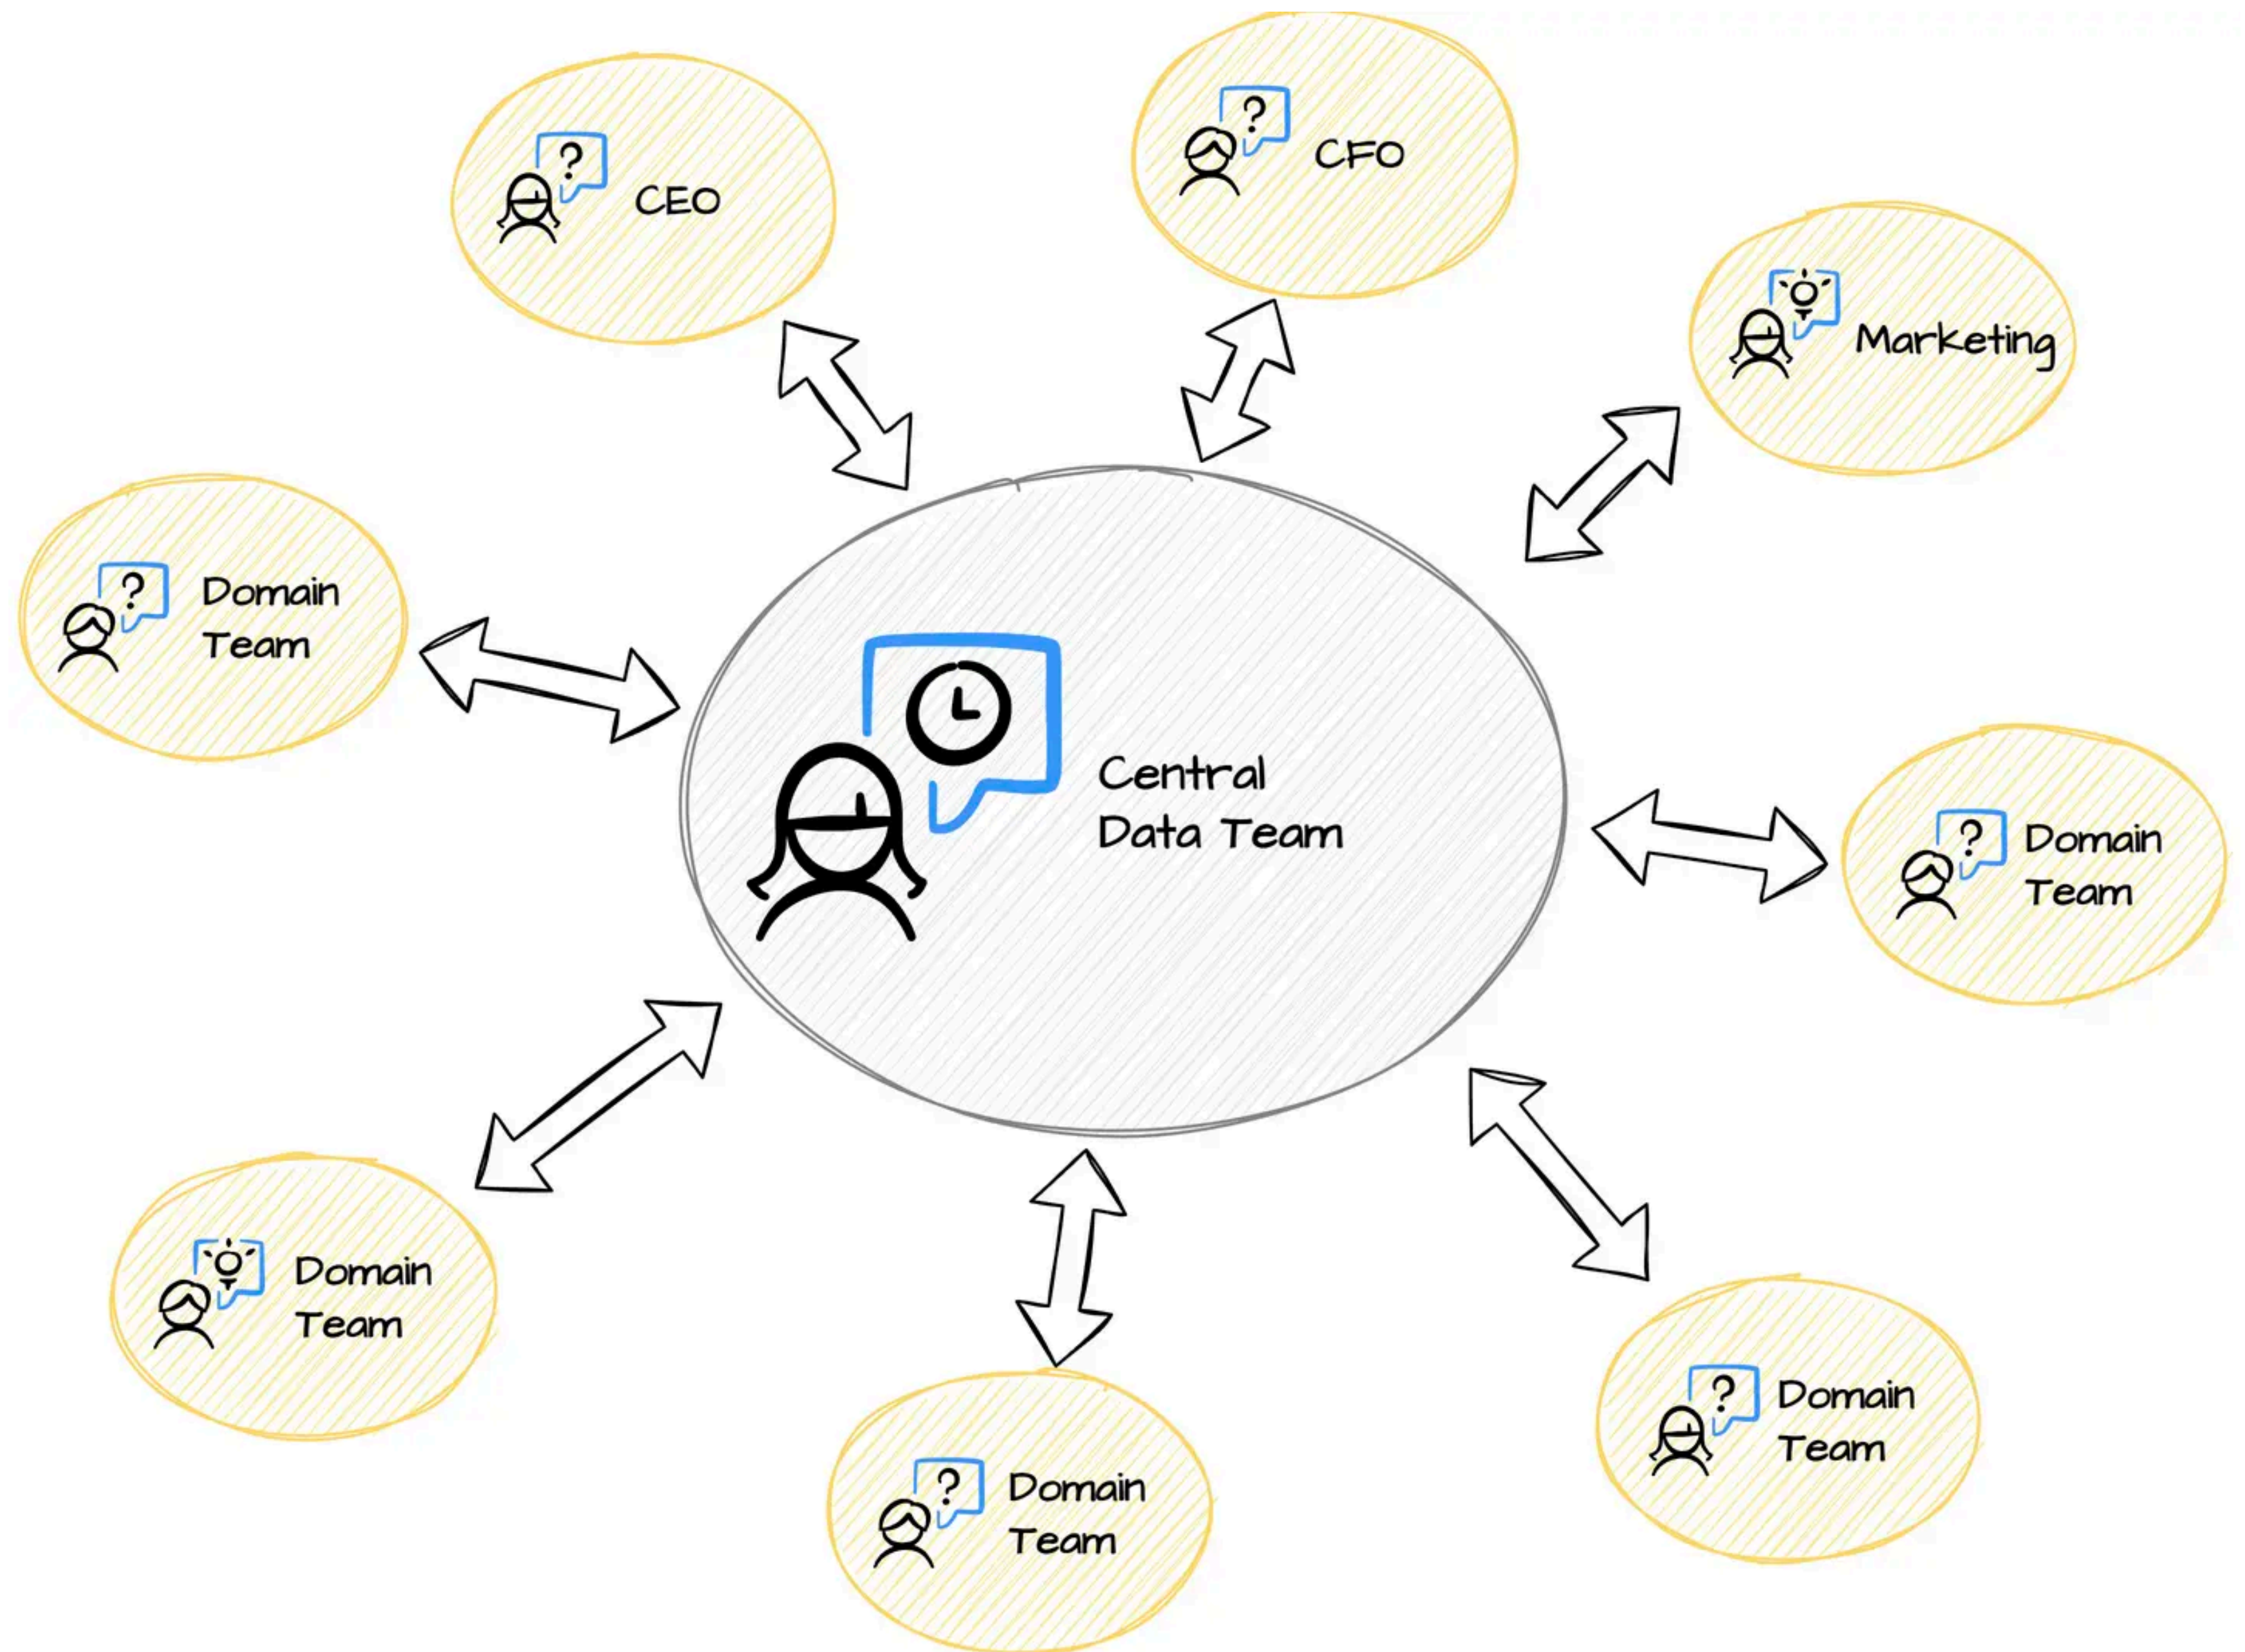
Self-serve
Data Platform


Federated
Governance

Context Mapping

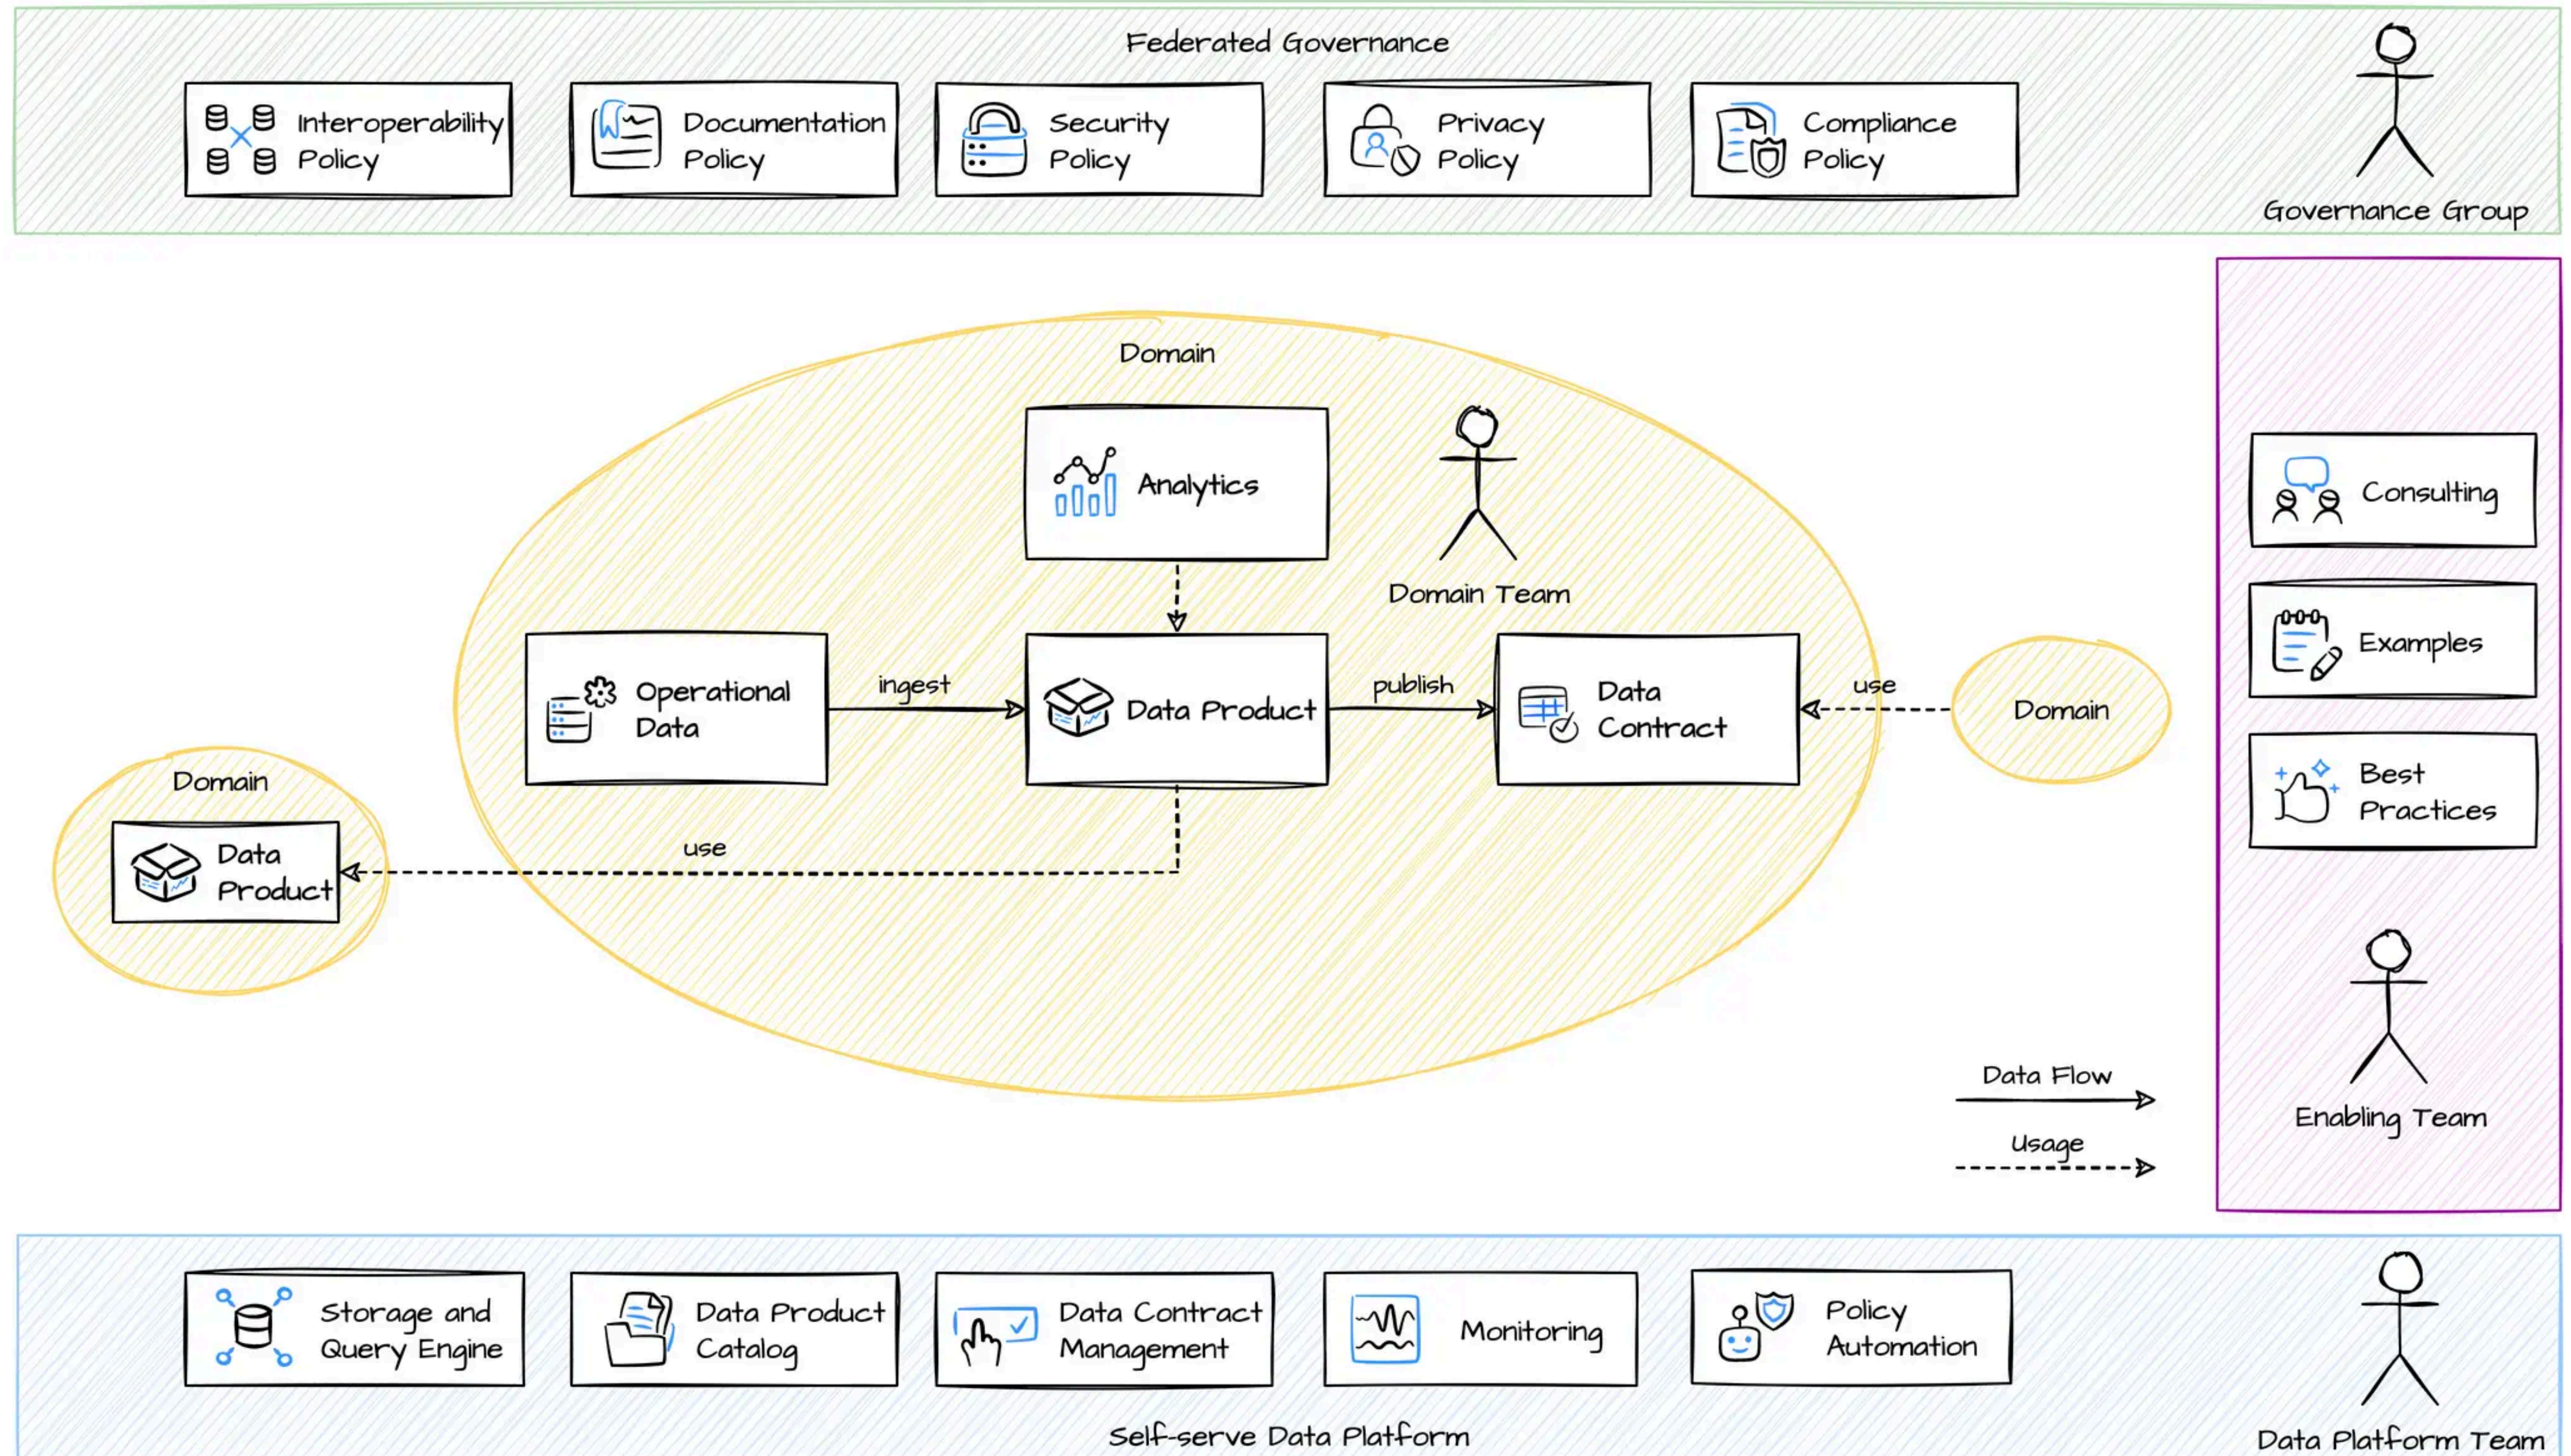
Governance
Group

Policy
Automation

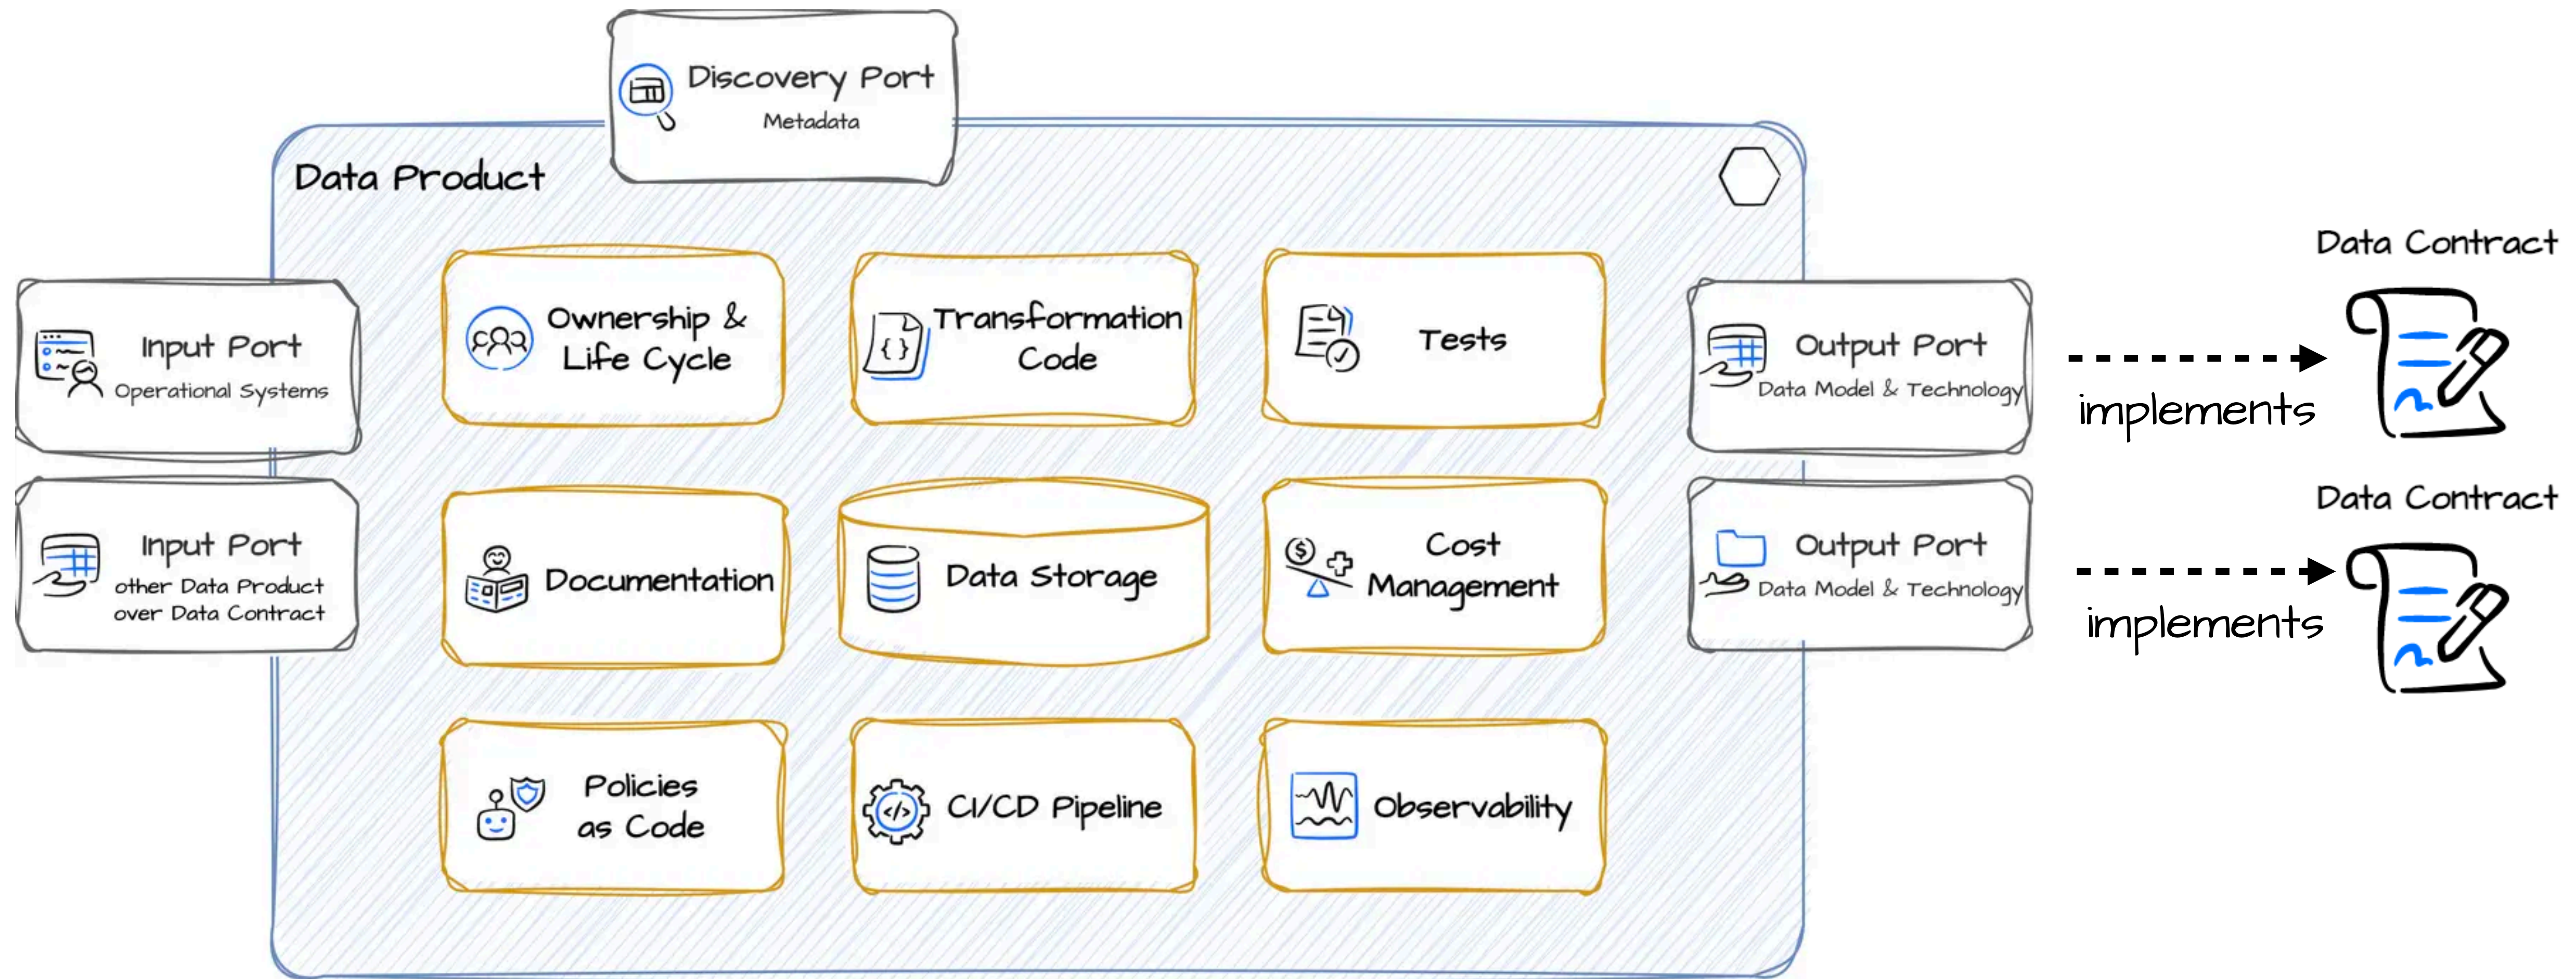




Data Mesh Architecture




Data Products are Modules



Example

Data Product Orders for Team Recommendations



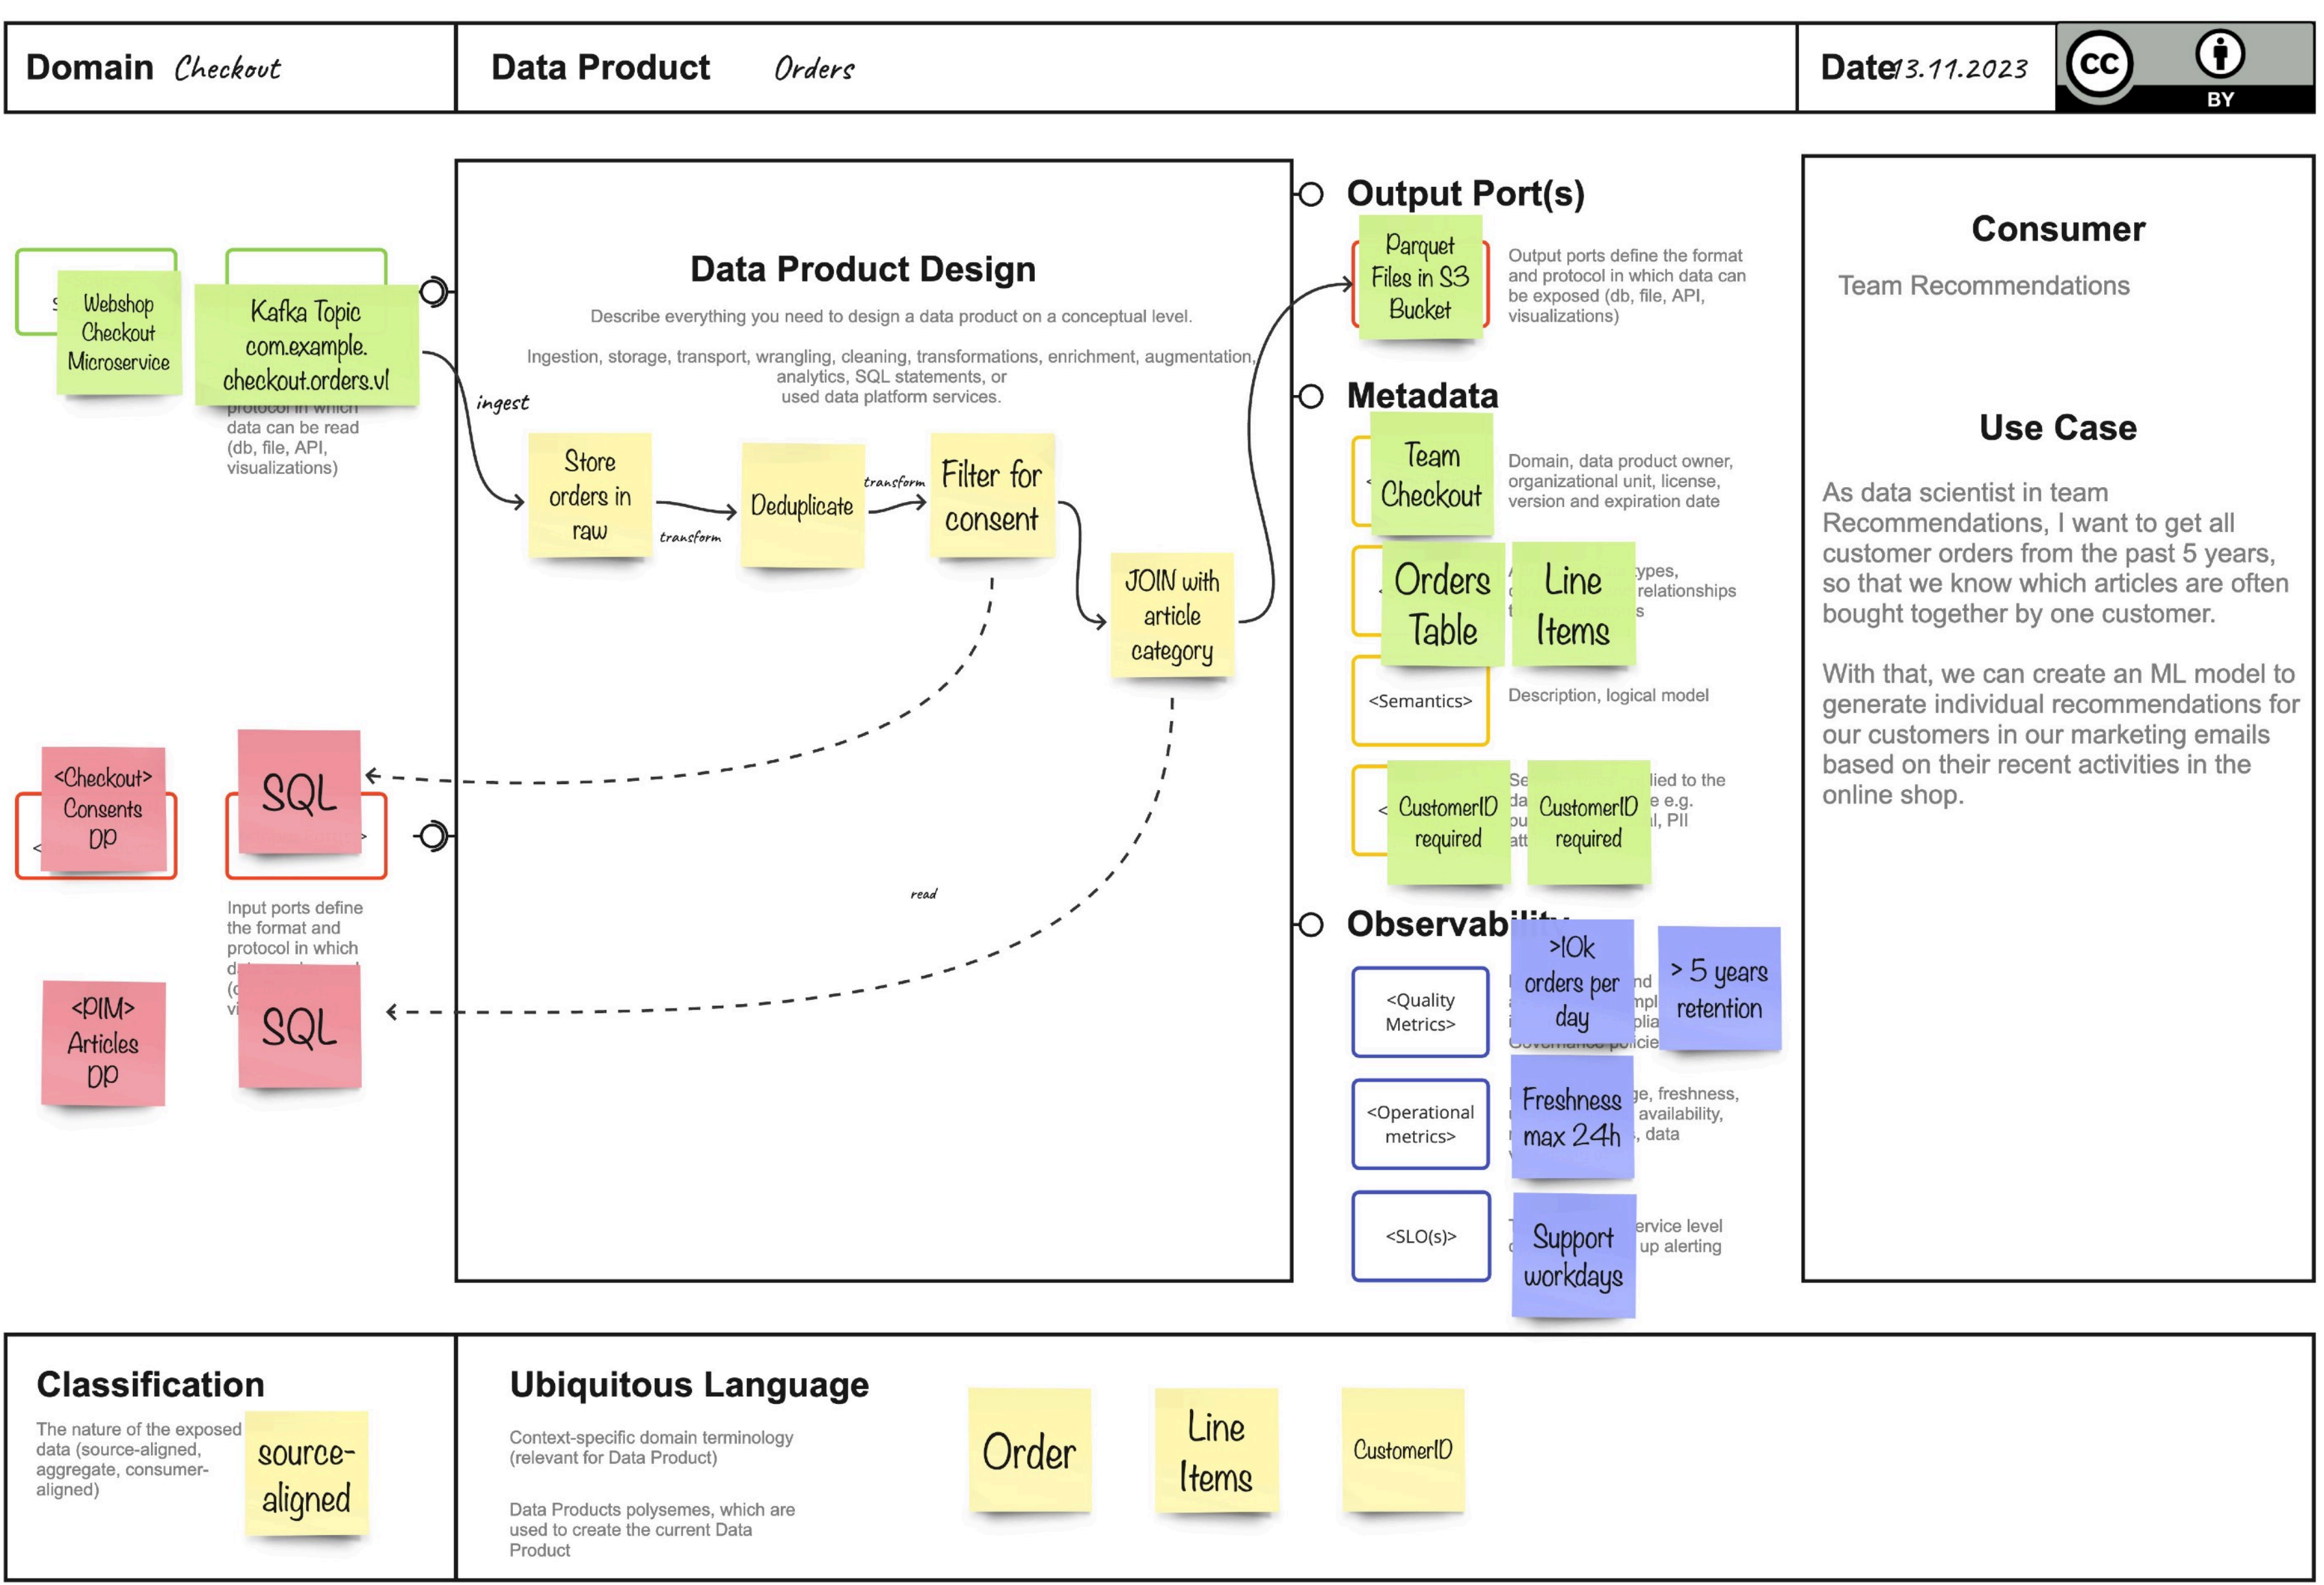
Drive Data Product Development by Business Case

"As data scientist in team Recommendations, I want to get all webshop orders from the past 5 years, so that we know which articles (categories) are often bought together by one customer.

With that, we can create an ML model to generate individual recommendations for our customers in our marketing emails based on their recent activities in the online shop."



Sophia, Data Scientist



Work with Examples

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	
ws_sold_date	ws_sold_time	ws_ship_date	ws_item_sk	ws_bill_customer	ws_bill_cder	ws_bill_header	ws_bill_address	ws_ship_customer	ws_ship_cde	ws_ship_header	ws_ship_address	ws_web_page	ws_web_site	ws_ship_month	ws_warehouse	ws_promo_sku	ws_order_number	
2451383	73313	2451482	4591	83074	596485	1096	40907	85919	41329	1140	1351	43	4	4	5	6	1	
2451383	73313	2451411	3566	83074	596485	1096	40907	85919	41329	1140	1351	28	7	3	2	271	1	
2451383	73313	2451413	7286	83074	596485	1096	40907	85919	41329	1140	1351	58	28	10	5	300	1	
2451383	73313	2451393	2755	83074	596485	1096	40907	85919	41329	1140	1351	2	7	6	1	63	1	
2451383	73313	2451502	2516	83074	596485	1096	40907	85919	41329	1140	1351	56	16	2	5	18	1	
2451383	73313	2451421	16966	83074	596485	1096	40907	85919	41329	1140	1351	50	7	2	2	185	1	
2451383	73313	2451457	10402	83074	596485	1096	40907	85919	41329	1140	1351	56	16	16	3	293	1	
2451383	73313	2451430	1735	83074	596485	1096	40907	85919	41329	1140	1351	25	19	14	3	202	1	
2451383	73313	2451458	15464	83074	596485	1096	40907	85919	41329	1140	1351	13	26	8	5	49	1	
2452625	34964	2452675	8925	42296	436090	2684	37278	57428	1447713	3907	30901	12	6	20	2	107	2	
2452625	34964	2452716	11041	42296	436090	2684	37278	57428	1447713	3907	30901	55	30	9	1	253	2	
2452625	34964	2452715	645	42296	436090	2684	37278	57428	1447713	3907	30901	36	3	6	2	171	2	
2452625	34964	2452729	12453	42296	436090	2684	37278	57428	1447713	3907	30901	19	27	18	5	289	2	
2452625	34964	2452741	13831	42296	436090	2684	37278	57428	1447713	3907	30901	6	12	19	2	75	2	
2452625	34964	2452725	9559	42296	436090	2684	37278	57428	1447713	3907	30901	49	15	11	1	54	2	
2452625	34964	2452743	8085	42296	436090	2684	37278	57428	1447713	3907	30901	12	13	5	1	294	2	
2452625	34964	2452702	12081	42296	436090	2684	37278	57428	1447713	3907	30901	9	13	16	3	267	2	
2452625	34964	2452638	15739	42296	436090	2684	37278	57428	1447713	3907	30901	1	13	16	2	221	2	
2451754	1529	2451774	17431	74800	1219525	3450	3375	84180	1487225	5550	41475	14	20	7	5	246	3	
2451754	1529	2451819	4694	74800	1219525	3450	3375	84180	1487225	5550	41475	56	11	3	5	237	3	
2451754	1529	2451836	2189	74800	1219525	3450	3375	84180	1487225	5550	41475	17	19	1	2	223	3	
2451754	1529	2451769	151	74800	1219525	3450	3375	84180	1487225	5550	41475	5	20	6	4	72	3	
2451754	1529	2451763	10793	74800	1219525	3450	3375	84180	1487225	5550	41475	44	2	18	2	276	3	
2451754	1529	2451863	12782	74800	1219525	3450	3375	84180	1487225	5550	41475	8	26	3	2	94	3	
2451754	1529	2451868	15769	74800	1219525	3450	3375	84180	1487225	5550	41475	23	2	9	4	99	3	
2451754	1529	2451833	5549	74800	1219525	3450	3375	84180	1487225	5550	41475	5	23	1	5	69	3	
2451754	1529	2451858	12475	74800	1219525	3450	3375	84180	1487225	5550	41475	5	23	8	4	23	3	
2451754	1529	2451785	5324	74800	1219525	3450	3375	84180	1487225	5550	41475	47	23	20	1	299	3	
2451754	1529	2451839	15686	74800	1219525	3450	3375	84180	1487225	5550	41475	47	26	17	2	7	3	
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	---	---	-	-	---	-	

Data Contract: Schema

Schema

dbt **yaml**

Model

Source

```
models:
- name: orders
  description: A successful sale in the web shop
  columns:
    - name: order_id
      data_type: string
      description: Primary key of the order
    - name: billing_customer_id
      data_type: string
      description: Customer ID of the billing customer
    - name: shipment_customer_id
      data_type: string
      description: Customer ID of customer to ship the order to
    - name: sold_timestamp
      data_type: timestamp
      description: The timestamp of the final confirmation step in the web form.
    - name: total_amount
      data_type: bigint
```


Data Contract: Schema

Schema

dbt yaml

Model

Source

orders

A successful sale in the web shop

order_id	string	Primary key of the order
billing_customer_id	string	Customer ID of the billing customer
shipment_customer_id	string	Customer ID of customer to ship the order to
sold_timestamp	timestamp	The timestamp of the final confirmation step in the web form.
total_amount	bigint	The total order amount in the smallest unit of the currency (such as Eurocents)
total_currency	string	The ISO 4217 currency code.

line_items

The articles in the order. One row per article.

order_id	string	Primary key of the order
sku	string	Stock keeping unit

Data Contract: Quality

Quality

SodaCL

yaml

checks for orders:

- freshness(order_timestamp) < 24h
- row_count > 500000
- duplicate_count(order_id) = 0


checks for line_items:

- row_count > 500000

Drive Data Product Development by Business Case

"As data scientist in team Recommendations, I want to get all webshop orders from the past 5 years, so that we know which articles (categories) are often bought together by one customer.

With that, we can create an ML model to generate individual recommendations for our customers in our marketing emails based on their recent activities in the online shop."



Ah, good to know.
That is OK for us.

Sophia, Data Scientist



"Sorry, we cannot give you **all** orders, but only orders when the customer expressed their consent for analytical use. That are around 80% of all orders.

John, Product Owner

Data Contract: Terms

Terms

Terms and conditions of the data contract

Usage

The data can be used for analytical and data science use cases, as the customer has expressed their consent.

Limitations

As the dataset is filtered, these data set cannot be used to aggregate financial KPIs.

Not suited for real-time use cases.

Billing

\$1000 per month

Notice Period

3 months

Data Contract: Info

Info

Information about the data contract

Title

Web Orders With Consent V1

Version

1.0.0

Description

All orders made through the web channel.

Filtered for orders where customers have expressed consent for analytical use.

Owner

John Doe

Contact

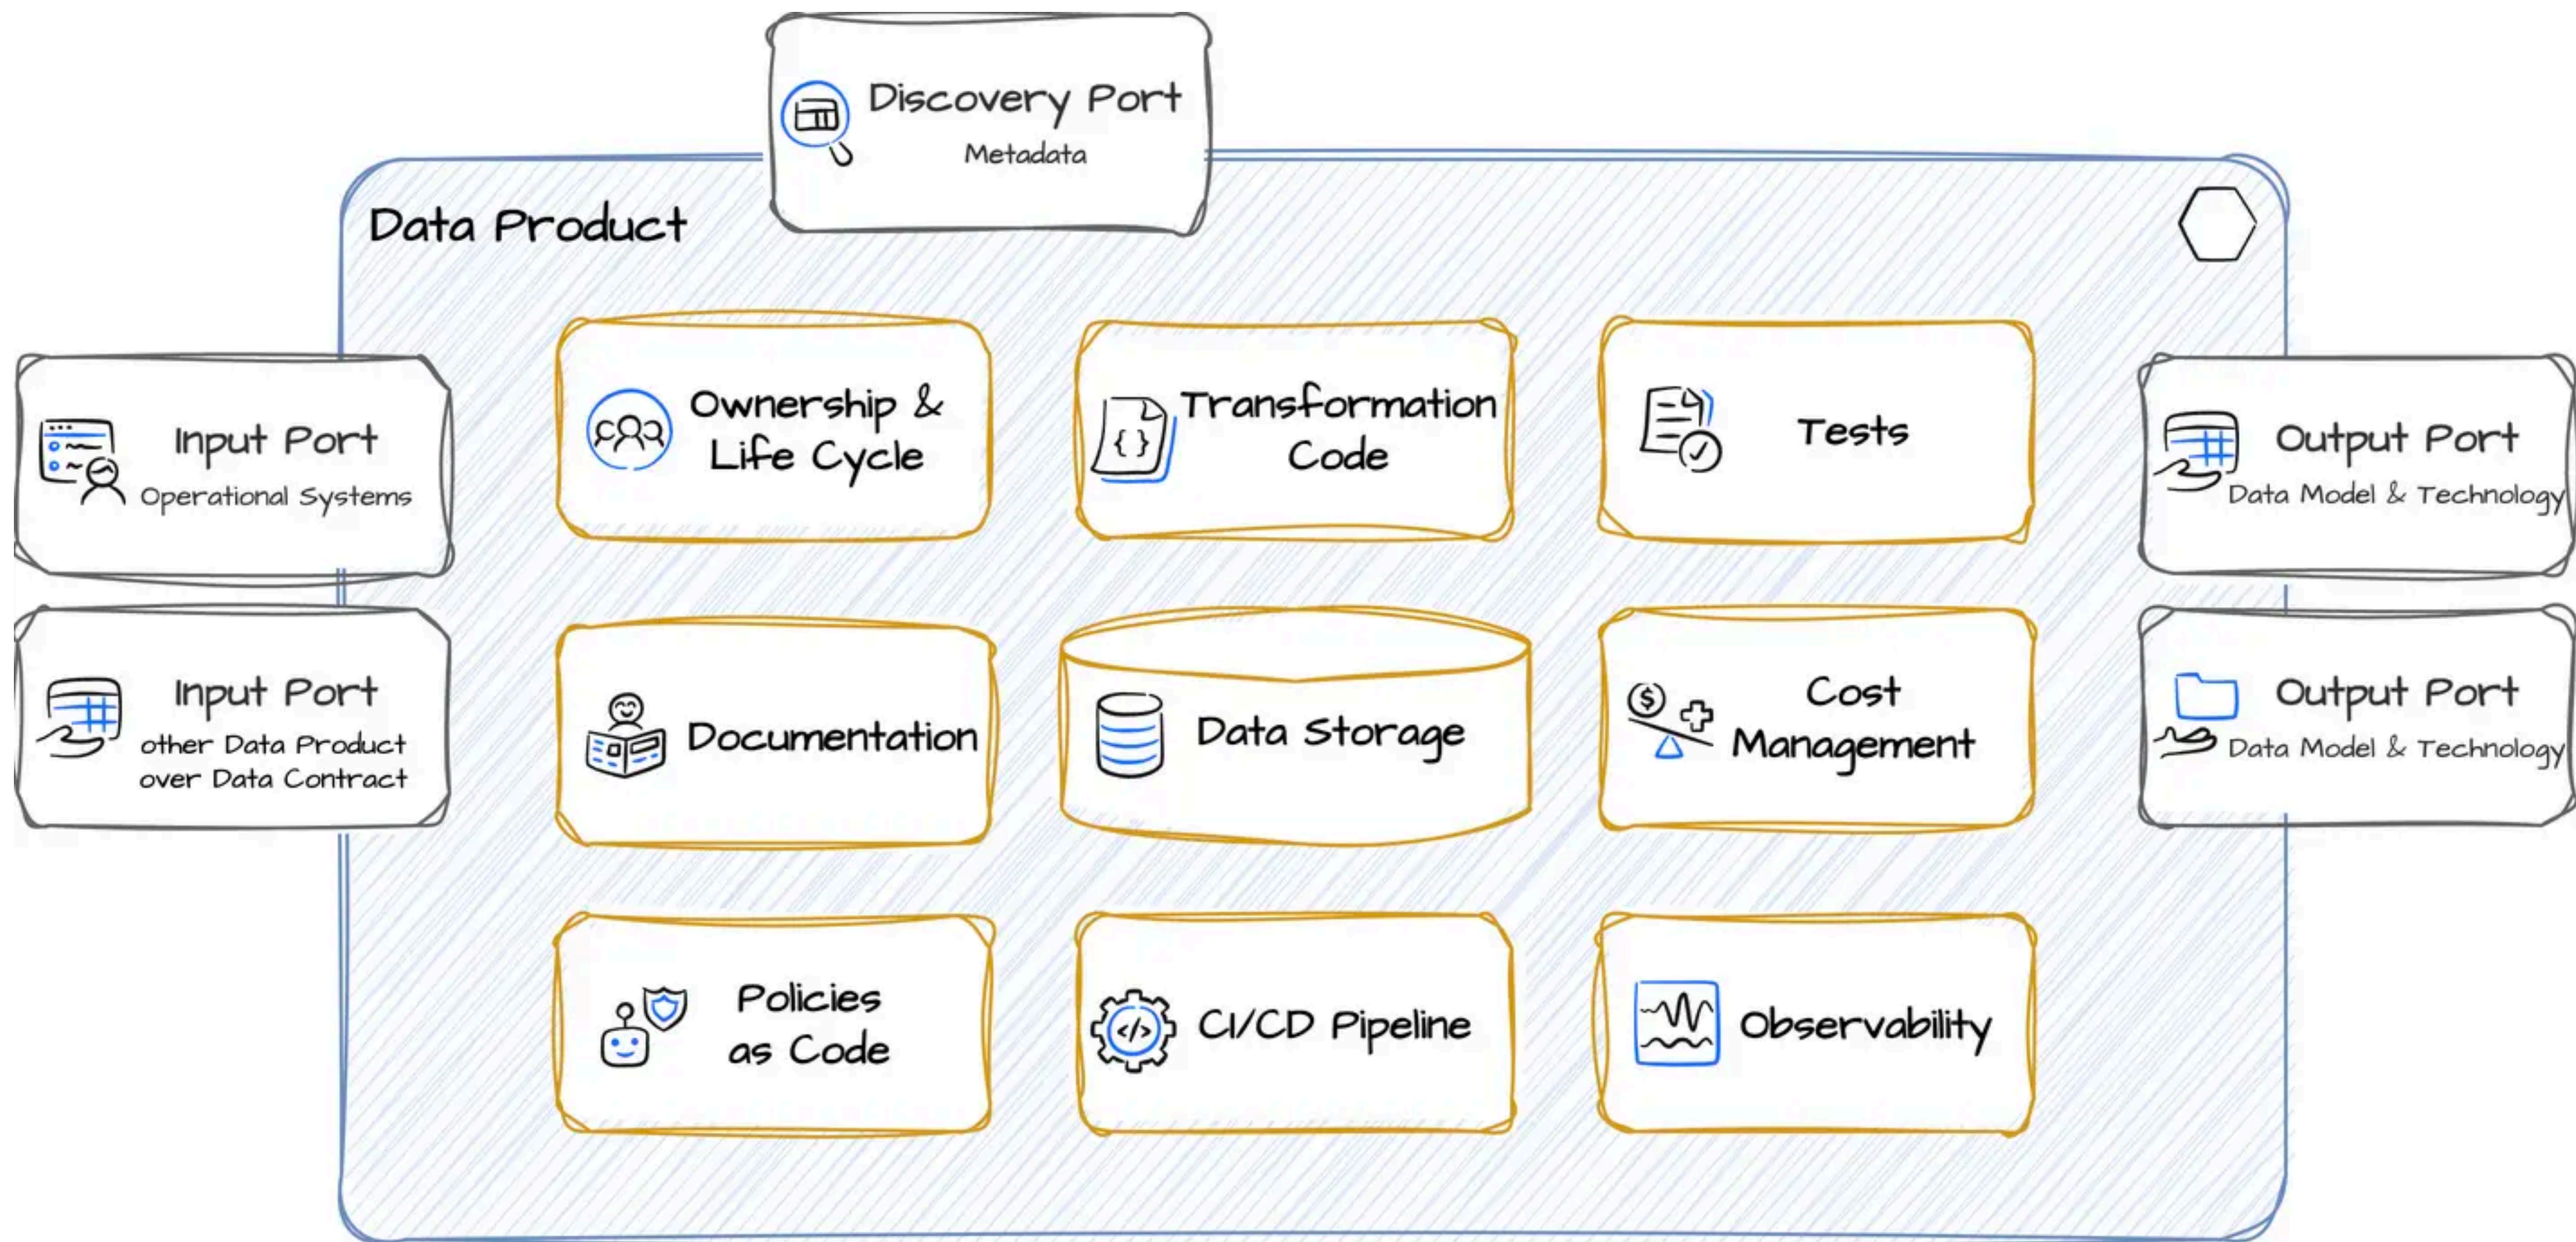
<https://teams.example.com/datacontracts/web-orders-with-consent-v1>

Data Contract as YAML

```
dataContractSpecification: 0.9.0
id: web-orders-with-consent-v1
info:
  title: Web Orders With Consent V1
  version: 1.0.0
  description: "All orders made through the web channel.\r\nFiltered for orders where customers have expressed consent for analytical use."
  owner: John Doe
  contact:
    url: https://teams.example.com/datacontracts/web-orders-with-consent-v1
terms:
  usage: "The data can be used for analytical and data science use cases, as the customer has expressed their consent."
  limitations: "As the dataset is filtered, these data set cannot be used to aggregate financial KPIs.\r\nNot suited for real-time use cases."
  billing: $1000 per month
  noticePeriod: P3M
schema:
  type: dbt
  specification: |
    models:
      - name: orders
        description: A successful sale in the web shop
        columns:
          - name: order_id
            data_type: string
            description: Primary key of the order
          - name: billing_customer_id
            data_type: string
```

datacontract.com

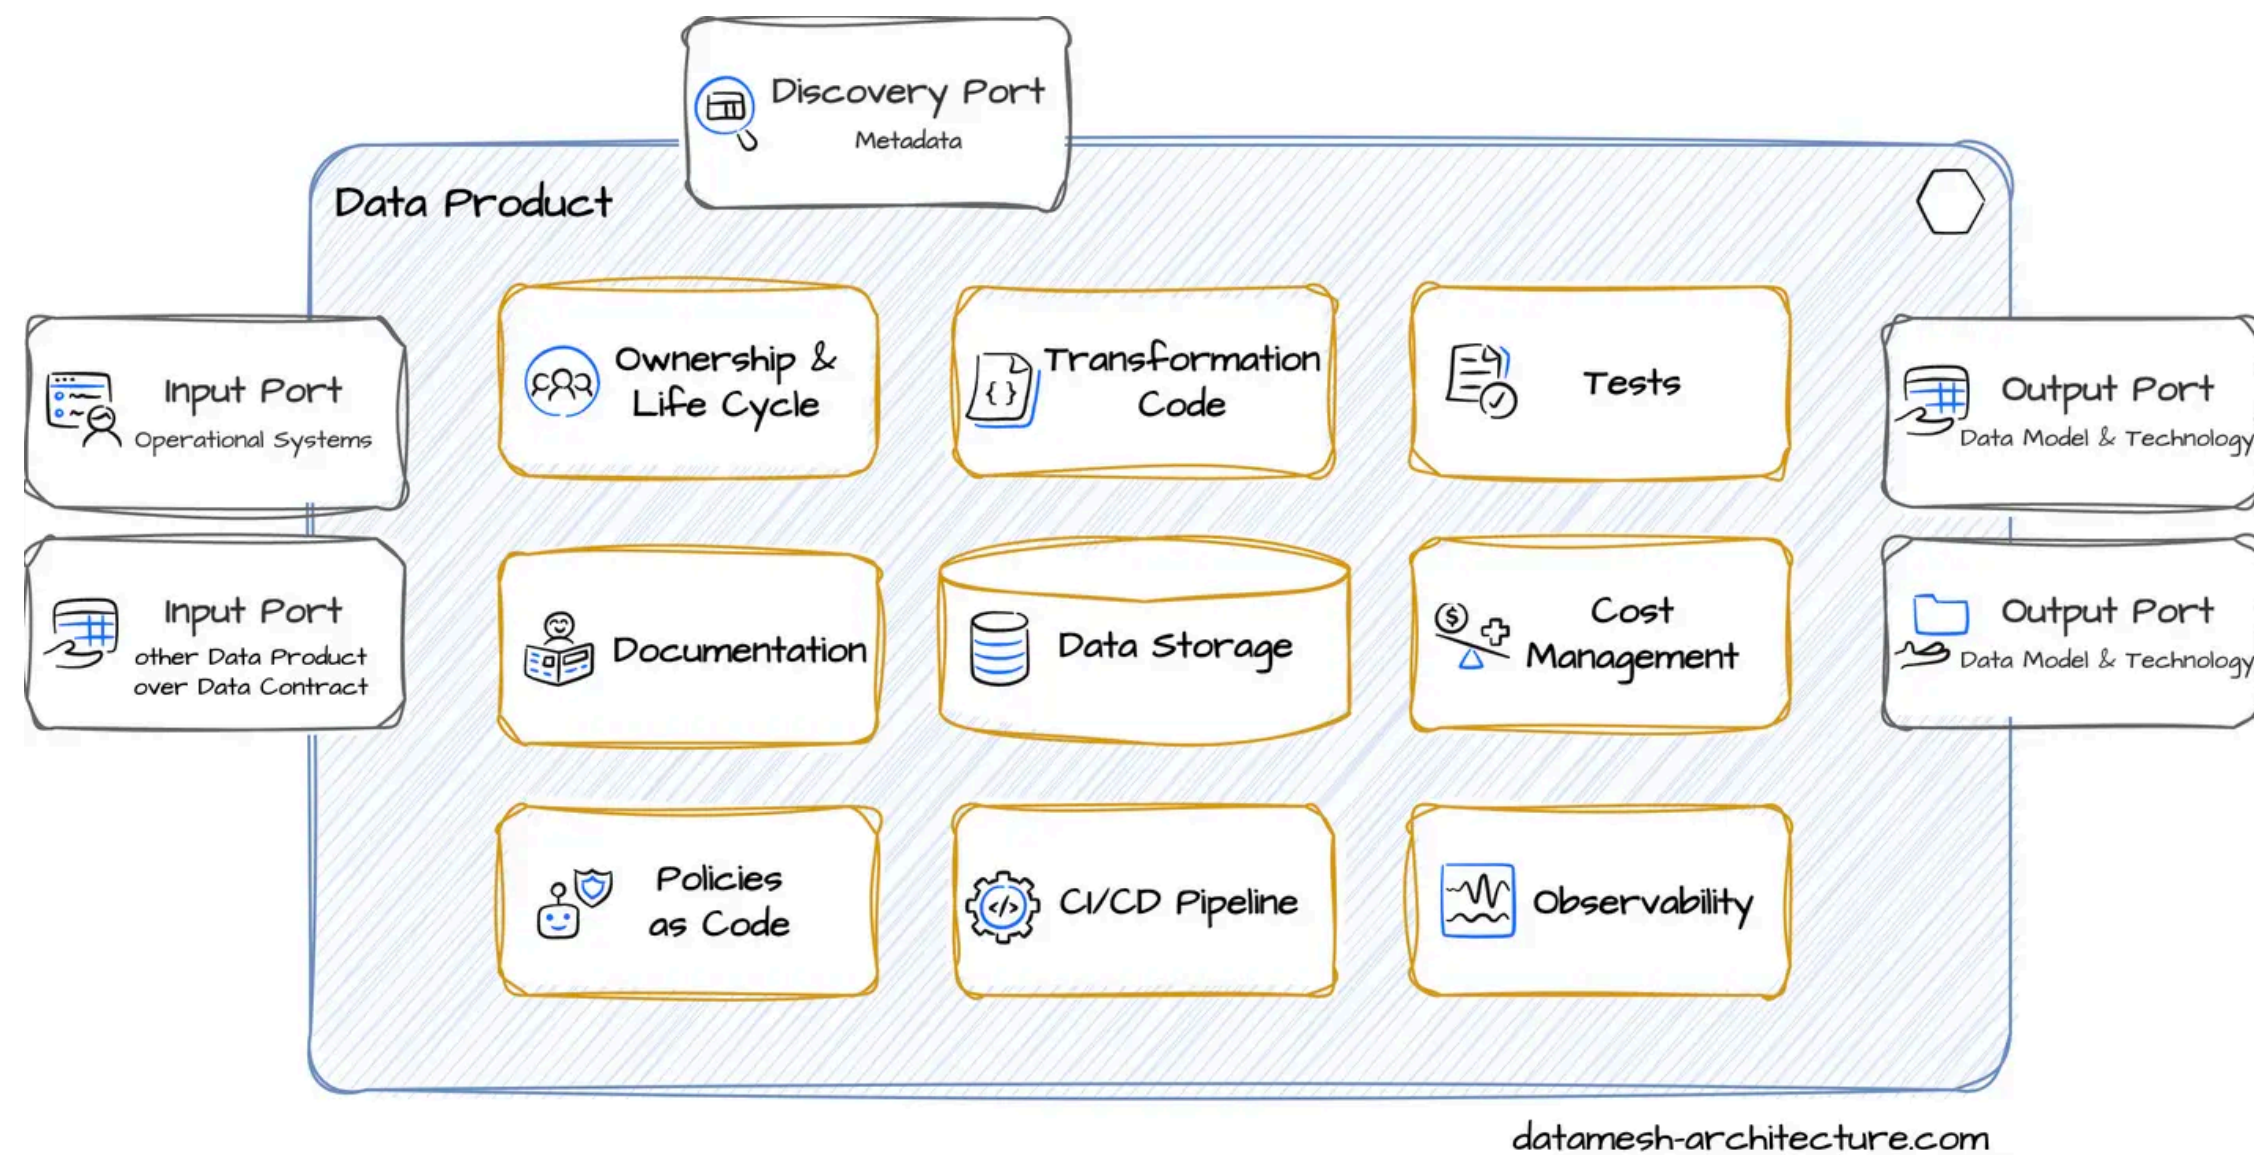
How to implement a Data Product?



Implementation depends on Tech Stack

Stack	Storage	Query Engine	Framework
AWS	S3	Athena (SQL)	Lambda / Step Functions
Google Cloud	BigQuery	BigQuery (SQL)	dbt
Azure (MS Fabric)	OneLake	Spark	Fabric notebook
Databricks	Deltalake	Spark	Databricks Asset Bundles
Snowflake	Snowflake	Snowflake (SQL)	dbt
On-Premise	S3 compliant	Presto (SQL)	dbt
Java	S3 compliant	Java	Spring Cloud Data Flow
DuckDB	S3 compliant	DuckDB (SQL)	dbt

1 Git Project per Data Product



===



1 Repository

Input Port: Kafka -> S3



confluent.cloud/environments/env-y732k/clusters/lkc-j557mm/connectors/sinks/S3_SINKConnector_0

CONFLUENT Stream Catalog LEARN

ENVIRONMENTS > DATAMESHEXAMPLE-FULFILLMENT-PROD > CLUSTER_FULFILLMENT

Cluster overview
Topics
Data integration
Clients
Connectors
API keys
Stream lineage
ksqlDB

Edit "S3_SINKConnector_0"

Overview Settings Pause Delete

1. Edit connection 2. Test and verify

Which topics do you want to get data from?

topics
inventory x

+ Create topic

How should we connect to your data?

Connector class S3_SINK

Name S3_SINKConnector_0

Input messages
Input Kafka record value format* JSON

Kafka Cluster credentials
Kafka Cluster Authentication mode KAFKA_API_KEY

Kafka API Key
Kafka API Secret

Connector Summary

Name S3_SINKConnector_0

Connector Class S3_SINK

Topics inventory

Max Tasks 1

input.data.format JSON

kafka.auth.mode KAFKA_API_KEY

s3.bucket.name datameshexample-fulfillment-inv...

output.data.format JSON

time.interval DAILY

flush.size 1000

Cancel Next

console.aws.amazon.com/s3/buckets/datameshexample-fulfillment-inventory?region=eu-central-1&pr...

aws Services Search for services, features, blogs, docs, and more [Option+S]

Amazon S3 Buckets datameshexample-fulfillment-inventory topics/ inventory/ year=2022/ month=04/ day=20/ hour=00/ Copy S3 URI

hour=00/

Objects Properties

To enable sorting in the table below, use the search to reduce the size of the list to 999 objects or fewer.

Objects (999+)

Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

Copy S3 URI Copy URL Download Open Delete Actions Create folder Upload

Find objects by prefix

	Name	Type	Last modified	Size	Storage class
<input type="checkbox"/>	inventory+0+0001327496.json	json	April 20, 2022, 15:50:35 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001328496.json	json	April 20, 2022, 15:50:36 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001329496.json	json	April 20, 2022, 15:50:38 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001330496.json	json	April 20, 2022, 15:50:38 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001331496.json	json	April 20, 2022, 15:50:38 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001332496.json	json	April 20, 2022, 15:50:40 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001333496.json	json	April 20, 2022, 15:50:40 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001334496.json	json	April 20, 2022, 15:50:40 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001335496.json	json	April 20, 2022, 15:50:42 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001336496.json	json	April 20, 2022, 15:50:42 (UTC+02:00)	94.9 KB	Standard
<input type="checkbox"/>	inventory+0+0001337496.json	json	April 20, 2022, 15:50:42 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001338496.json	json	April 20, 2022, 15:50:44 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001339496.json	json	April 20, 2022, 15:50:44 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001340496.json	json	April 20, 2022, 15:50:44 (UTC+02:00)	95.0 KB	Standard
<input type="checkbox"/>	inventory+0+0001341496.json	json	April 20, 2022, 15:50:46 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001342496.json	json	April 20, 2022, 15:50:46 (UTC+02:00)	95.1 KB	Standard
<input type="checkbox"/>	inventory+0+0001343496.json	json	April 20, 2022, 15:50:48 (UTC+02:00)	95.0 KB	Standard

Feedback Looking for language selection? Find it in the new Unified Settings

© 2022, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Stream operational data into analytical data platform to preserve historic states and enable efficient JOIN and aggregate operations.

5

EXPLORER

...

5

FULFILLMENT_DBT

> analyses

> dbt_packages

> logs

> macros

> models

> analytics

> data_products

> staging/inventory

! src_inventory.yml

stg_inventory__inven...

! stg_inventory.yml

> seeds

> snapshots

> target

> tests

inventory_history_test...

.gitignore

! dbt_project.yml

i README.md

> OUTLINE

> TIMELINE

stg_inventory__inventory_history.sql

inventory_history.sql

inventory_latest.sql

inventory__avai

...

models > staging > inventory > stg_inventory__inventory_history.sql

23

24 deduplicated as (
25 SELECT *
26 FROM (
27 SELECT *,
28 ROW_NUMBER() OVER (PARTITION BY event_id ORDER BY event_time DESC) row_number
29 FROM parsed)
30 WHERE row_number = 1
31),
32 final as (
33 select
34 available,
35 location,
36 sku,
37 updated_at
38 from deduplicated
39)
40 select * from final

Terminal (^`)

PROBLEMS

OUTPUT

DEBUG CONSOLE

TERMINAL


> zsh + ▾ □ □ ✕ ^ ×

jochen@Jochens-MacBook-Pro-2 fulfillment_dbt % dbt run

⊗ 0 ⚠ 0 AWS

Ln 31, Col 3 Spaces: 4 UTF-8 LF SQL 🔍 🔔

Tests: Test Engine to ensure Quality

 Soda

Checks

Datasets


Agreements

Incidents

Get Started

Ask SodaGPT

?



JR

Checks

Filters

LAST SCANNED DATE

LAST RESULT

NAME

OWNER

DATASET

paxstats (paxstats2)

Dataset: paxstats (paxstats2) ✕

SORT BY Last Result ▼ DESC ▼

CHECK	VALUE	DATASET	LAST EVALUATED	INCIDENT	
❗ Verify 2-digit IATA	3	paxstats paxstats2/paxstats	about 17 hours ago	-	⋮
❗ Validate terminal ID	27	paxstats paxstats2/paxstats	about 17 hours ago	4 incidents	⋮
⚠ Abnormally large PAX count	659,837	paxstats paxstats2/paxstats	about 17 hours ago	-	⋮
✅ Schema Check		paxstats paxstats2/paxstats	about 17 hours ago	-	⋮
✅ anomaly score for row_count <...	15,007	paxstats paxstats2/paxstats	about 17 hours ago	-	⋮

Deployment: Terraform Module

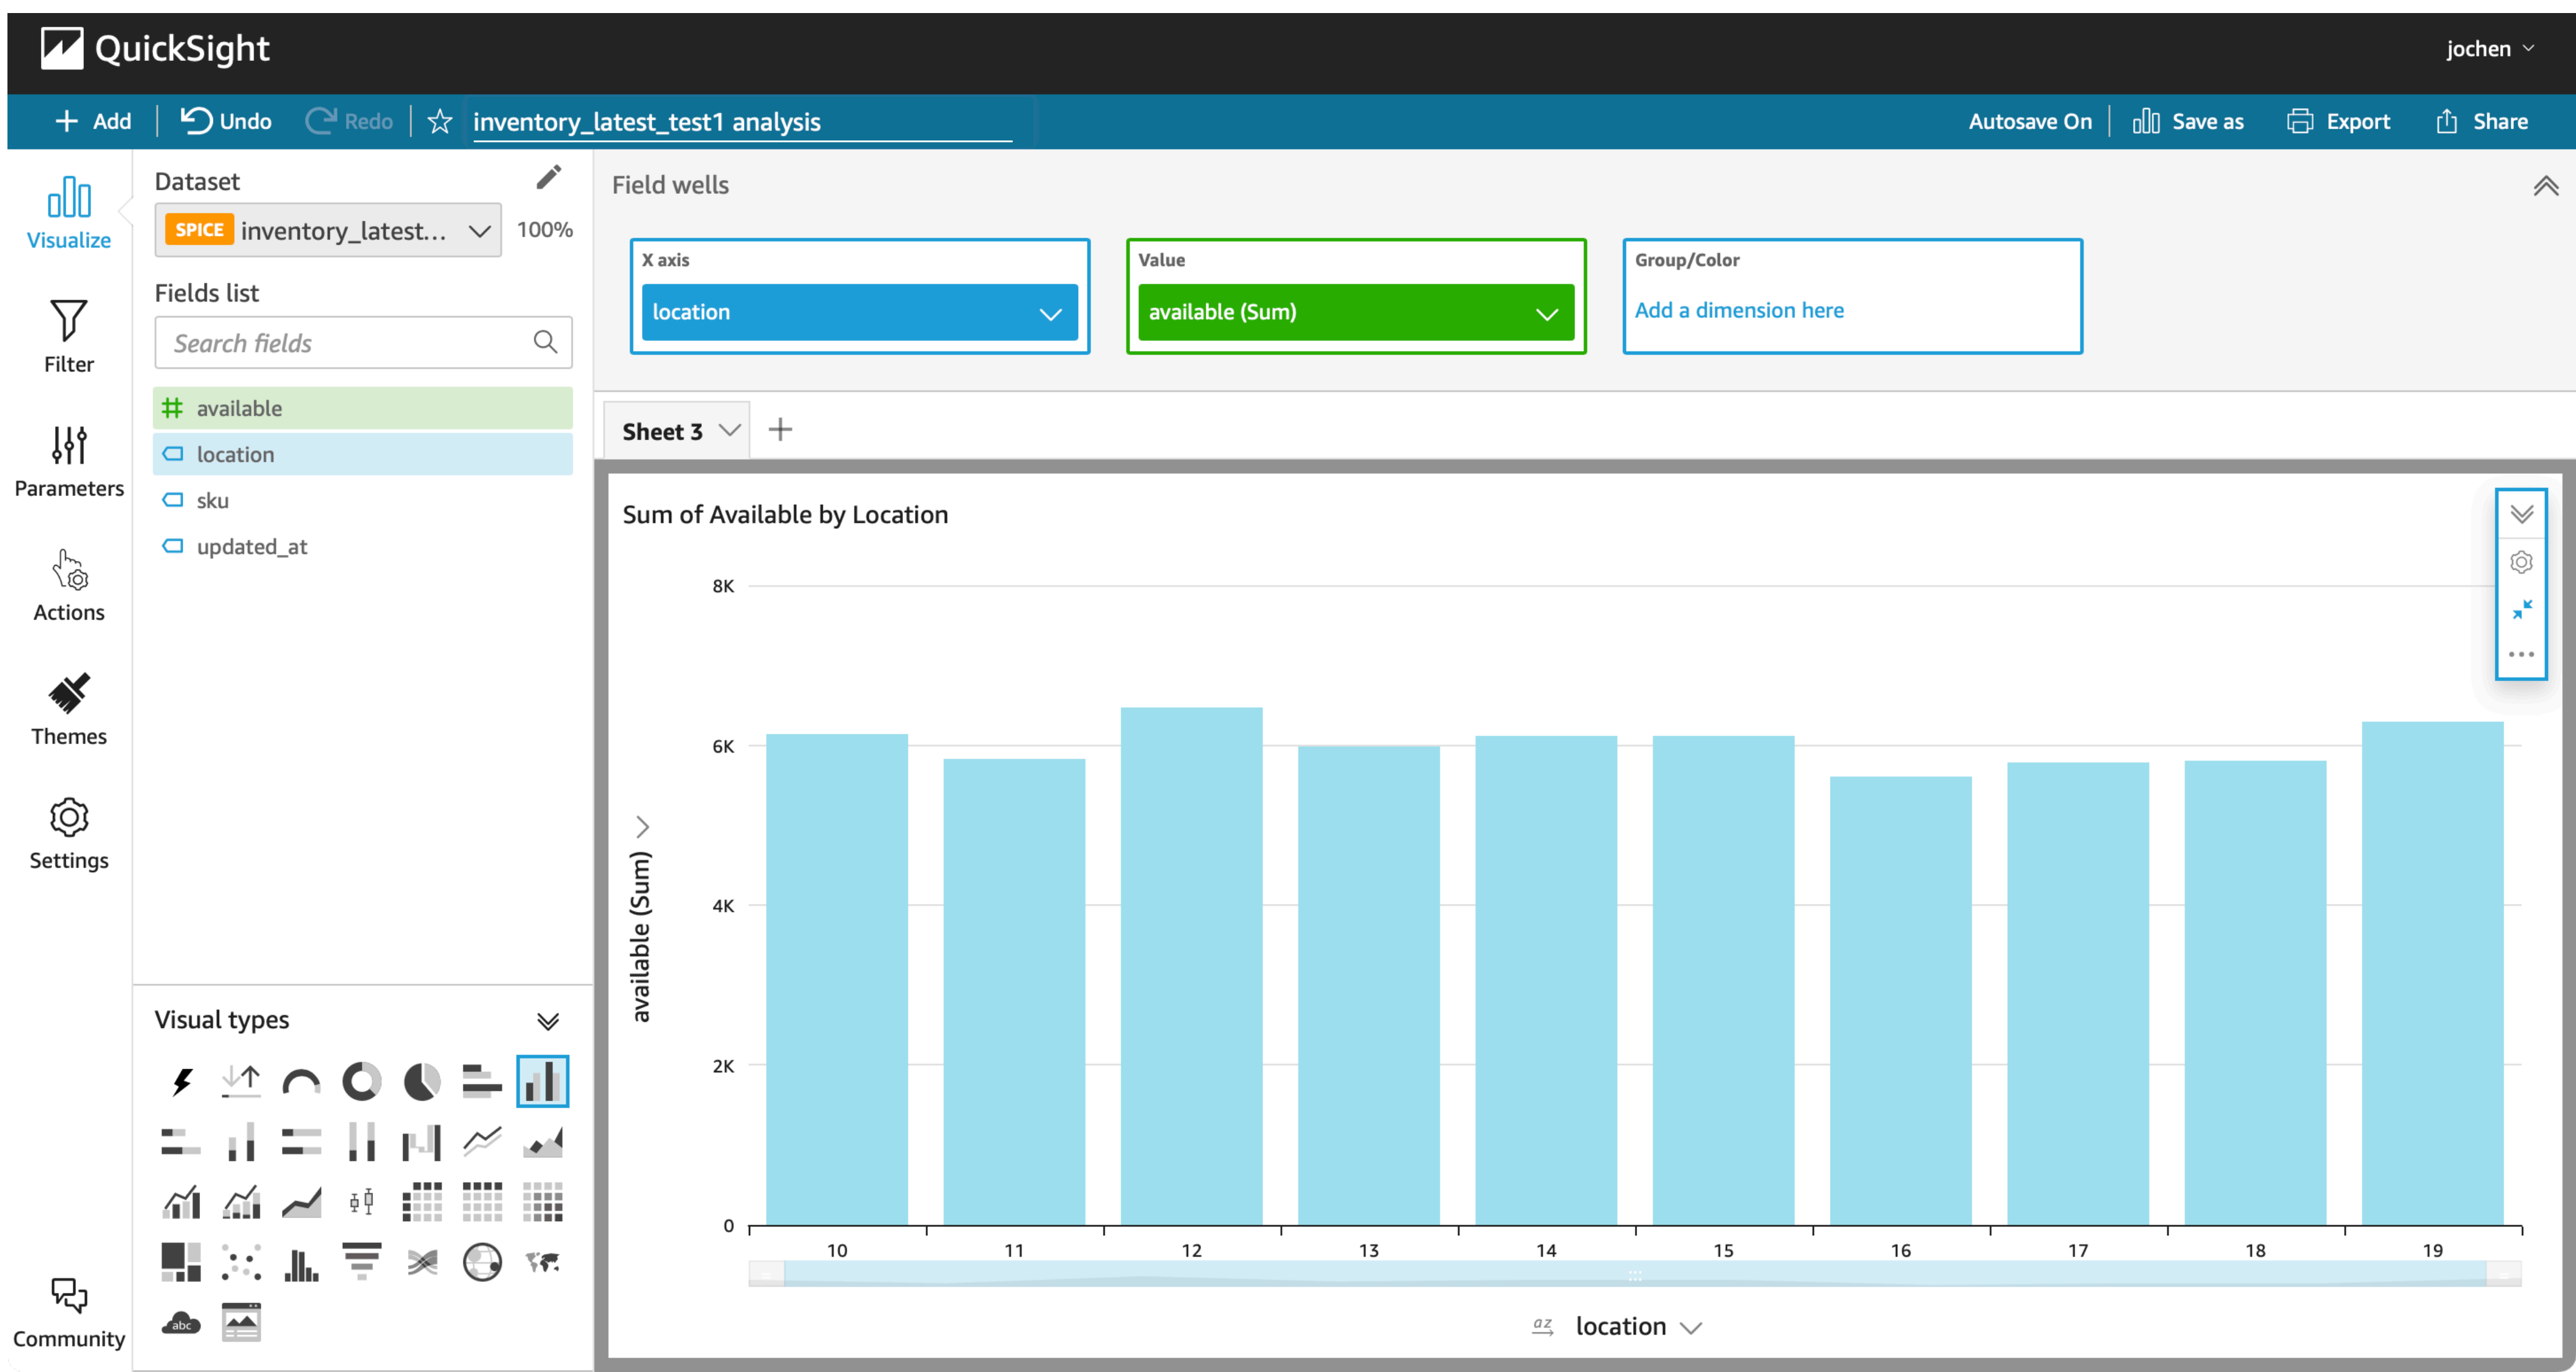
Configured and
deployed
by domain team

```
module "orders_data_product" {  
  source = "git@github.com:datamesh-architecture/terraform-dataproduct-aws-athena.git"  
  
  domain    = "checkout"  
  name      = "orders"  
  schedule  = "0 0 * * ? *" # Run at 00:00 am (UTC) every day  
  
  input = [{  
    source = "s3://example-bucket/orders"  
  }, {  
    source = "s3://example-bucket/consents"  
  }]  
  
  transform = {  
    query = "sql/orders-web-sales-filtered.sql"  
  }  
  
  output = {  
    format = "PARQUET"  
    schema = "schema/orders-web-sales-filtered.schema.json"  
  }  
}
```

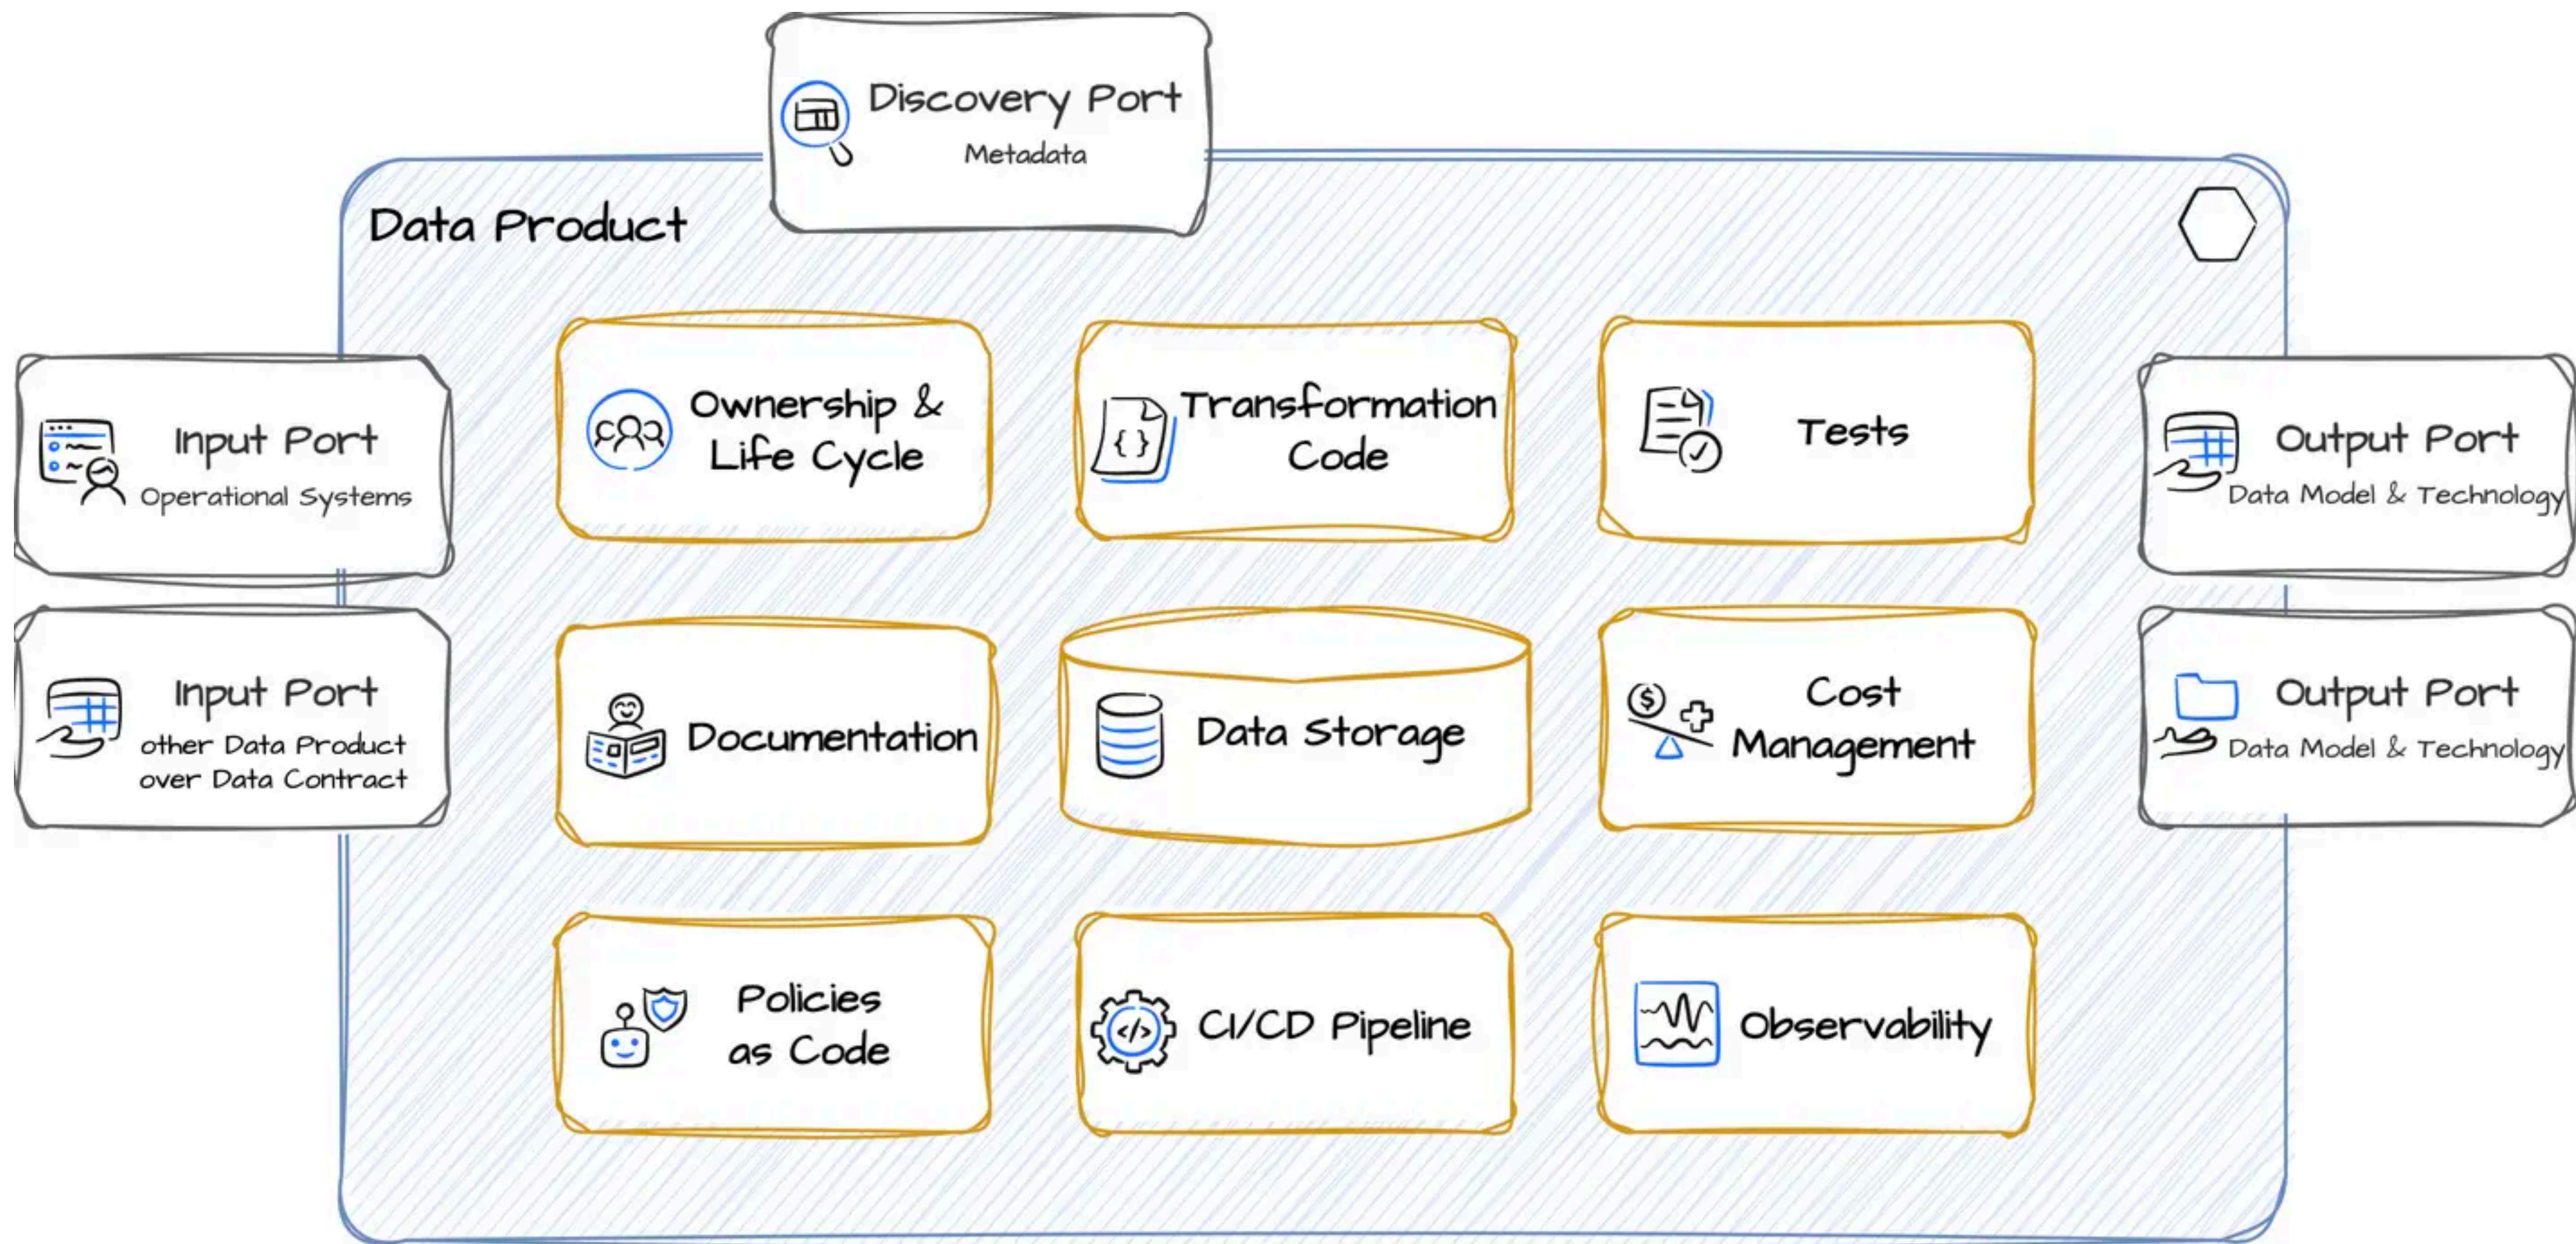
Terraform module
provided
by data platform
team

Example:
<https://github.com/datamesh-architecture/terraform-aws-dataproduct-aws-athena>

Analytics Tooling



How to promote my data product?



Meta Data

dataProductSpecification: 0.0.1

id: urn:dataproduct:checkout:orders

info:

title: Orders

description: Successful customer orders in the webshop.

owner: checkout

status: active

archetype: source-aligned

maturity: managed

inputPorts:

– id: kafka-orders

name: Orders Kafka

description: Order Updates in Kafka

sourceSystemId: checkout-orders-service

type: Kafka Topic

Discovery: Find Data Products

Data Mesh Manager

 Data Products

 Data Contracts

 Global Policies

 Teams

ACME

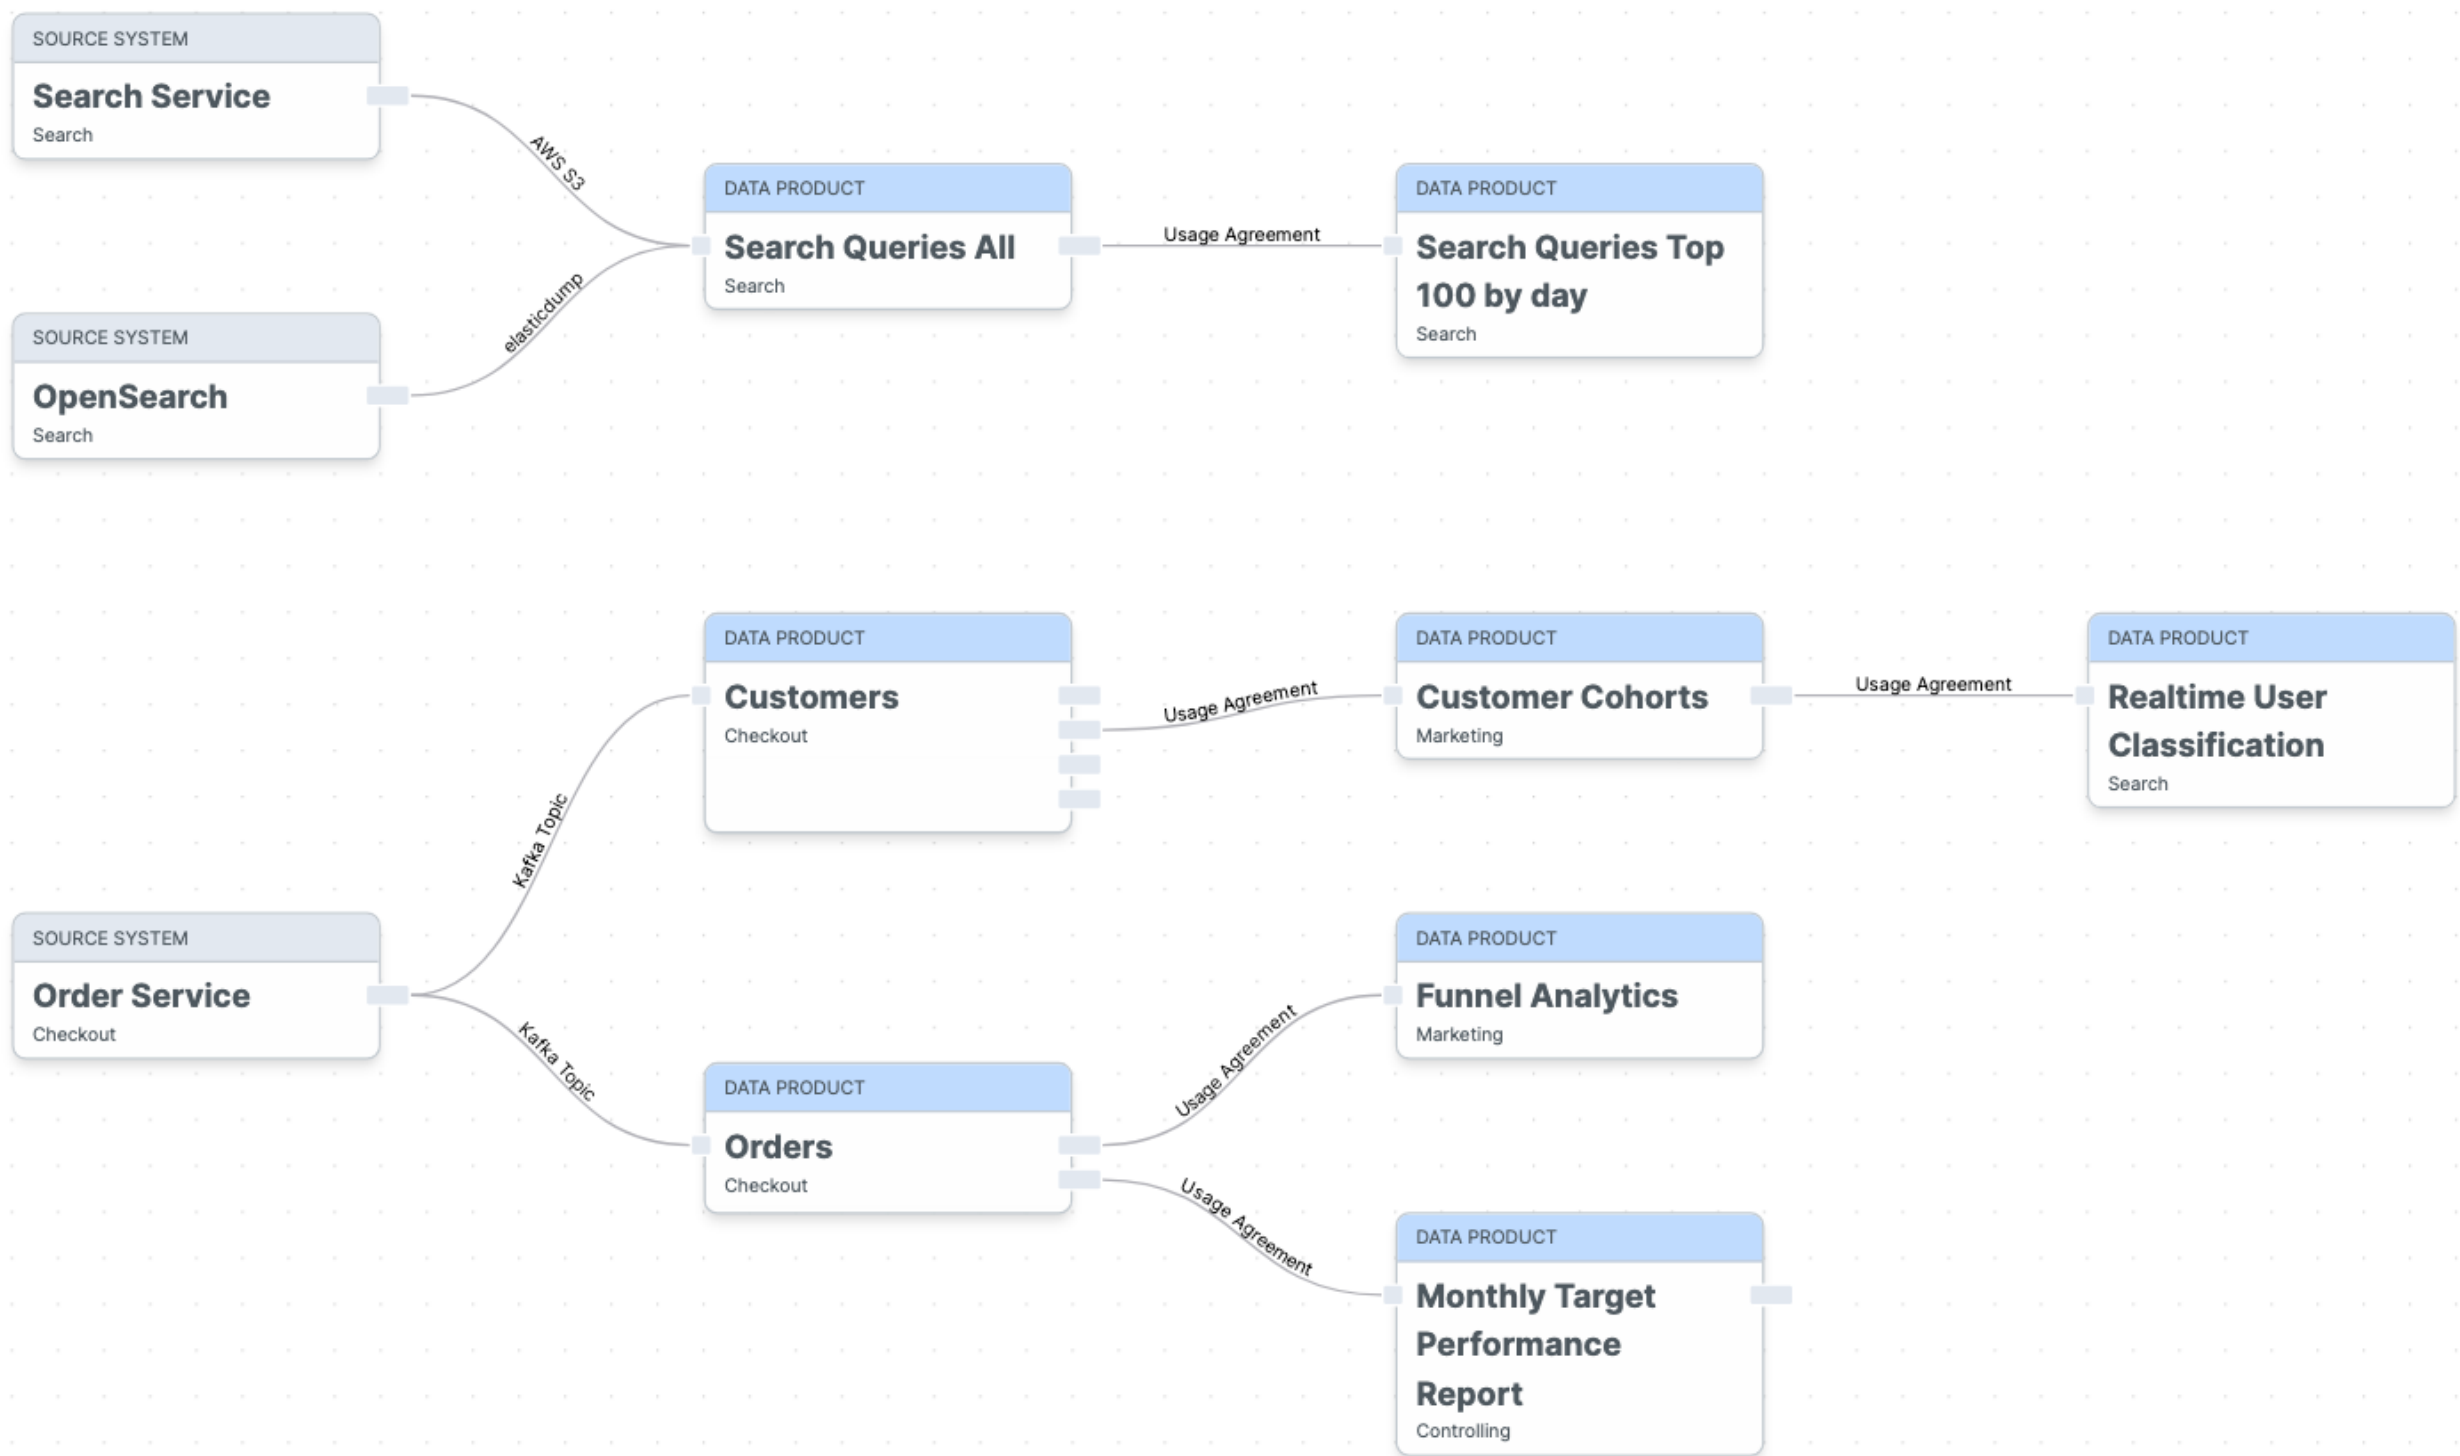


 > Data Mesh Map

Data Mesh Map

[View as list](#)

[+ Register Data Product](#)



Discovery: Find Data Products

Data Mesh Manager

 Data Products

 Data Contracts

 Global Policies

 Teams

ACME



 > Data Products

Data Products

[View as map](#)

[+ Register Data Product](#)

 Search

Owner 

Status 

Sort 

Customers Checkout	s3_customers_history_npii_v1	Data Contract	
	s3_customers_history_pii_v1	Data Contract	
	snowflake_customers_latest_npii_v1	Data Contract	
	snowflake_customers_latest_pii_v1	Data Contract	1
Orders Checkout	snowflake_orders_npii_v2	Data Contract	1
	snowflake_orders_pii_v2	Data Contract	1
Article Profitability Analysis Controlling			
Monthly Target Performance Report Controlling	looker_target_performing_report		
Shelf Warmers Fulfillment	snowflake_fulfillment_shelf_warmers	Data Contract	

✓ 1 / 1 Policies

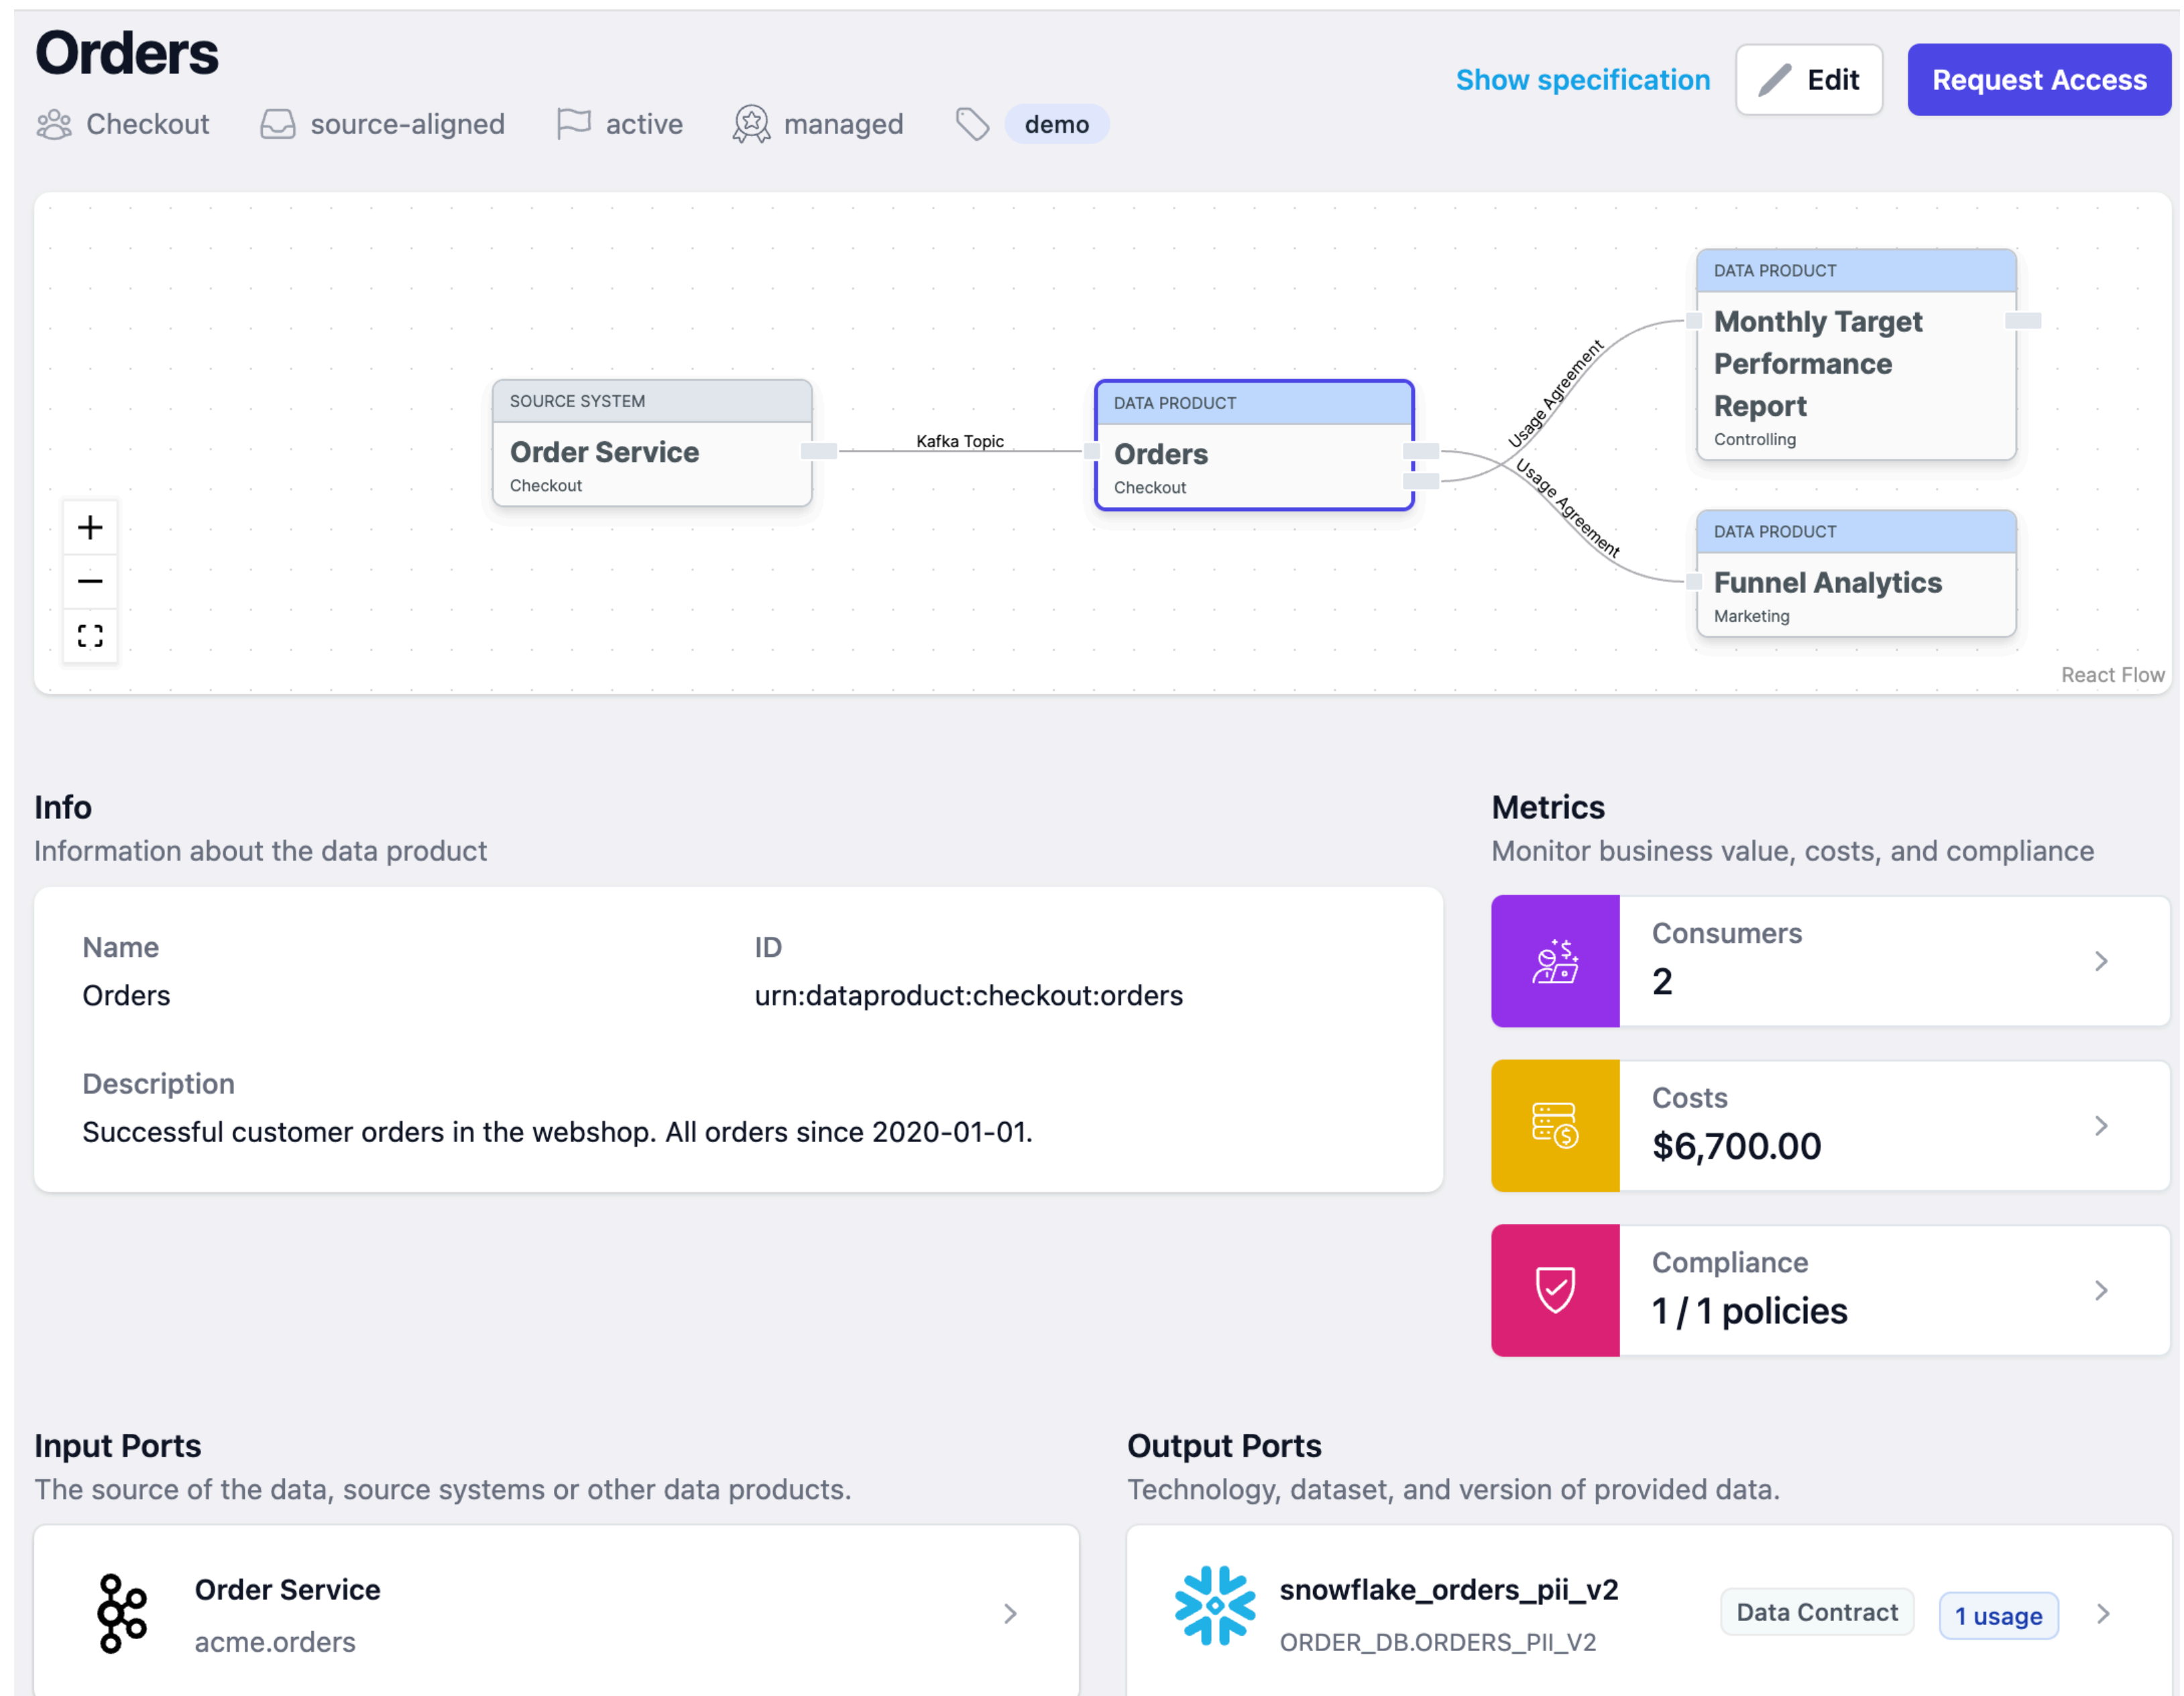
✓ 1 / 1 Policies

✓ 1 / 1 Policies

✓ 1 / 1 Policies

✓ 0 / 1 Policies

Discovery: Details



Data Usage Agreements

- One-to-one relationship between data provider and data consumer
- Automate permissions
- Data providers know their consumers
- Lifecycle: Can be terminated with notice period
- Data product evolution (e.g., for breaking changes)

Request Access

You are requesting access to the data product [Orders](#) on output port [web-orders-with-consent-v1](#) .
The system will create a data usage agreement for the team [Checkout](#) to approve.

Consumer

Required

Recommendations (Team Recommendations) ▾

Select your data product that wants to access and use the provided data.

Purpose

Required

Create an ML model to generate individual recommendations for our customers in our marketing emails based on their recent activities in the online shop.

Why do you want access and what do you want to do with the data?

Cancel

Customize

Request access

Data Product Controlling

Data Mesh Manager



Data Products



Data Contracts



Global Policies



Teams

ACME



[Home](#) > [Data Products](#) > [Search Queries All](#) > [Costs](#)

Cost Management

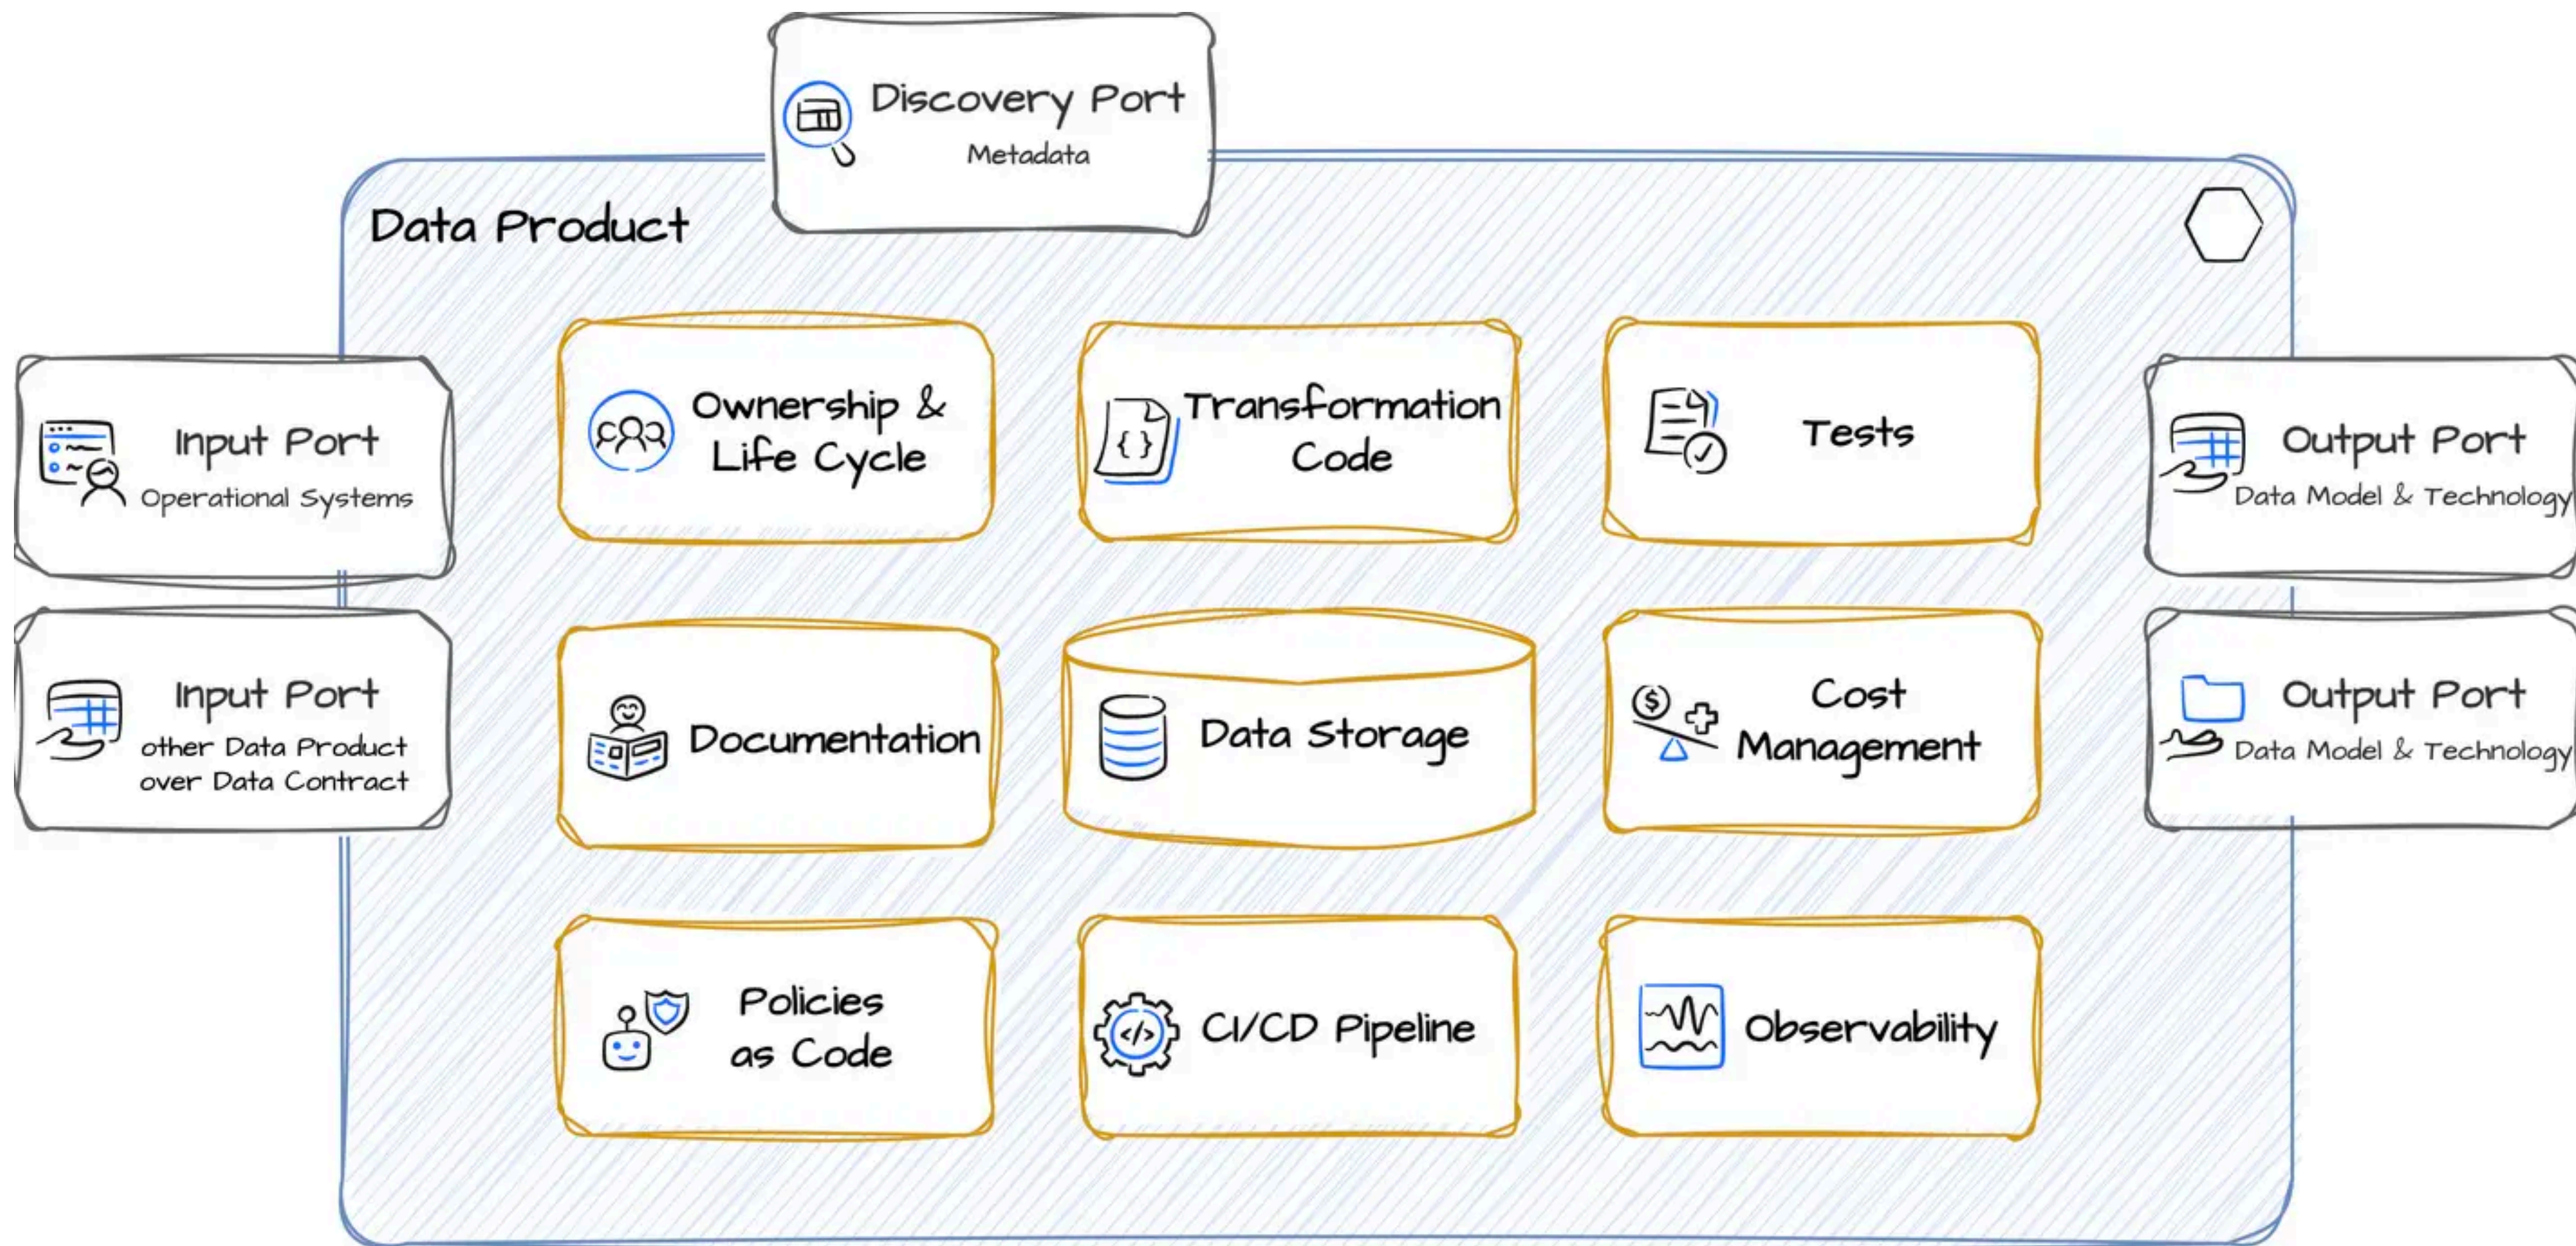
Expenses


Monthly expenses for running and maintaining the data product


Add expense

Name	Description	Category	Costs
Snowflake Compute	Average costs for Snowflake compute engine credits. Costs may vary depending on the autoscaling.	Data Platform	\$1,950.00
Maintenance	The cost of domain team's data developers.	Data Product Development	\$4,000.00
Snowflake Storage	Costs for Snowflake upfront storage capacity.	Data Platform	\$230.00
S3	S3 storage costs for storing search data	Data Platform	\$380.00
Total			\$6,560.00

We have implemented our first Data Product







Decentralized Data Architecture

Why?



Make qualified data-driven decisions in your domain

Use data to better understand your users and system behavior. Derive features from insights, qualify value, and fast iterations. Also qualified rejection of unnecessary tasks.

Do the right things, purpose, motivation



Build innovative services in your domain

Enhance your customer experience with data technologies, such as LLMs, visualizations, classifications, and ML models for predictions and recommendations.

Customer value through innovation




Provide data as business value for other domains

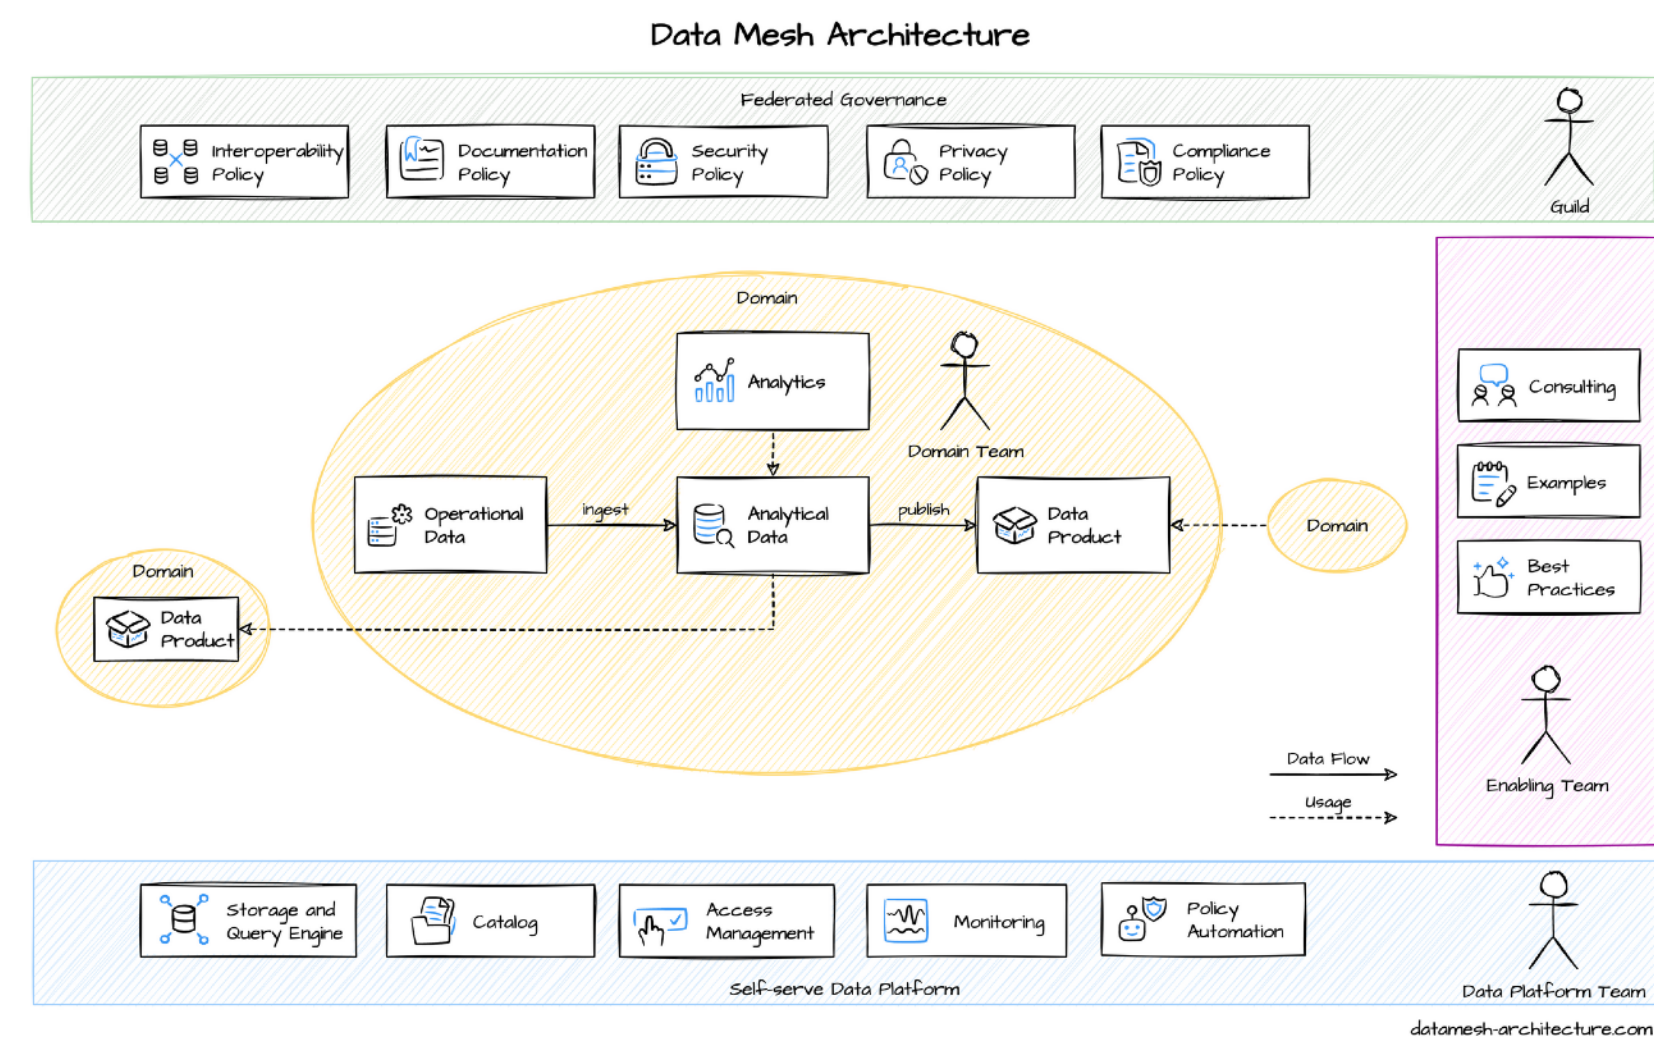
Domain data is valuable for other business units as reference data and to aggregate. Needs managed, explained, high-quality and easy accessible data as products.

Company success

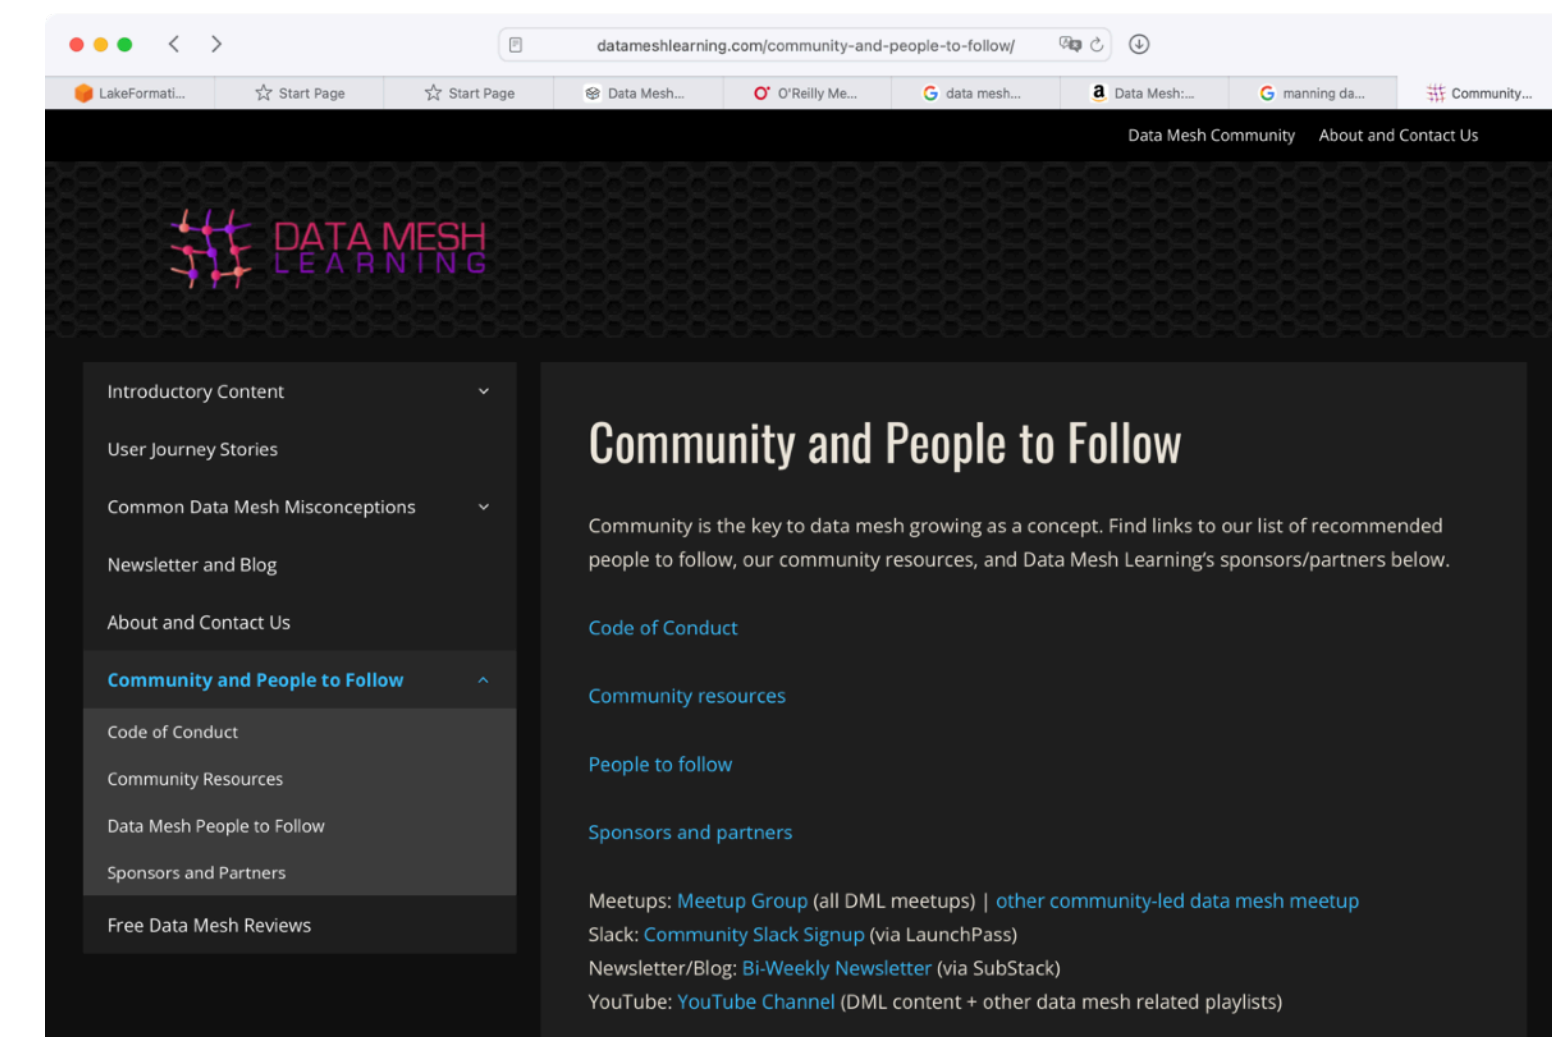
Learn more



oreilly.de/produkt/data-mesh



datamesh-architecture.com



datameshlearning.com
Podcast, newsletter, slack channel

Wir helfen, Data Mesh erfolgreich einzuführen

TRAINING



SOCREATORY

screatory.com/de/trainings/datamesh

TOOLING

Data Mesh Manager

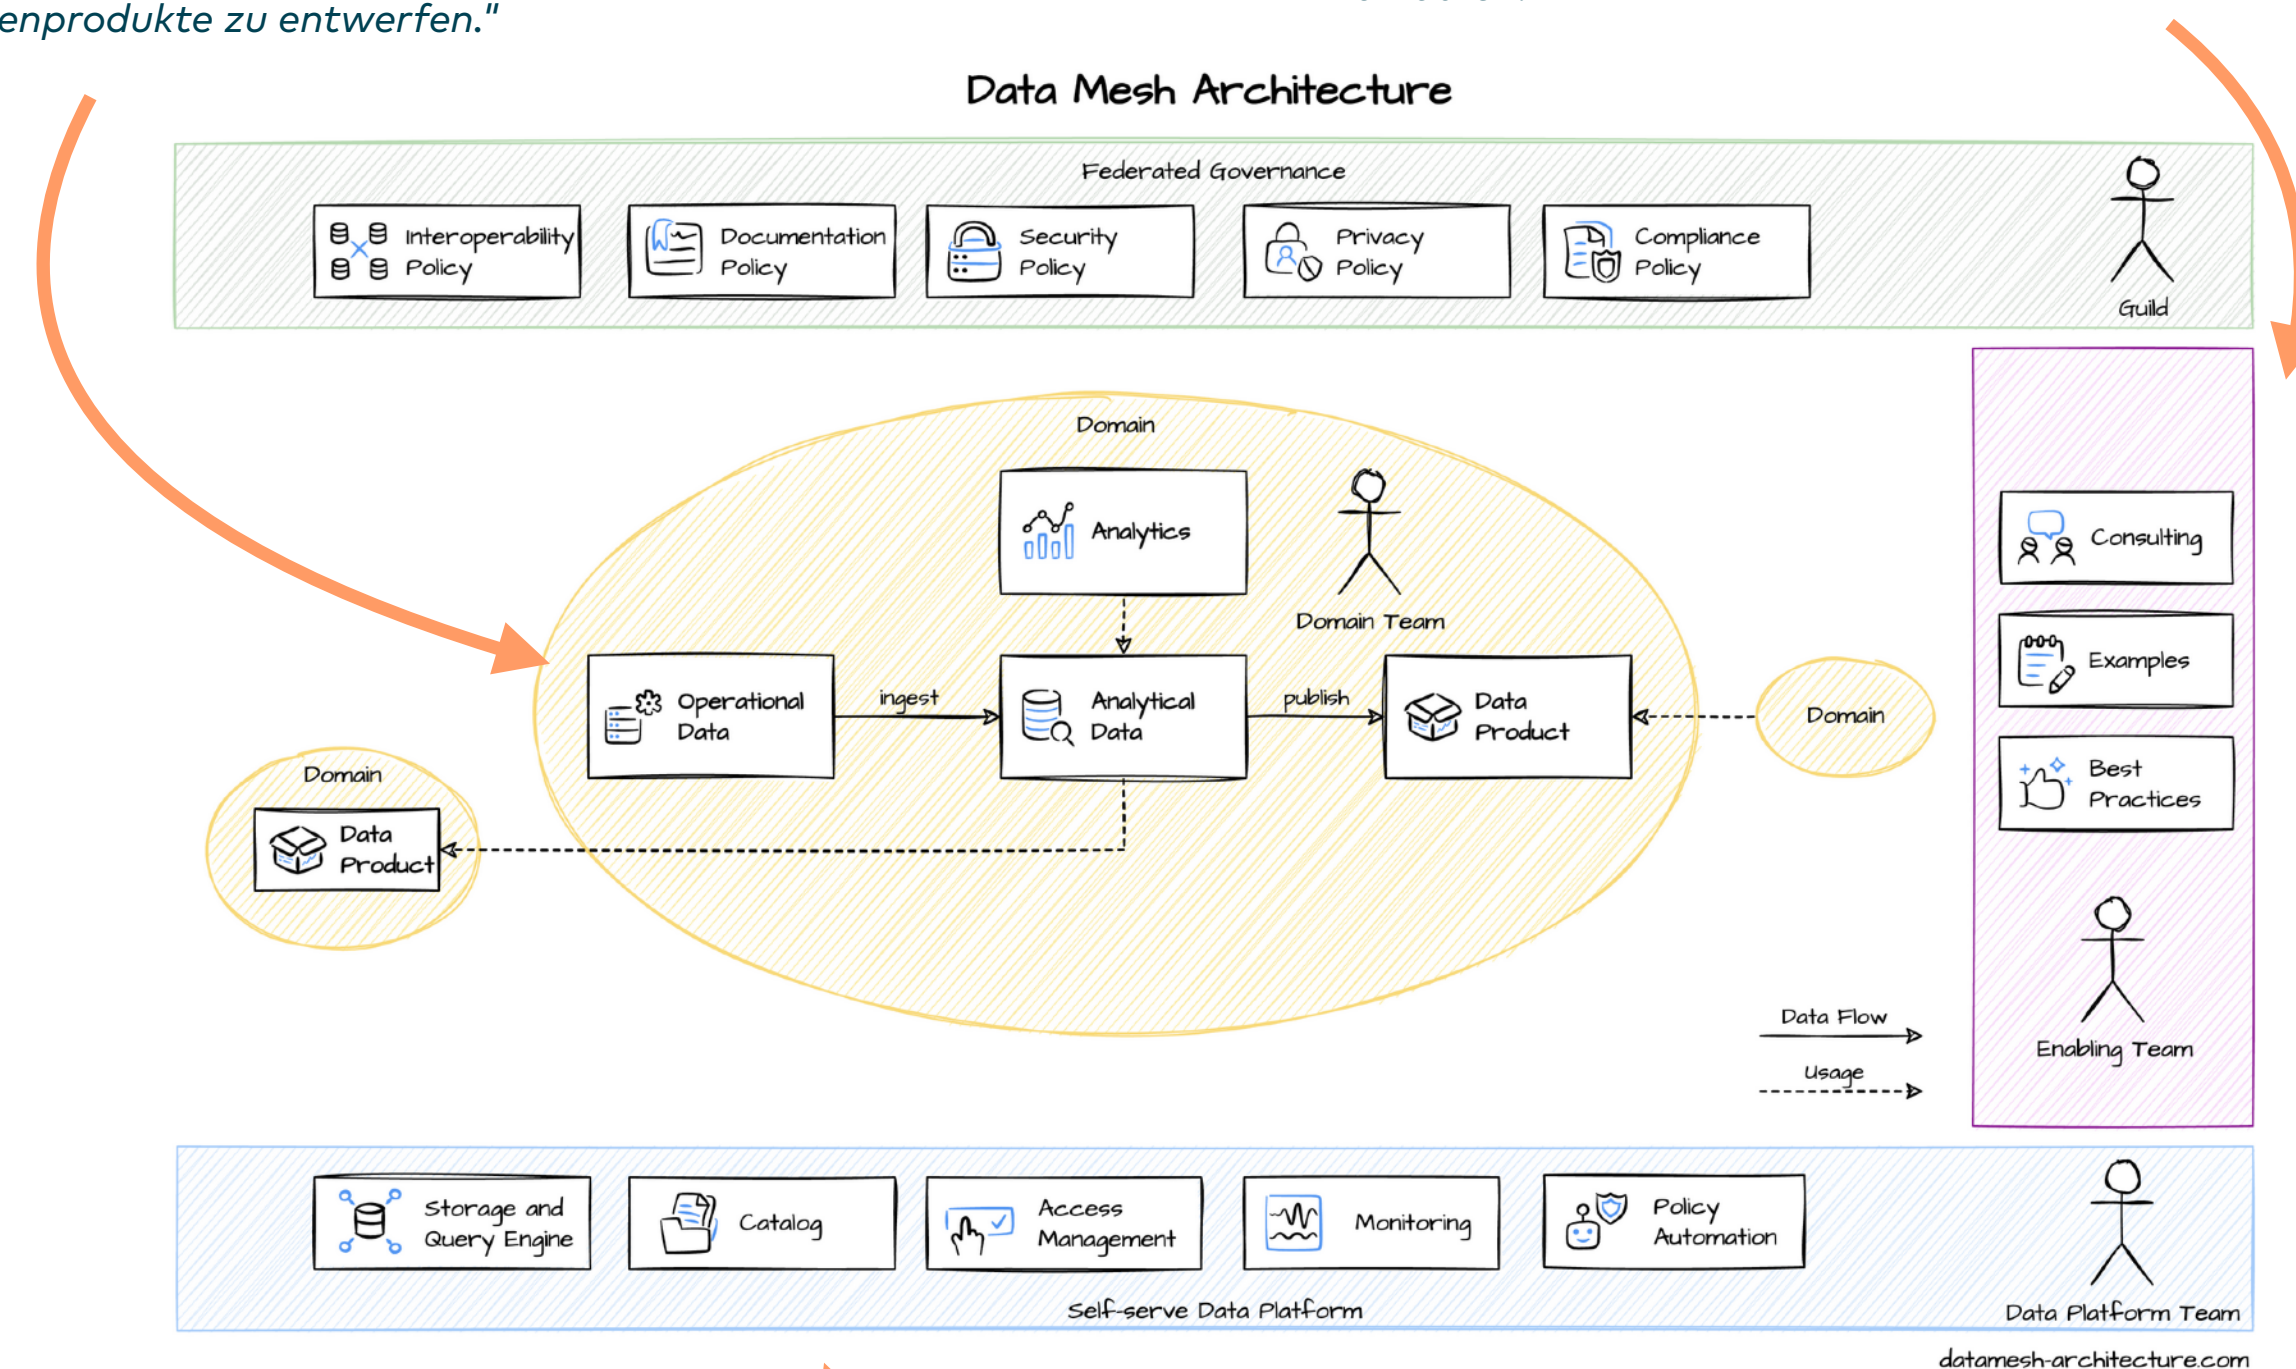
Data Contract Specification

Terraform Modules

CONSULTING

"Wir übernehmen **Umsetzungsverantwortung** und bieten **Trainings** für Entwickler:innen an, um mit entwicklungsnahe Werkzeugen wie dbt eigene Analysen durchzuführen und Datenprodukte zu entwerfen."

"Als Teil eines **Enabling Teams** helfen wir Data Mesh innerhalb eines Unternehmens aufzusetzen, zu promoten und es Engineers leicht zu machen, die Data Platform für ihre analytischen Daten zu nutzen."



"Wir helfen bei der **Auswahl passender Technologien und Produkte** für die Data Platform und setzen diese anhand von **Best Practices** auch gerne mit auf."

Data Mesh

Was ist ein Datenprodukt?



JOCHEN CHRIST
@JOCHEN_CHRIST