dojo

Building a business-critical data platform to process over £50bn in card transactions

dojo

www.dojo.tech

One of the largest Fintechs in Europe

Enabling 130,000 businesses to take card payments from 4 million consumers per day

Taking a card payment is complicated and highly regulated

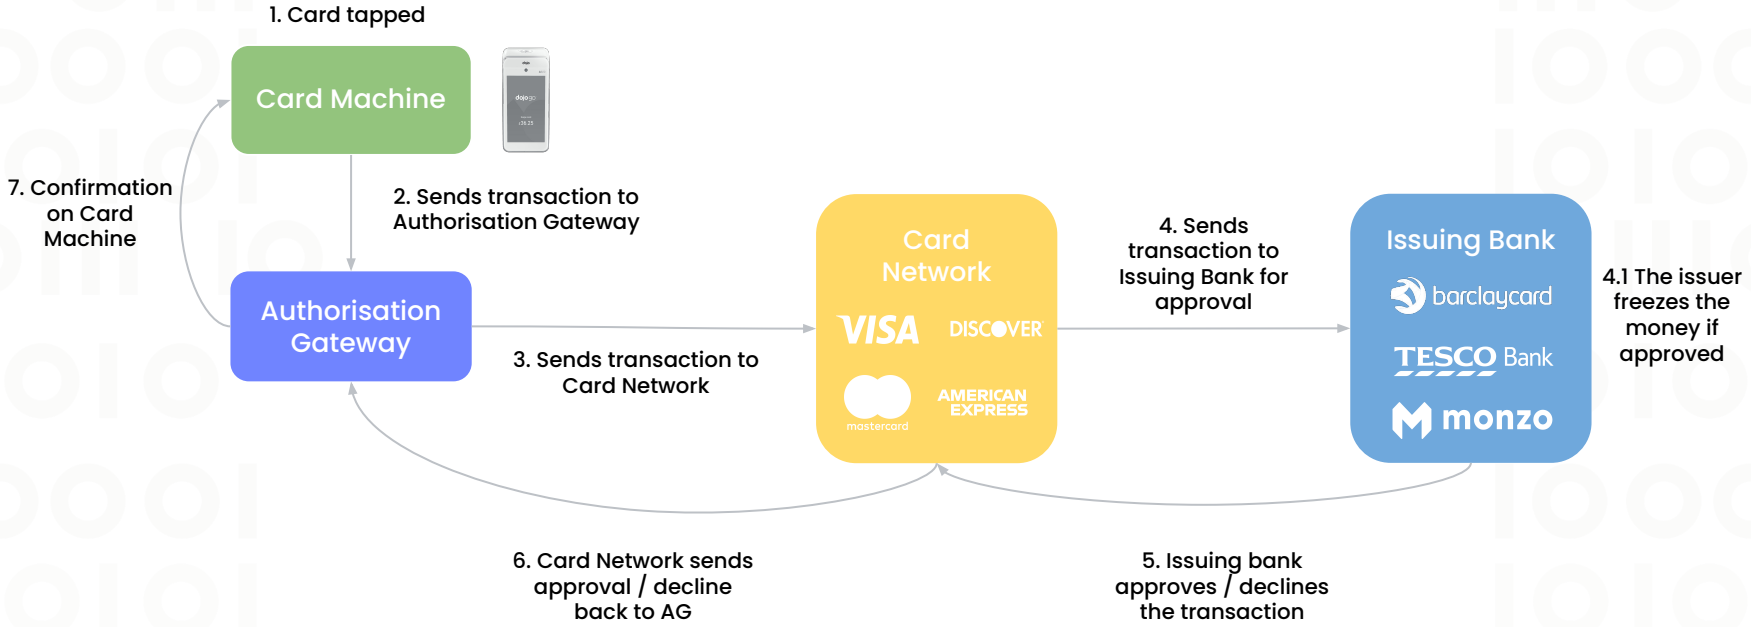

dojo

Challenges

Building a nuclear power station - it cannot fail

Regulatory

FCA FINANCIAL CONDUCT AUTHORITY

Safeguard customer funds from working capital at all times

PCI Security Standards Council

Data contains raw card numbers (PANs)

Complexity

Over 1000+ schemas
Vast number of different proprietary file formats that are fragile and frequently change

Multiple file formats and sizes
Files vary from few megabyte XML all the way through to multi-gigabyte proprietary files with tens of millions of rows.

Scalability

Unpredictable demand
Unpredictable surges in transaction volumes

Business Critical (99.99%)- Internally we refer to this as building a nuclear power station - it cannot fail.

dojo

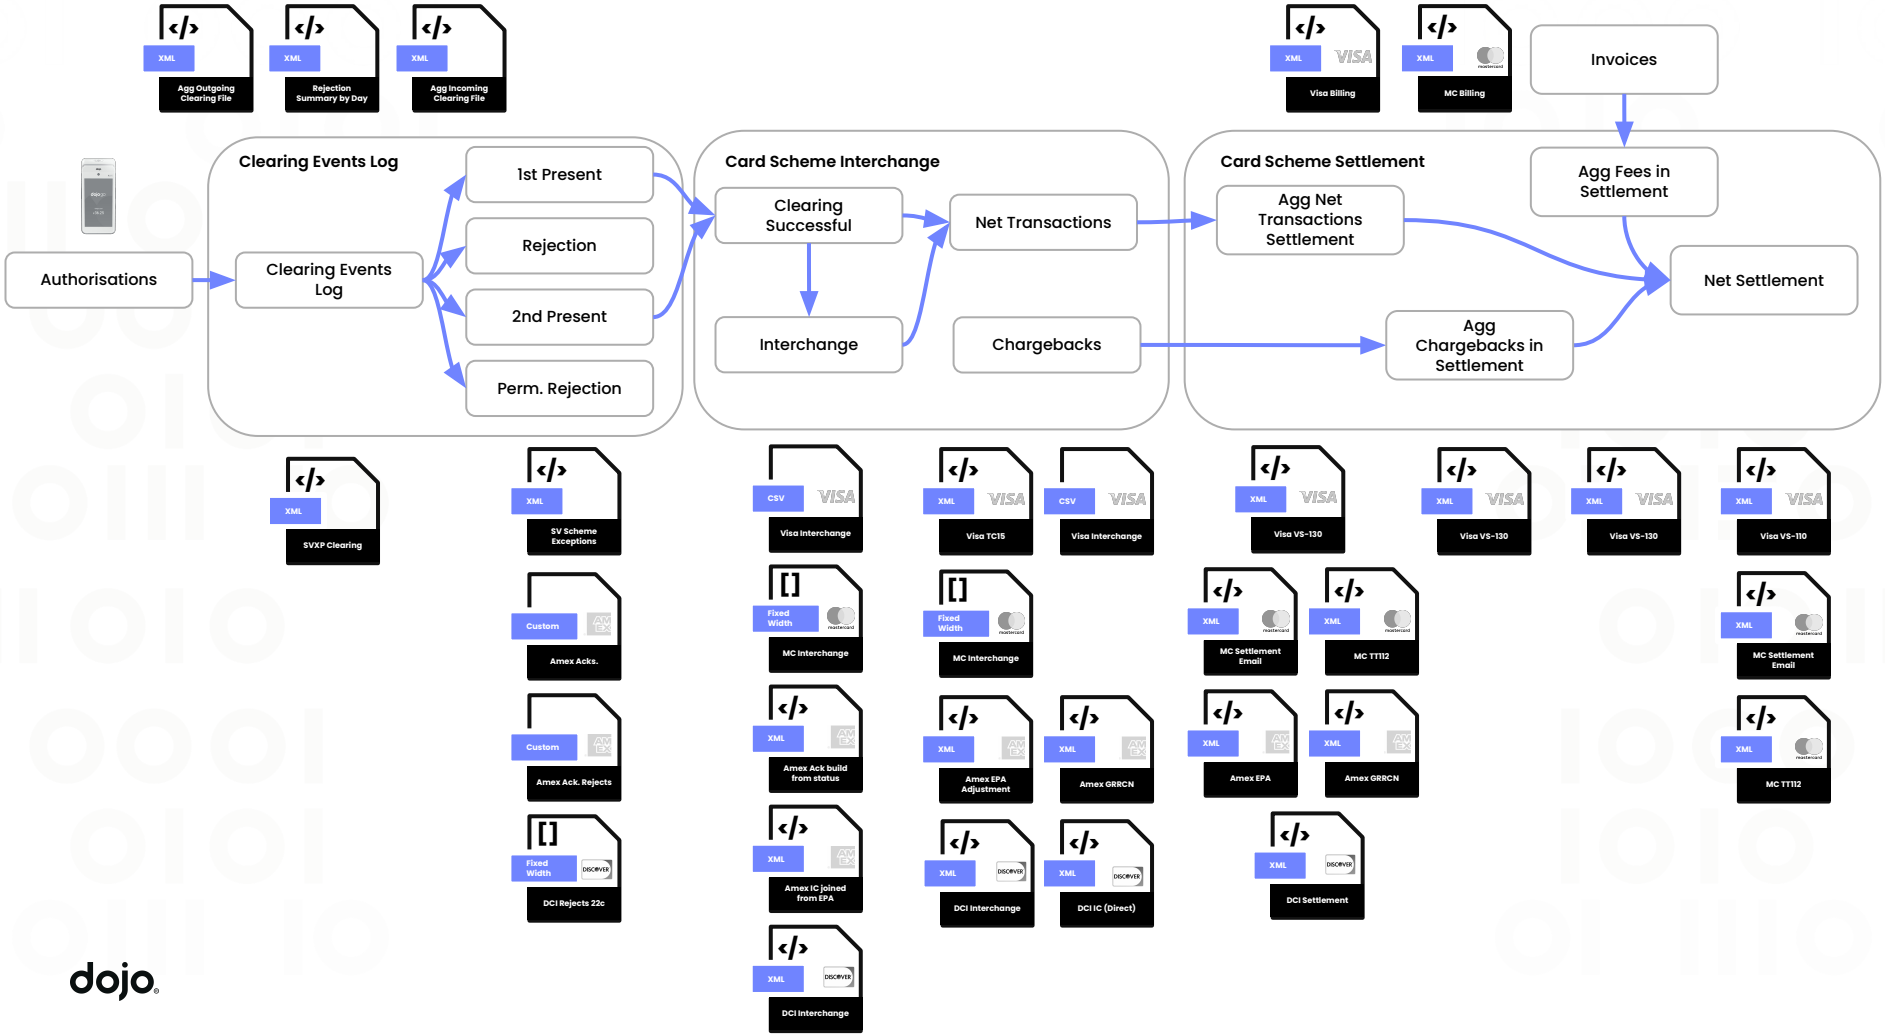
Agg Outgoing Clearing File
Rejection Summary by Day
Agg Incoming Clearing File
Visa Billing
MC Billing
Invoices
Clearing Events Log
1st Present
Rejection
2nd Present
Perm. Rejection
Authorisations
Clearing Events Log
Card Scheme Interchange
Clearing Successful
Net Transactions
Interchange
Chargebacks
Card Scheme Settlement
Agg Net Transactions Settlement
Agg Fees in Settlement
Net Settlement
Agg Chargebacks in Settlement
SVXP Clearing
SV Scheme Exceptions
Amex Acks.
Amex Ack. Rejects
DCI Rejects 22c
Visa Interchange
MC Interchange
Amex Ack build from status
Amex IC joined from EPA
DCI Interchange
Visa TC15
Visa Interchange
MC Interchange
Amex EPA Adjustment
Amex GRRCN
DCI Interchange
DCI IC (Direct)
Visa VS-130
Visa VS-130
Visa VS-130
MC Settlement Email
MC TT112
Amex EPA
Amex GRRCN
DCI Settlement
Visa VS-110
MC Settlement Email
MC TT112

Abstracting away the complexity

Transformation into a consistent file format - Avro

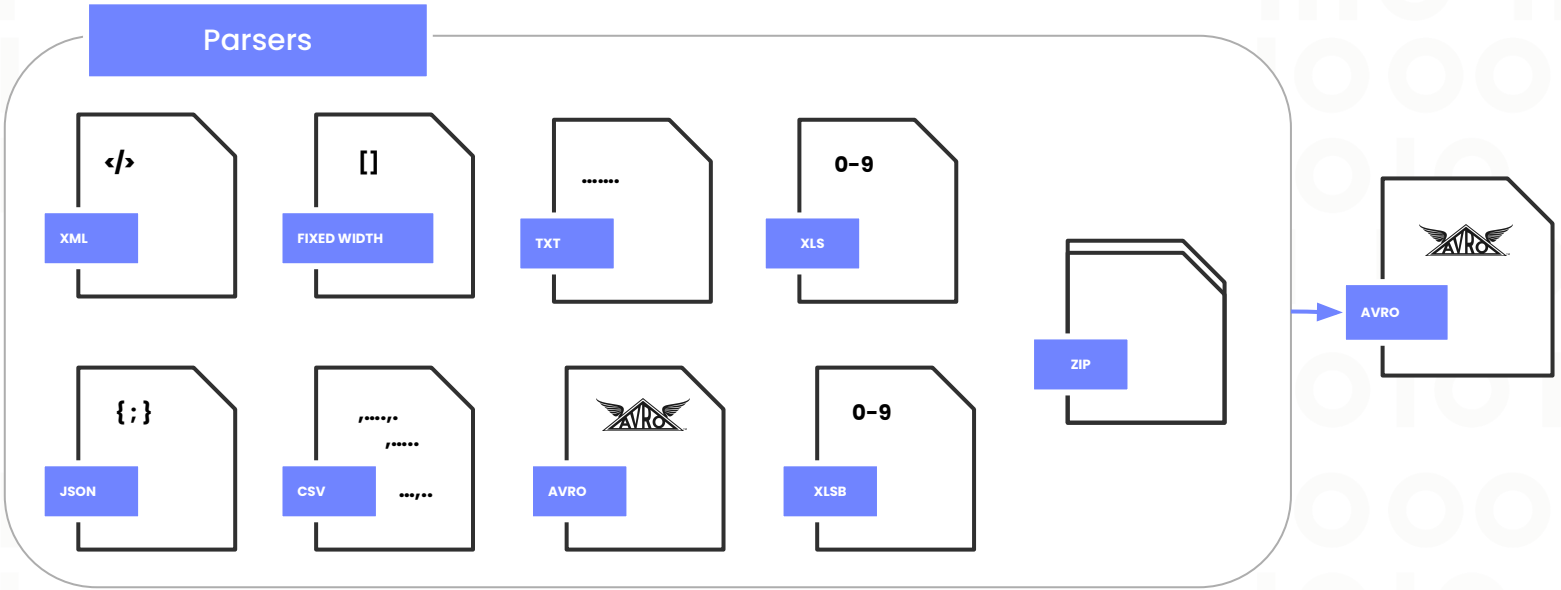

dojo

other Data Domains as well ...

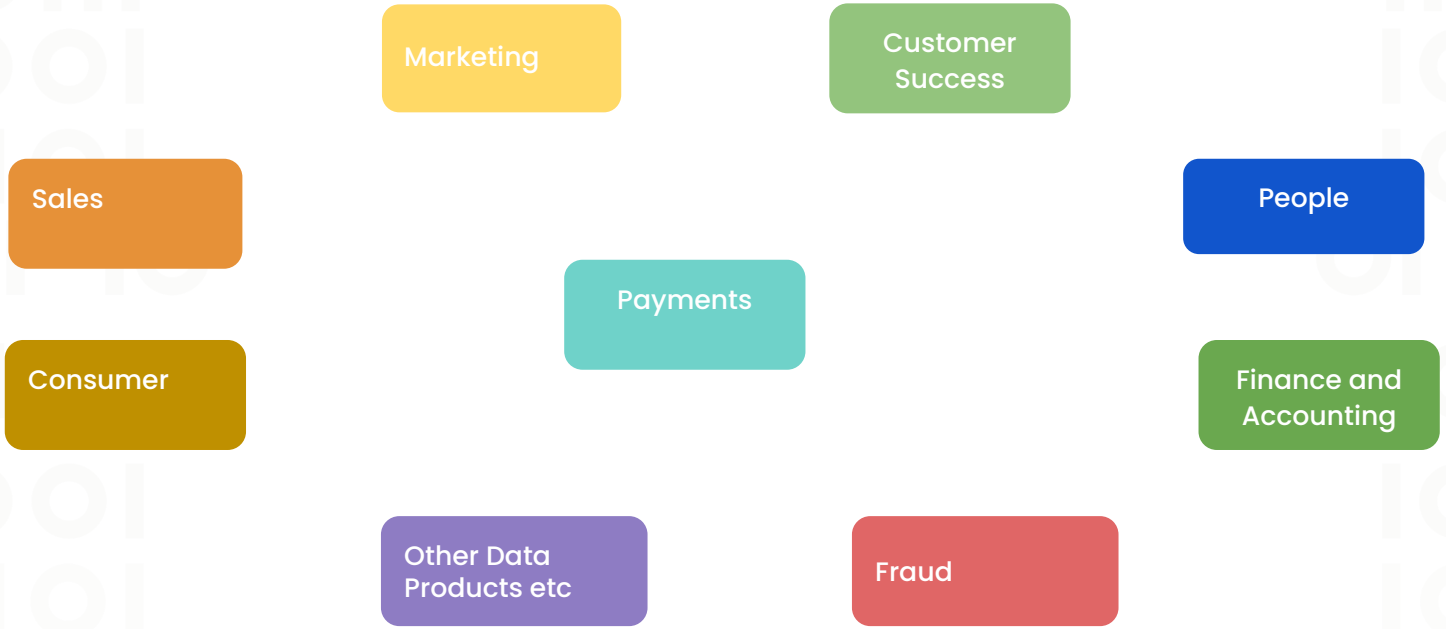

dojo

Data Infrastructure Generations

Challenges: Central Data Platform

Centralised Team Ownership

Siloed and Specialised Data Platform Team

Stretched Data Platform Team

Data Quality, Accountability and Democratisation

Data Quality Issues

Lack of Accountability when it comes to data issues

Delays in Data Access and Insights

Scalability

Data Volume and New Data Sources

Cost Efficiency

dojo

What's Next

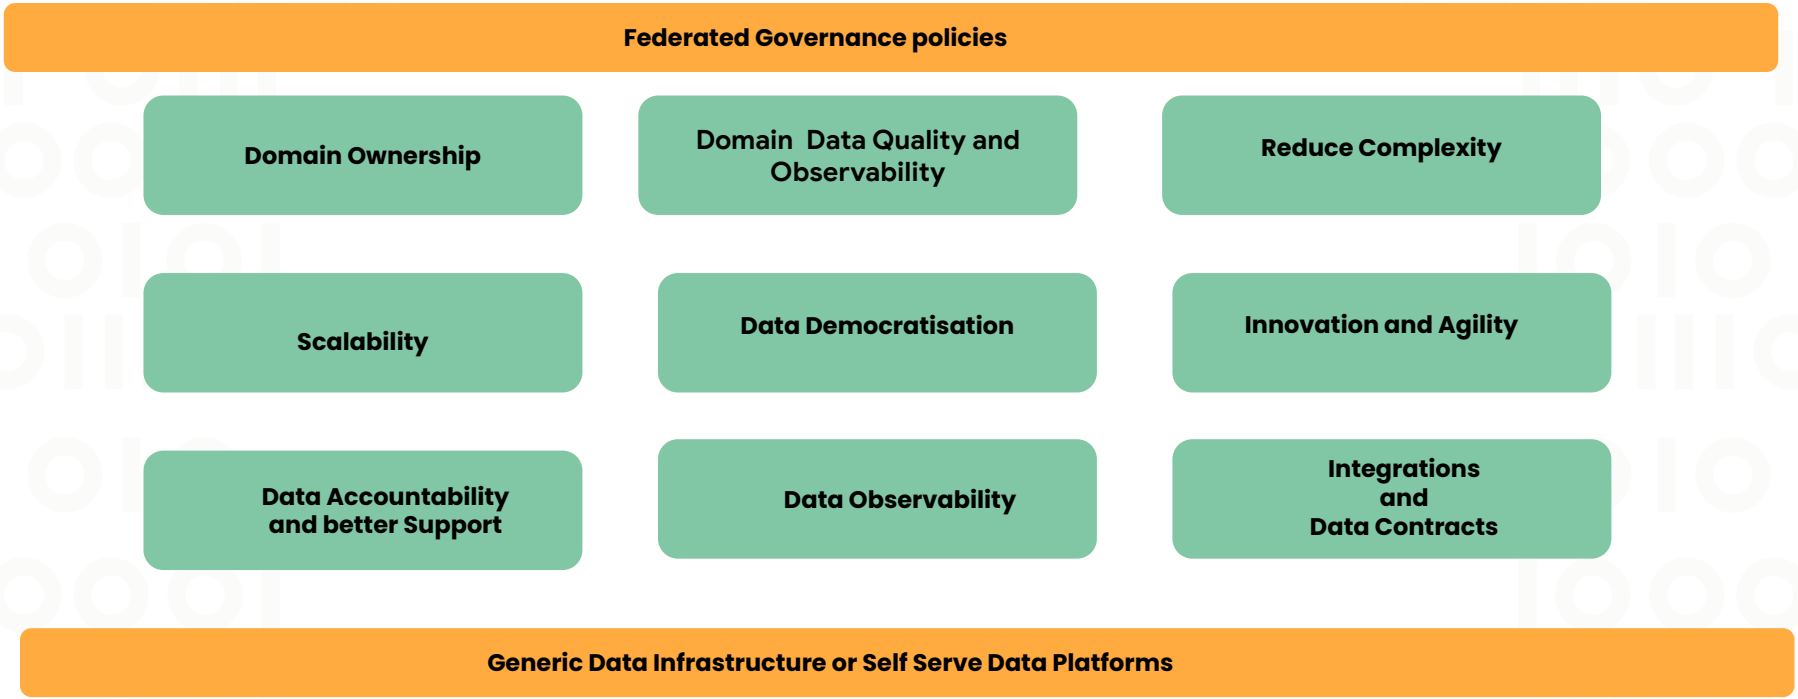

dojo

Data Mesh

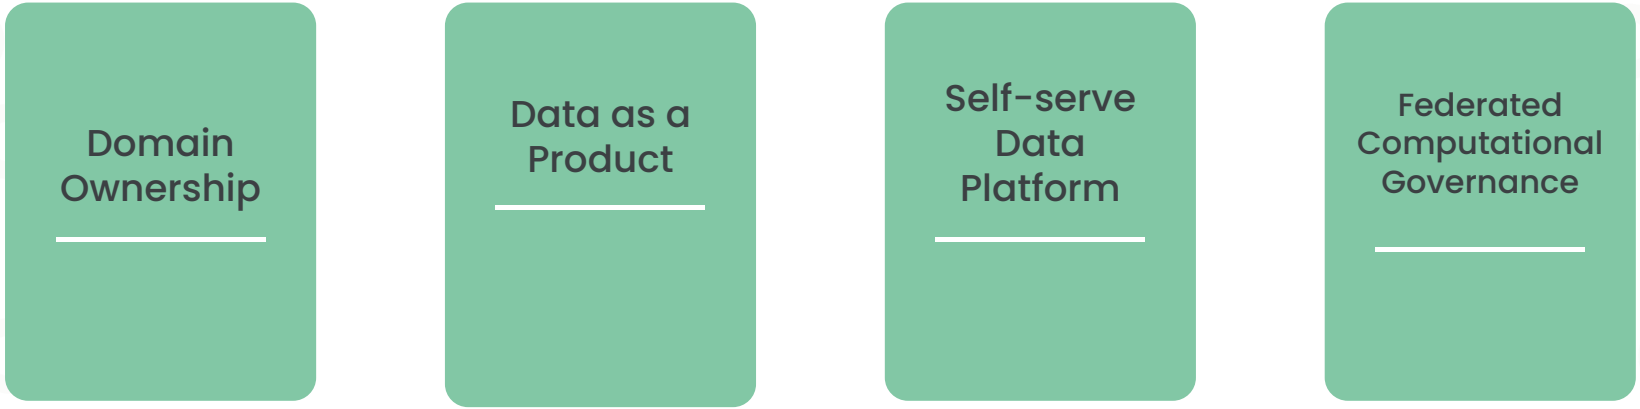
Domain Ownership
Data as a Product
Self-serve Data Platform
Federated Computational Governance

Modern Data Stack is Broken

What should we do ?

Build Small and Go Big

Open Source Tools

Early Feedbacks

Kubernetes is your friend

Managed Cloud Services can be handy

Cloud Agnostic

Platform Overview

Component based architecture

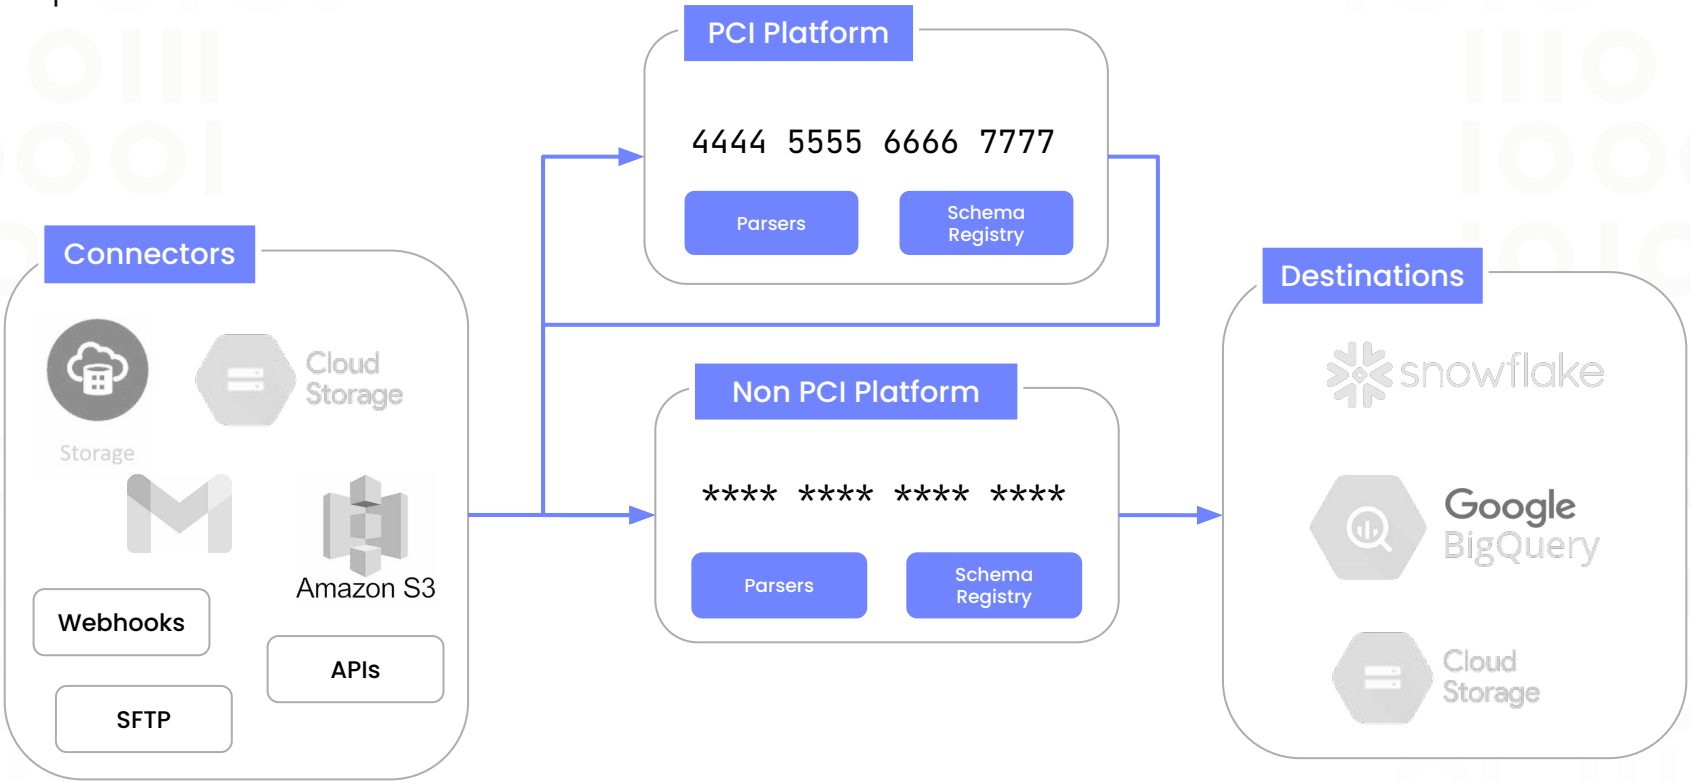

dojo

Connectors Overview

A unified approach to ingesting data from multiple sources

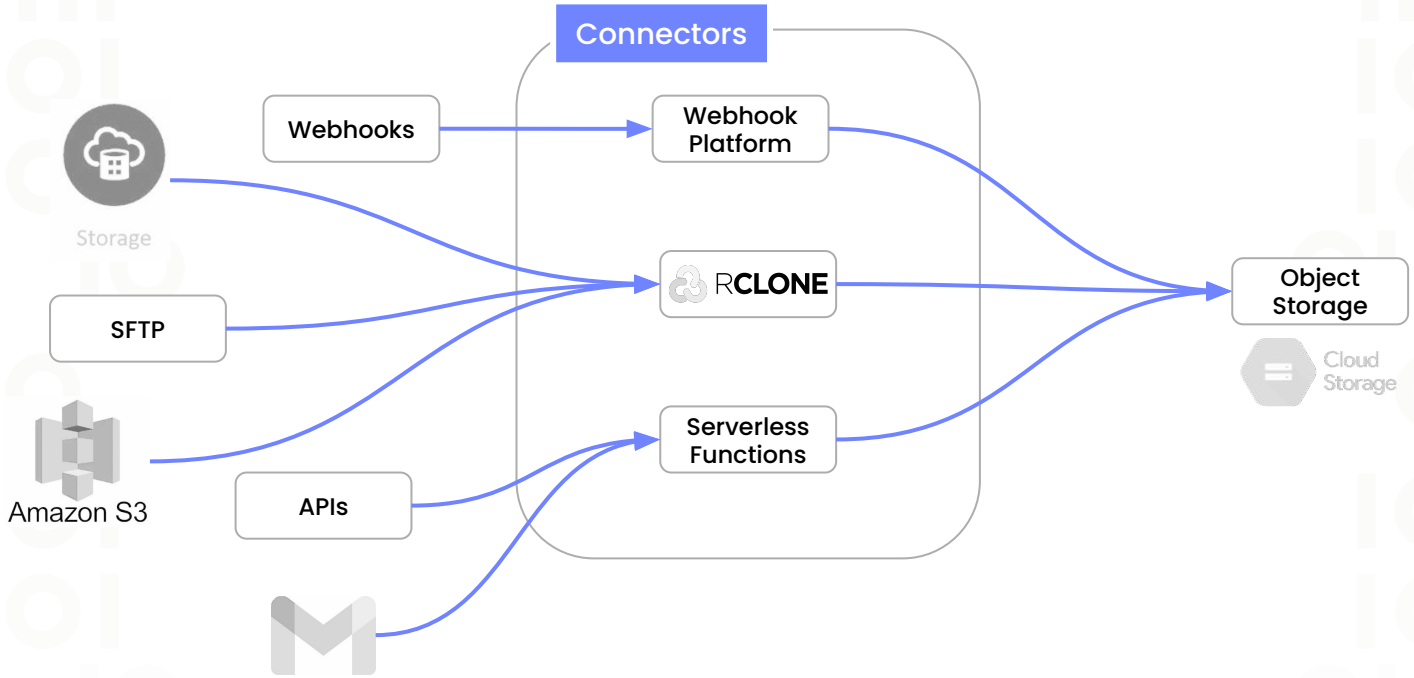

dojo

PCI DSS Level 1

Process, store, or transmit credit card or cardholder data maintain a secure environment

- Card data transmitted securely into the Data Platform
- Strong Encryption, Monitoring and Security testing of Data Platform
- Yearly Audits to verify the security of the Data Platform

PCI Security Standards Council®

dojo

PCI Management Process

Secure and prevent the egress of PCI sensitive data

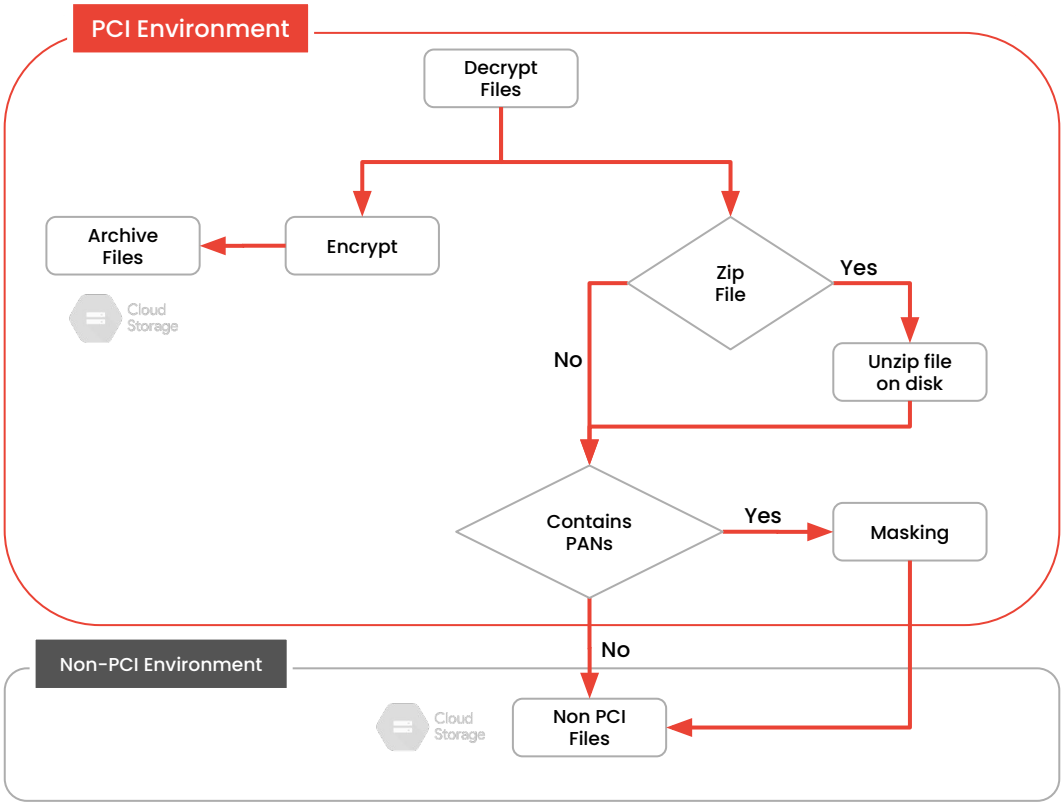

dojo

PCI Platform Overview

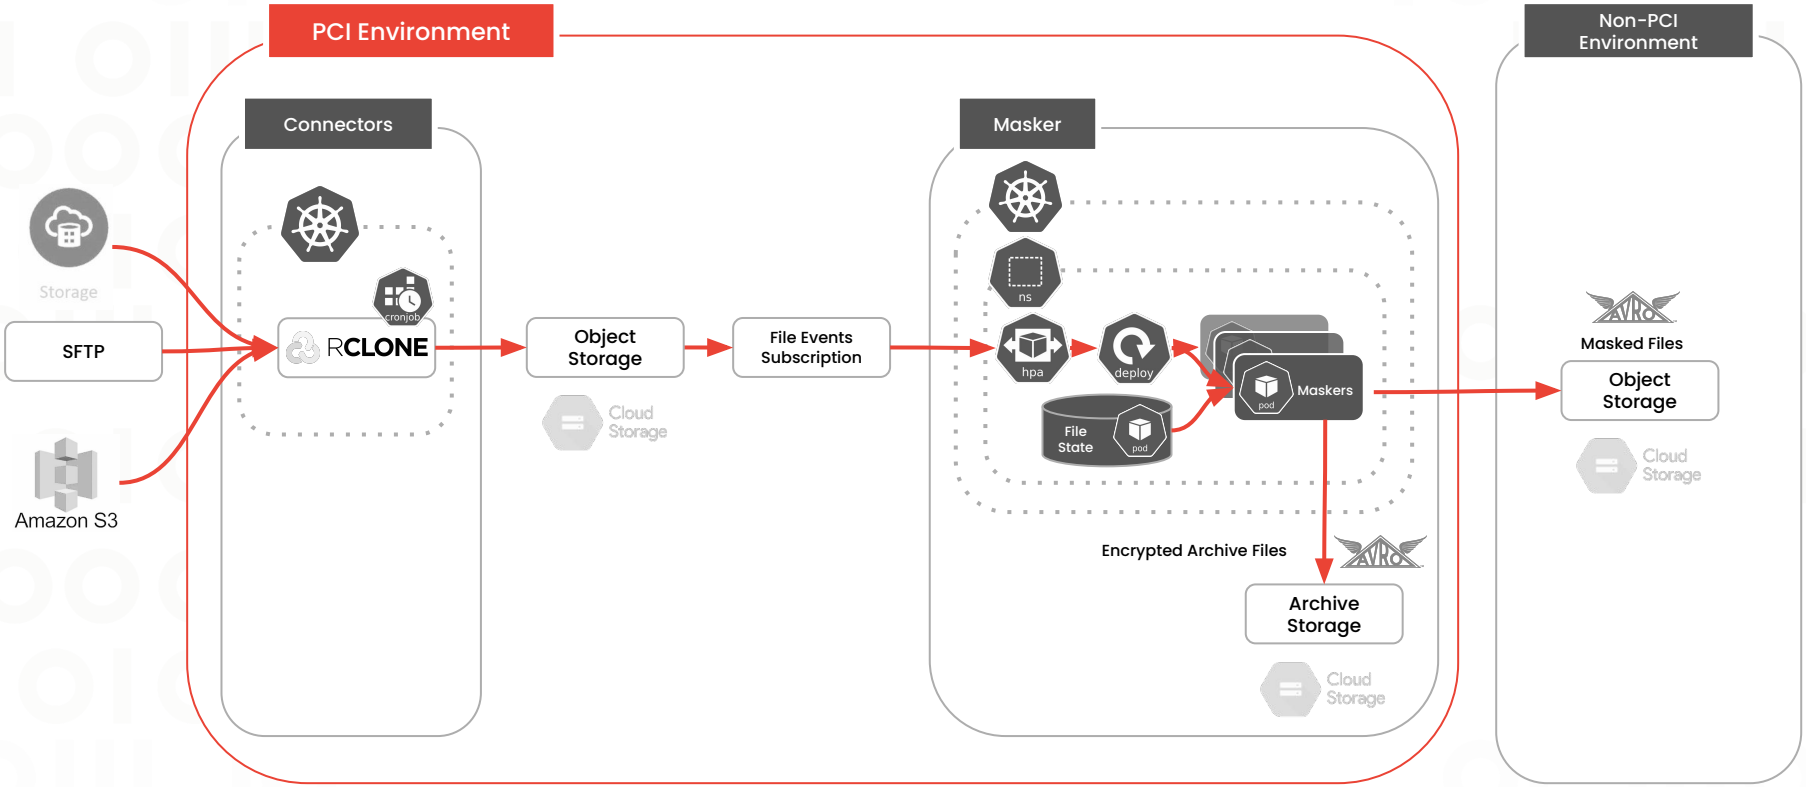

dojo

Transformation Platform Overview

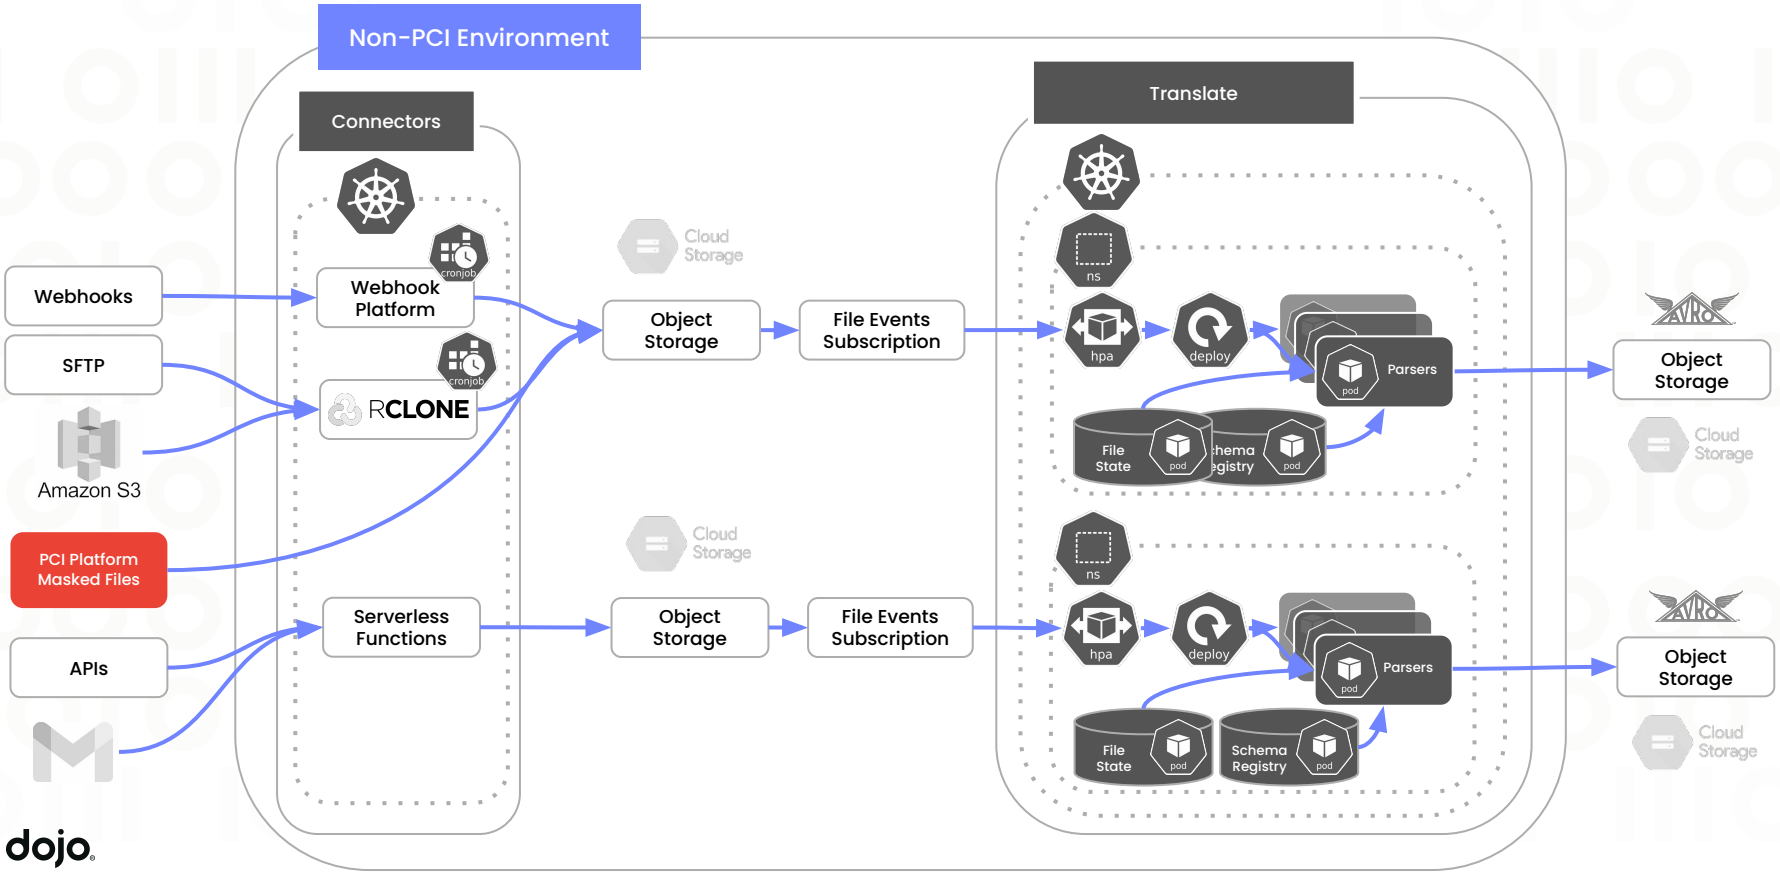

dojo

Autoscaling Challenges

Capacity planning is essential to scaling efficiently

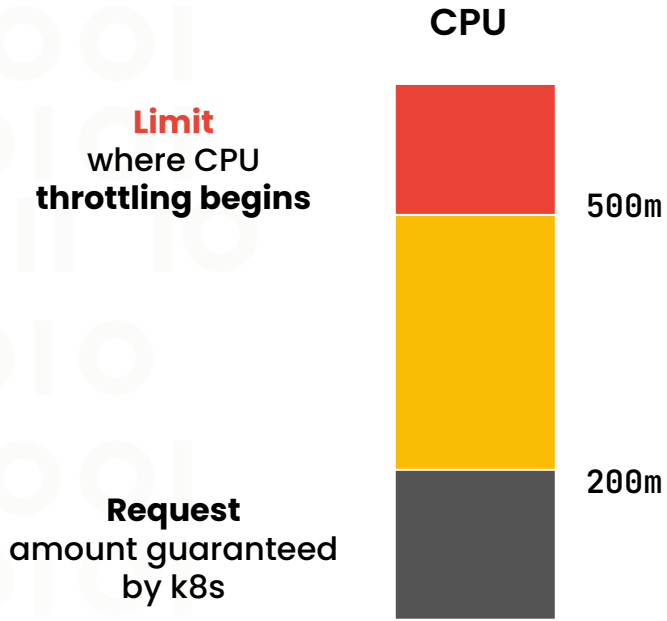

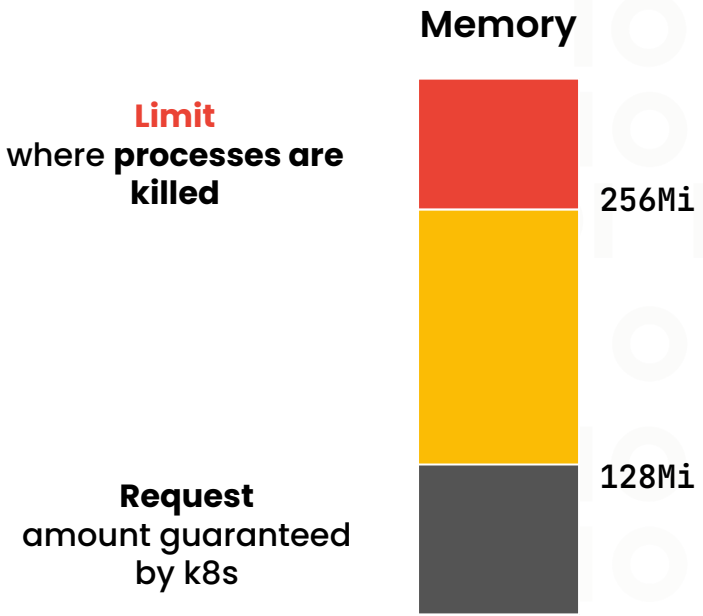

dojo

Autoscaling

Two approaches to autoscaling in Kubernetes - Horizontal (HPA) or Vertical (VPA)

Scale Out: Increase / decrease number of pods based on metrics

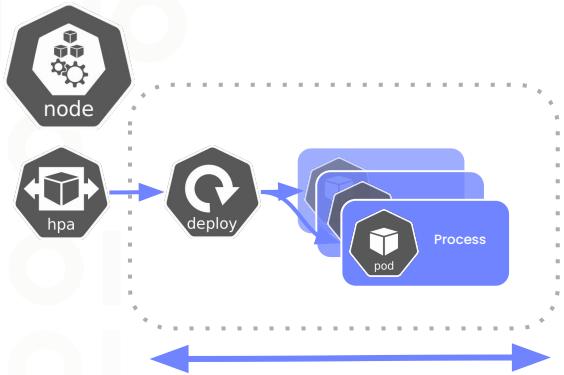

Scale Up: Increase resources assigned to workload

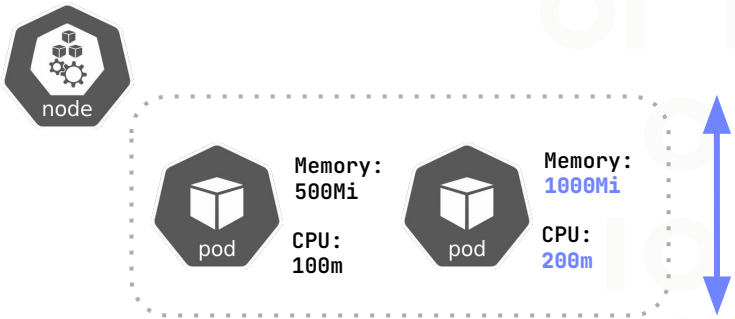

dojo

Autoscaling Triggers

Kubernetes provides three solutions depending on the use case

Resource Usage

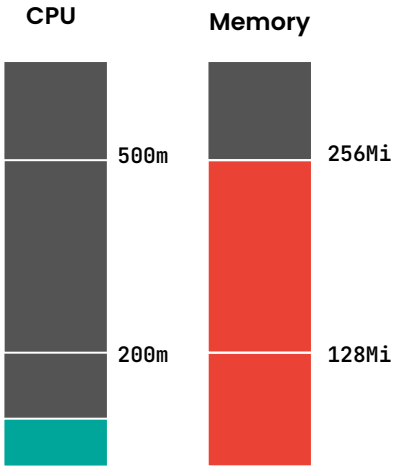

metrics.k8s.io

Custom Metrics

hits-per-second

custom.metrics.k8s.io

External Custom Metrics

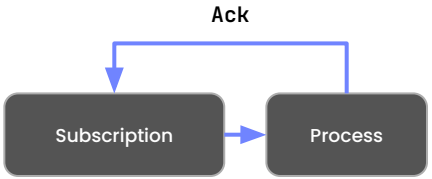

external.metrics.k8s.io

dojo

Using HPA to elastically scale

Pods required determined by number of unacknowledged messages in queue

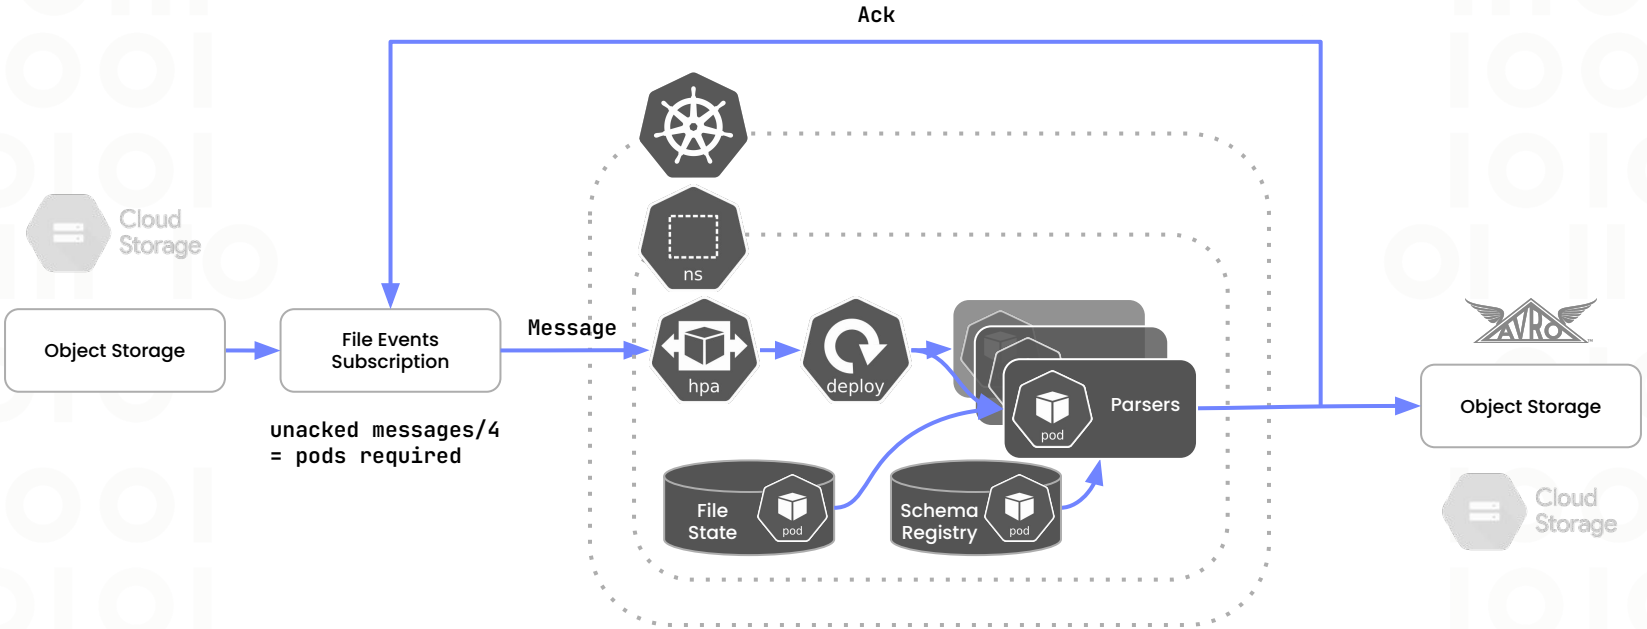

Cluster Autoscaler

New nodes provisioned based on pods in **`Pending`** state

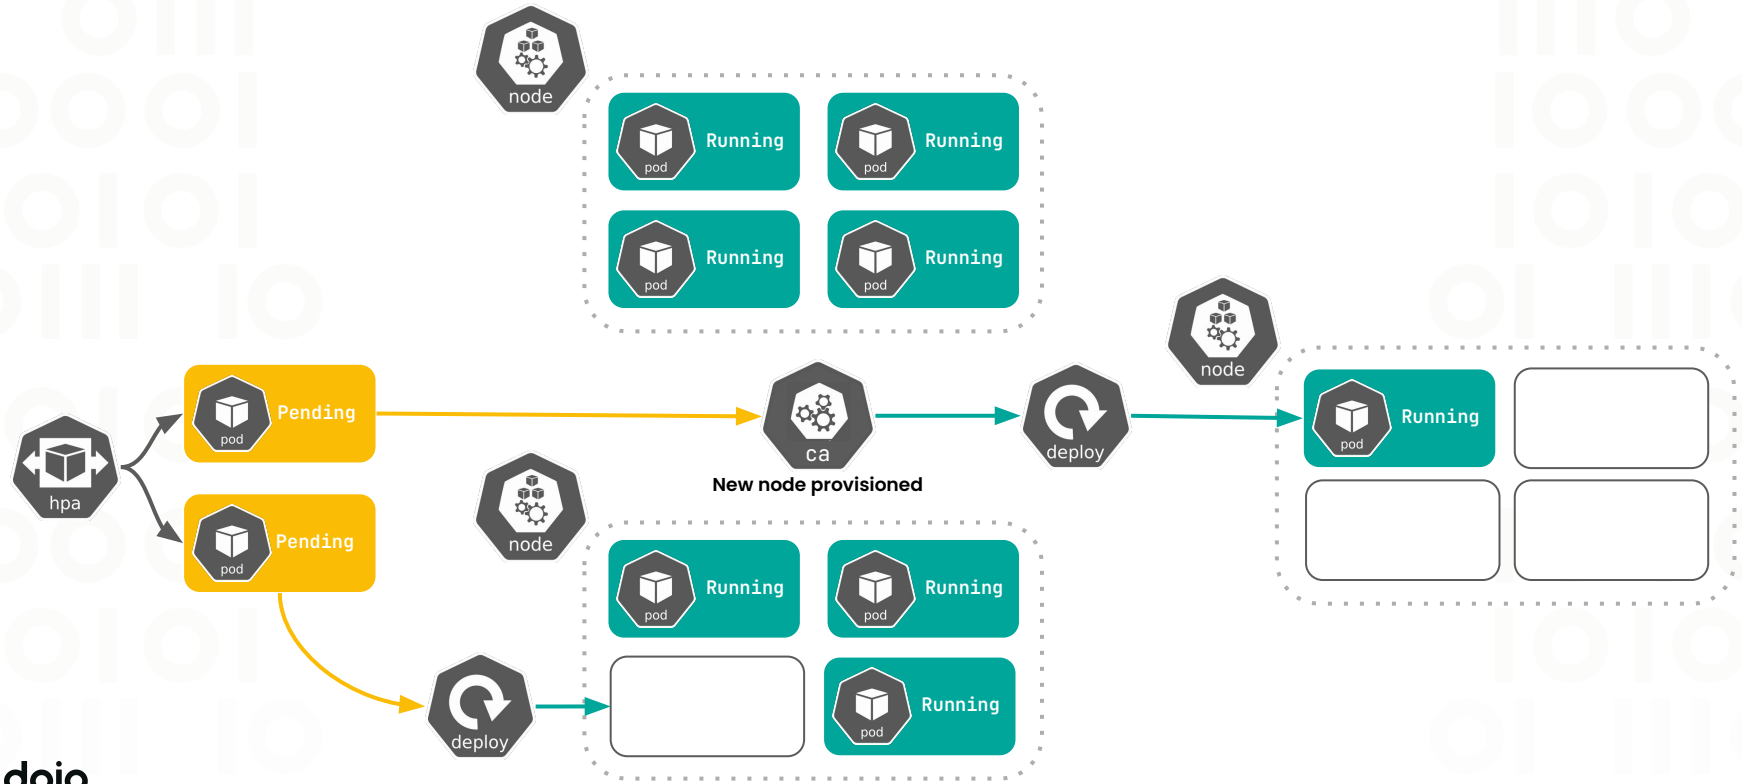

dojo

Using HPA to elastically scale

Pods required determined by number of unacknowledged messages in queue

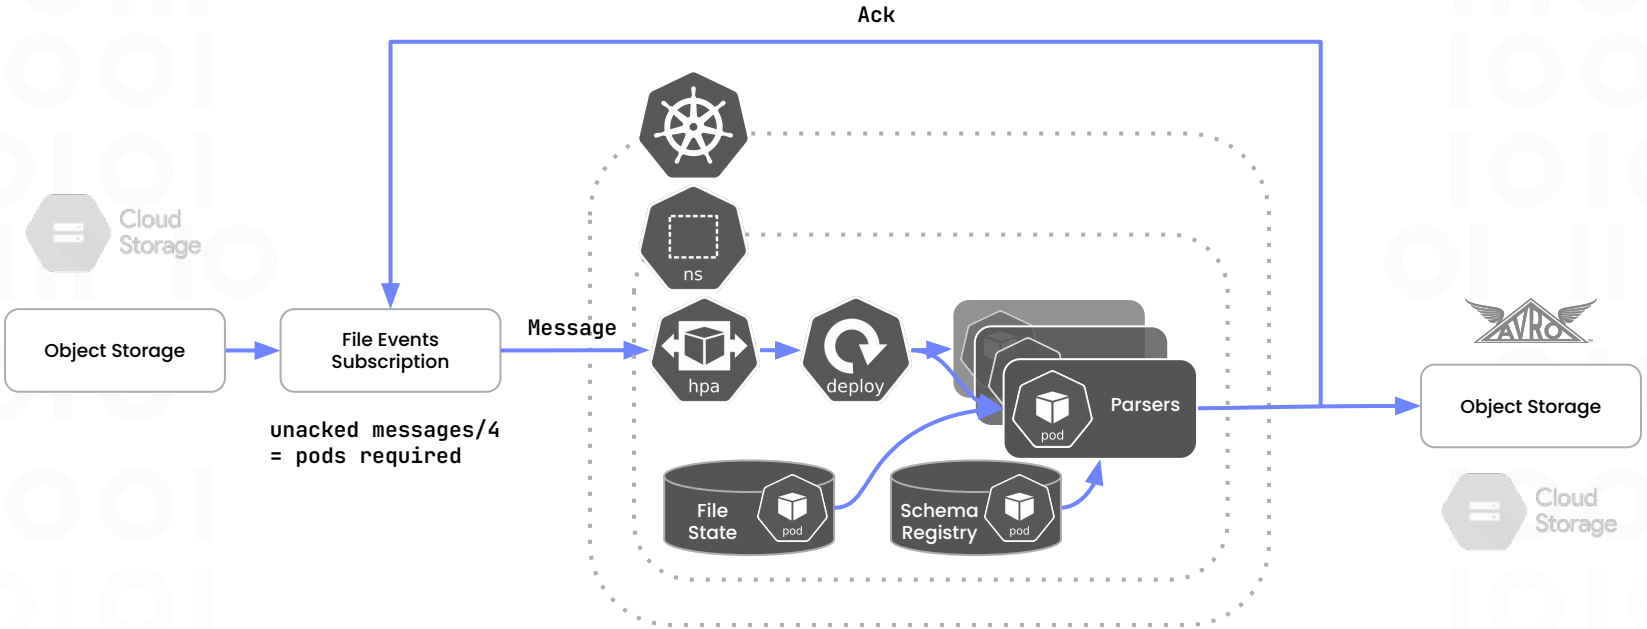

Schema Registry

A central source of metadata for file lifecycle management

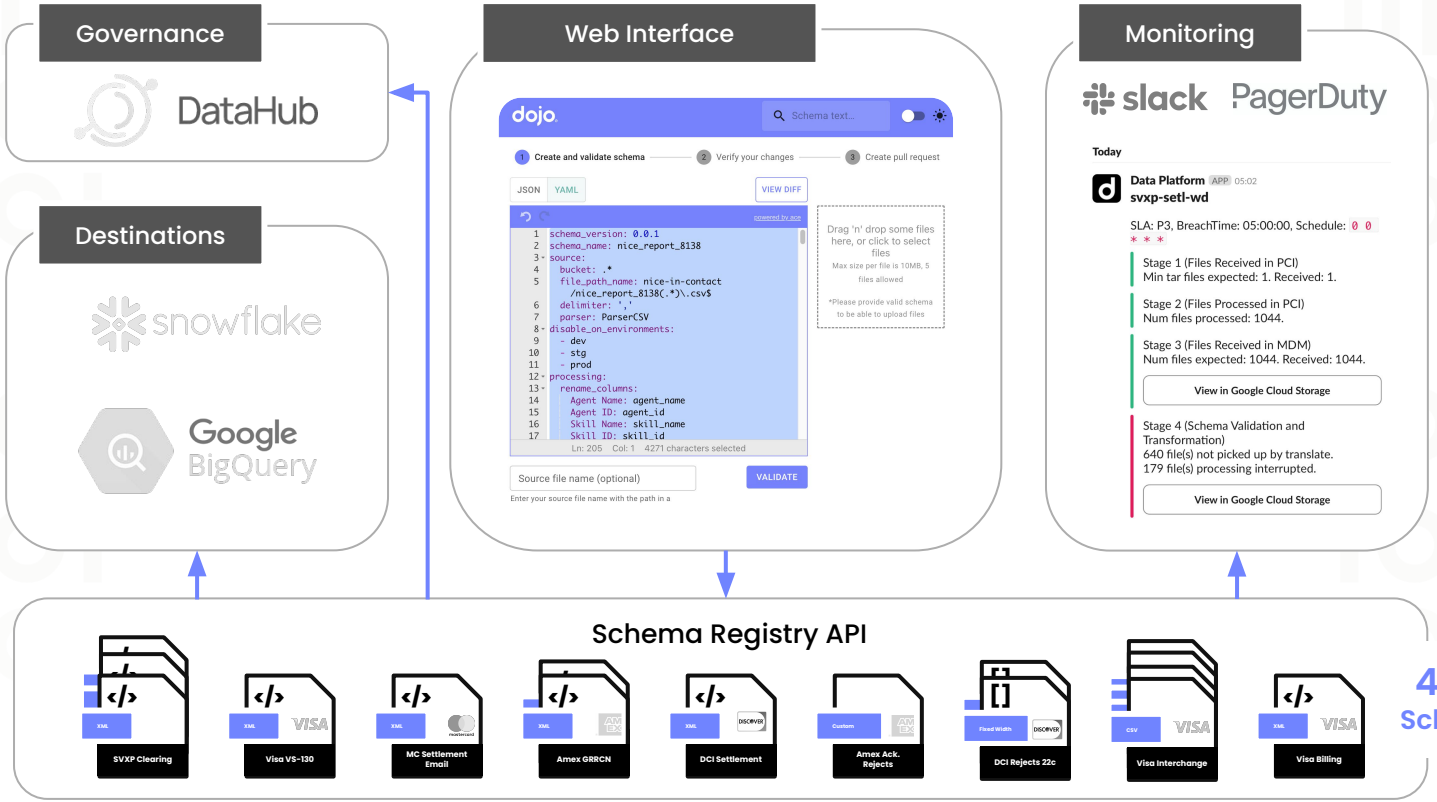

dojo.

State Management

State of every file recorded throughout the transformation process

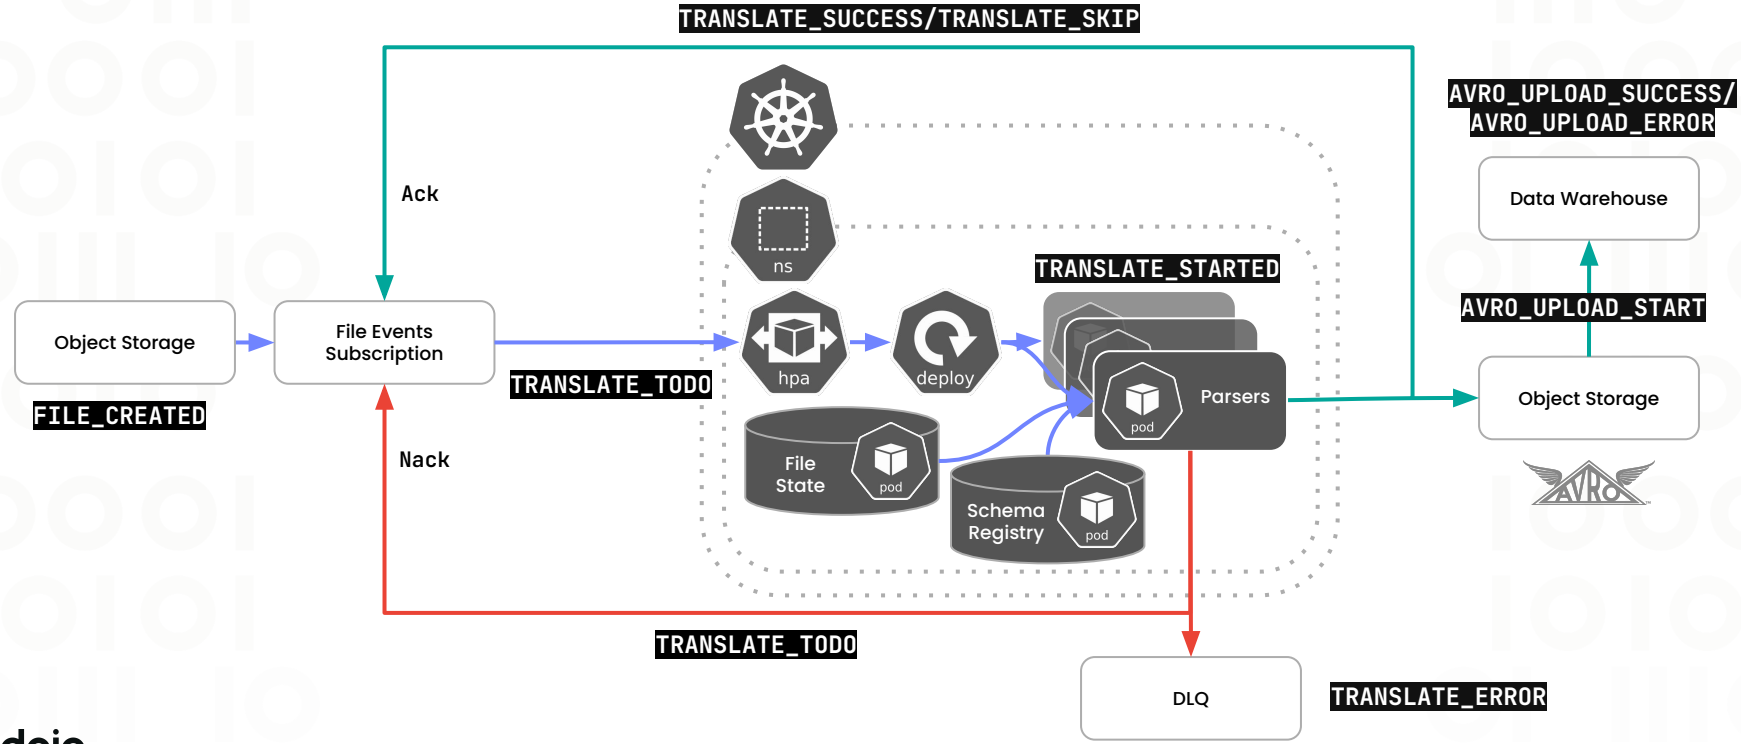

dojo

Real time File Monitoring

State of every file recorded throughout the transformation process

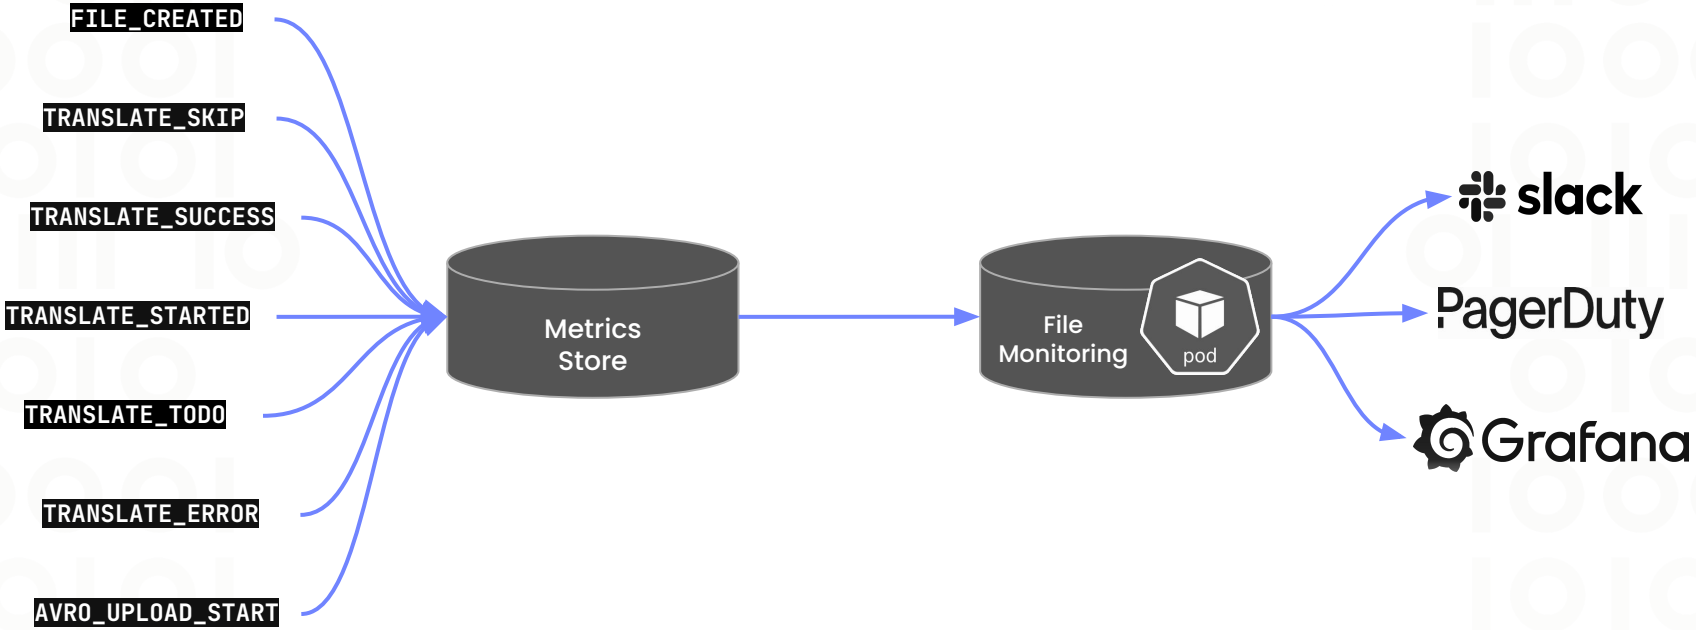

dojo

E2E File Monitoring

PagerDuty slack

P0 P1 P2 P3

Looker

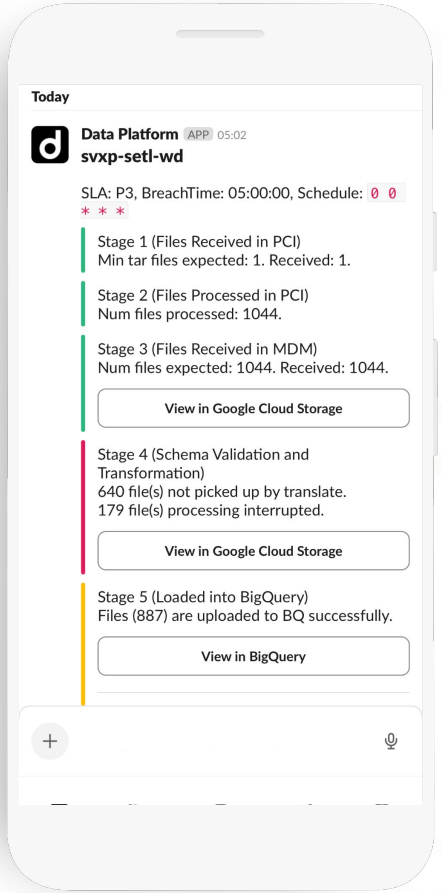

dojo

Analytics Platform

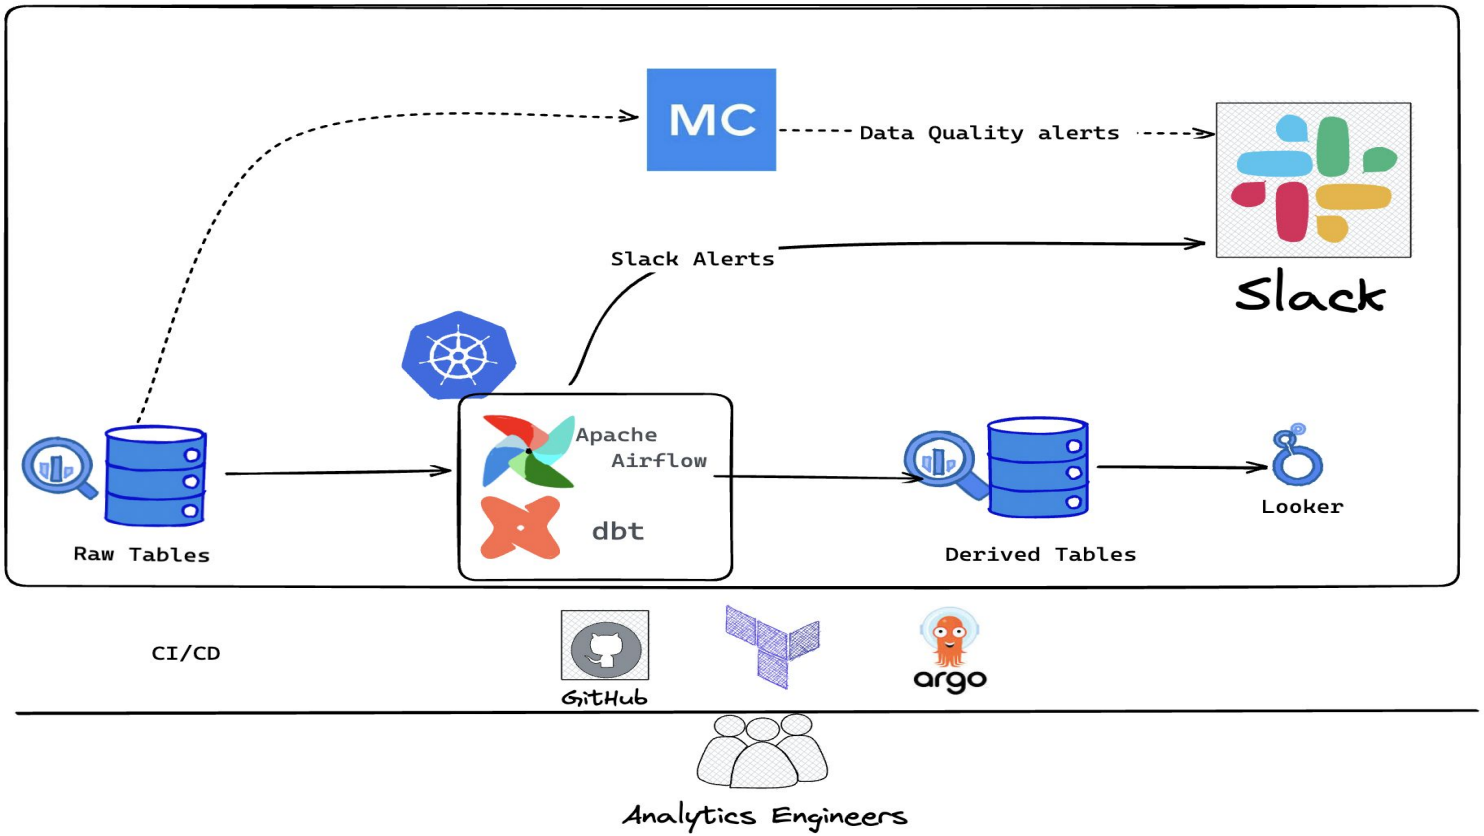

Infrastructure Observability

Monitoring and alerting on Data Platform resources

Challenges: Embracing Data Mesh

- Cultural Shift
- Talent and Skill Gaps
- Data Governance Complexity
- Data Mesh Tooling
- Data Privacy and Security
- Data Integrations
- ROI Measurement
- Data Cataloging and Data Discovery
- Legacy system Integrations

dojo

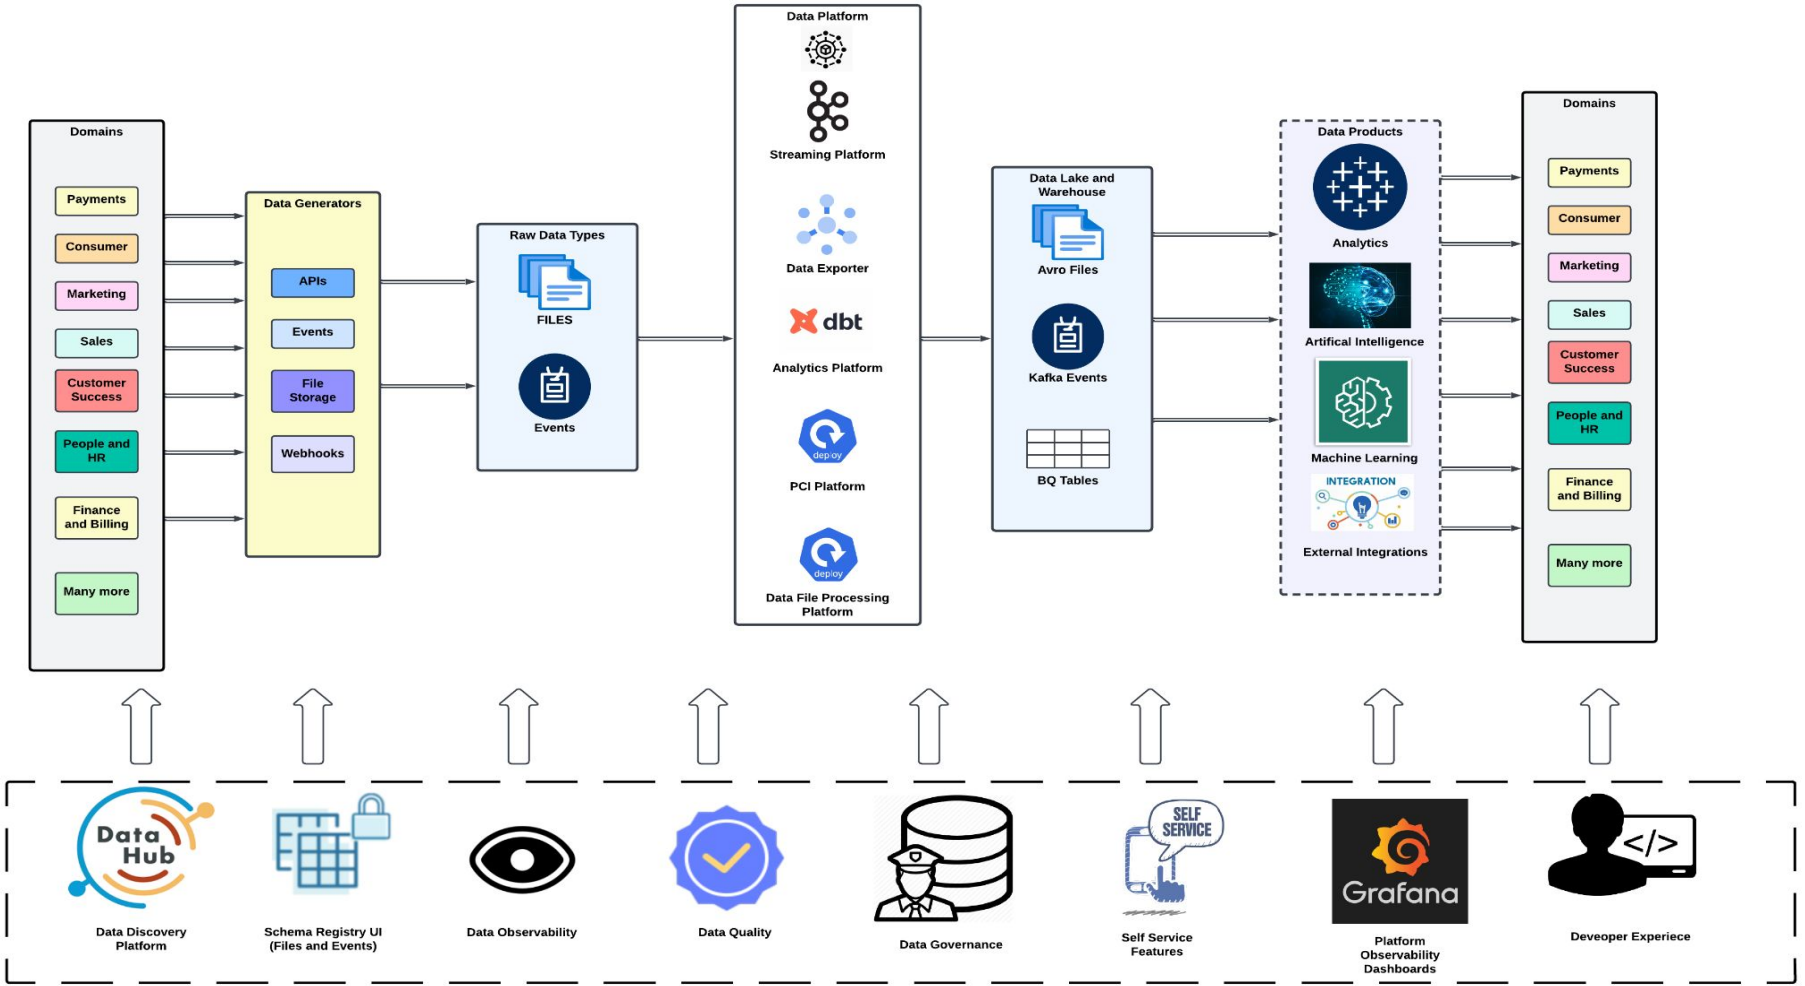
Domains
Payments
Consumer
Marketing
Sales
Customer Success
People and HR
Finance and Billing
Many more
Data Generators
APIs
Events
File Storage
Webhooks
Raw Data Types
FILES
Events
Data Platform
Streaming Platform
Data Exporter
dbt
Analytics Platform
deploy
PCI Platform
deploy
Data File Processing Platform
Data Lake and Warehouse
Avro Files
Kafka Events
BQ Tables
Data Products
Analytics
Artifical Intelligence
Machine Learning
INTEGRATION
External Integrations
Domains
Payments
Consumer
Marketing
Sales
Customer Success
People and HR
Finance and Billing
Many more
Data Hub
Data Discovery Platform
Schema Registry UI (Files and Events)
Data Observability
Data Quality
Data Governance
SELF SERVICE
Self Service Features
Grafana
Platform Observability Dashboards
</>
Deveoper Experiece

Thank you for your time!

https://www.dojo.careers

Sandeep Mehta

Data Platform Lead

sandeep-mehta26