



Will HAL Open the Pod Bay Doors?

An (Enterprise FI) Decisioning Platform Leveraging Machine Learning

October 11, 2018

INTRODUCTION



Sumit Daryani

Manager, Software Engineering

Sumit is working on a real-time machine learning decision platform to protect the banking platform and foster quick decisions to support the fraud strategy. Prior to Capital One, Sumit has been a full-stack engineer on a diverse number of projects scaling from the Financial to Technology space.




Niraj Tank

Sr. Manager, Software Engineering

Niraj is working on a team which has built a fast data streaming and decisioning platform for Capital One Bank. Niraj has been an engineer for past 21 years, his diverse experience ranges from developing products for startups to leading various large-scale integration services.

DECISIONING

The act or process of deciding; determination, as of a question or doubt, by making a judgment.



HAL 9000 ~ Fictional character

WHAT TO EXPECT FROM THIS SESSION

- Our journey in building a decisioning platform
- How to achieve operational excellence
- Architectural choices to sustain the growth
- Tools and Automation
- Identify techniques and pitfalls to avoid
- Q & A



OUR CHALLENGES

Problem Statement

- Rule & Model based solutions on COTS product
- Slow delivery to production
- Need for real-time risk assessment
- Proprietary tech stack
- Requires specialized skill sets



OUR MISSION

Machine Learning at the Core

- Embrace Open Source
- Speed to market
- Rules to ML Model based solutions
- Algorithms – champion/ challenger decision strategies
- Run over multiple iterations - refine - rinse - repeat
- Learn and improve from experience



DECISIONING PLATFORM: TAKE ONE

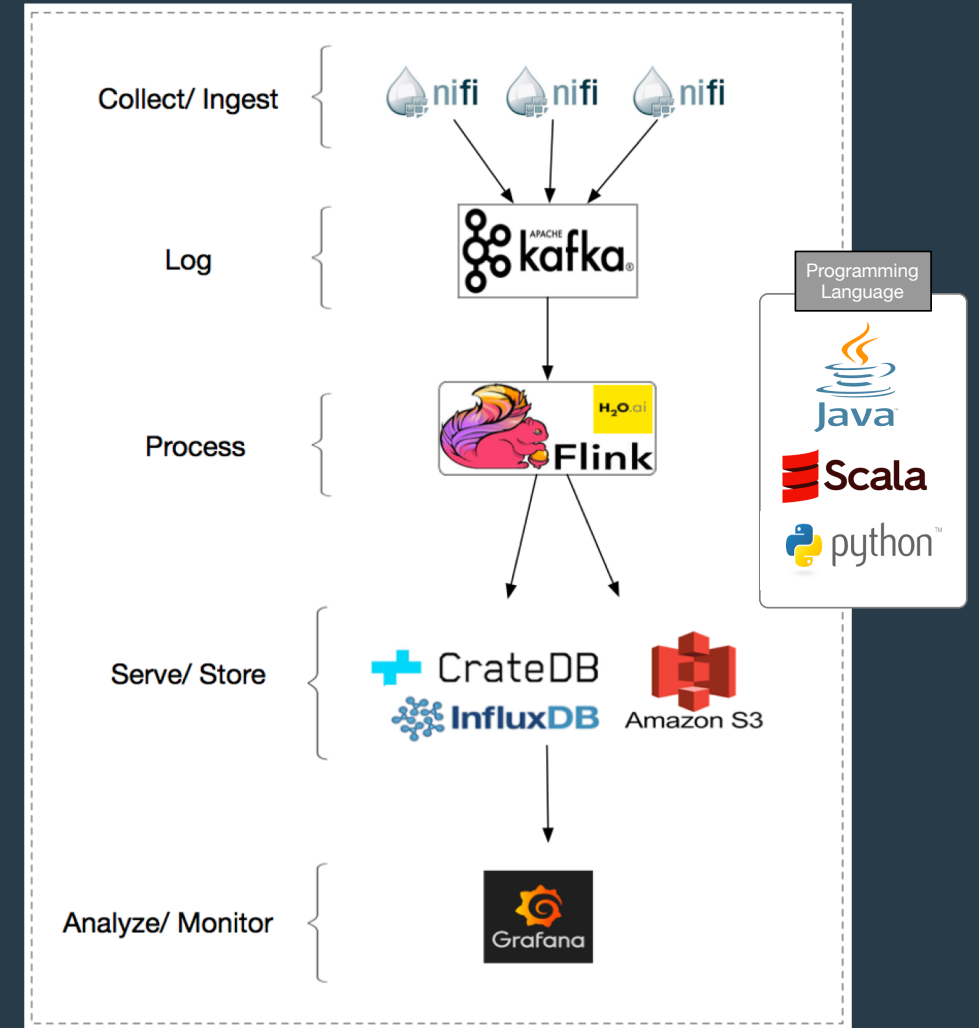
“I am the H.A.L 9000. You may call me Hal.” ~ HAL 9000

DECISIONING PLATFORM – Take 1

- Batch and Micro-batch use case
- Rule based to ML based models

Open source software:

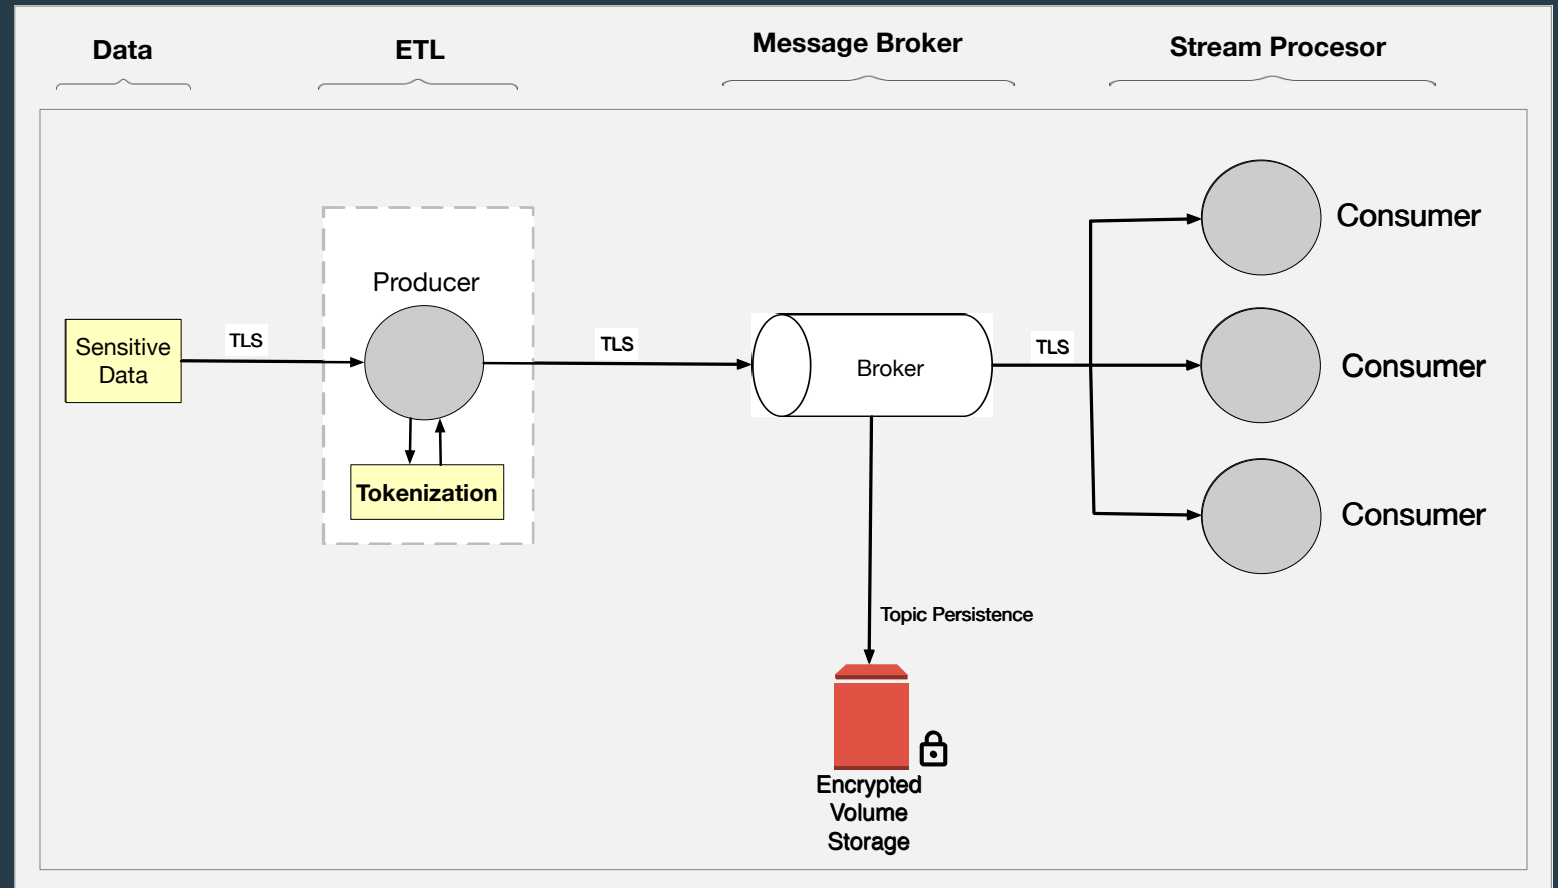
- Acquiring data = custom ETL using Apache Nifi
- Stream processing window aggregations = Flink
- Message bus = Kafka
- Real-time DB = CrateDB
- Monitoring = Grafana
- Analytics = SQL over JDBC



DECISIONING PLATFORM – Take 1

Sensitive Data Protection

- Responsibility to handle customer data
- Data in-transit and data at rest



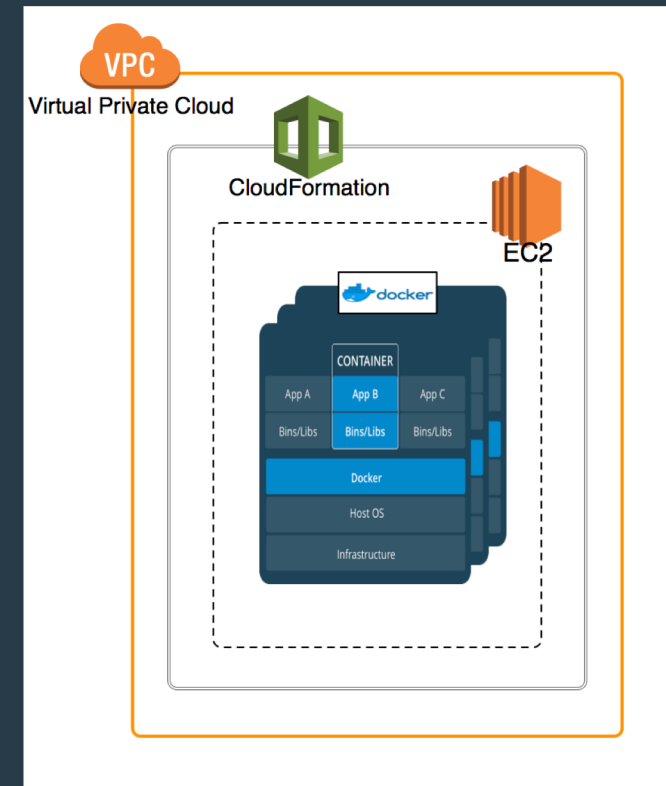
DECISIONING PLATFORM – Take 1

Deployment on AWS Cloud

- CloudFormation
- Docker Compose

End State

- Pattern supported: Micro Batch
- Use Case: 1
- Time To Market: 5 months
- Customer: Business user
- Models supported: 1



DECISIONING PLATFORM: TAKE TWO

“I am completely operational, and all my circuits are functioning perfectly.” ~ HAL 9000

DECISIONING PLATFORM – Take 2

- Setting the stage for enterprise level infrastructure
- Automated deployments
- Business Analytics
- Simple Data Redundancy
- Monitoring Dashboards and Alerts

Start of an Enterprise scale platform: **Kubernetes**

- Container Orchestration
- Maximize resource utilization
- Greater Computing Capacity
- Kubernetes Stateful Sets

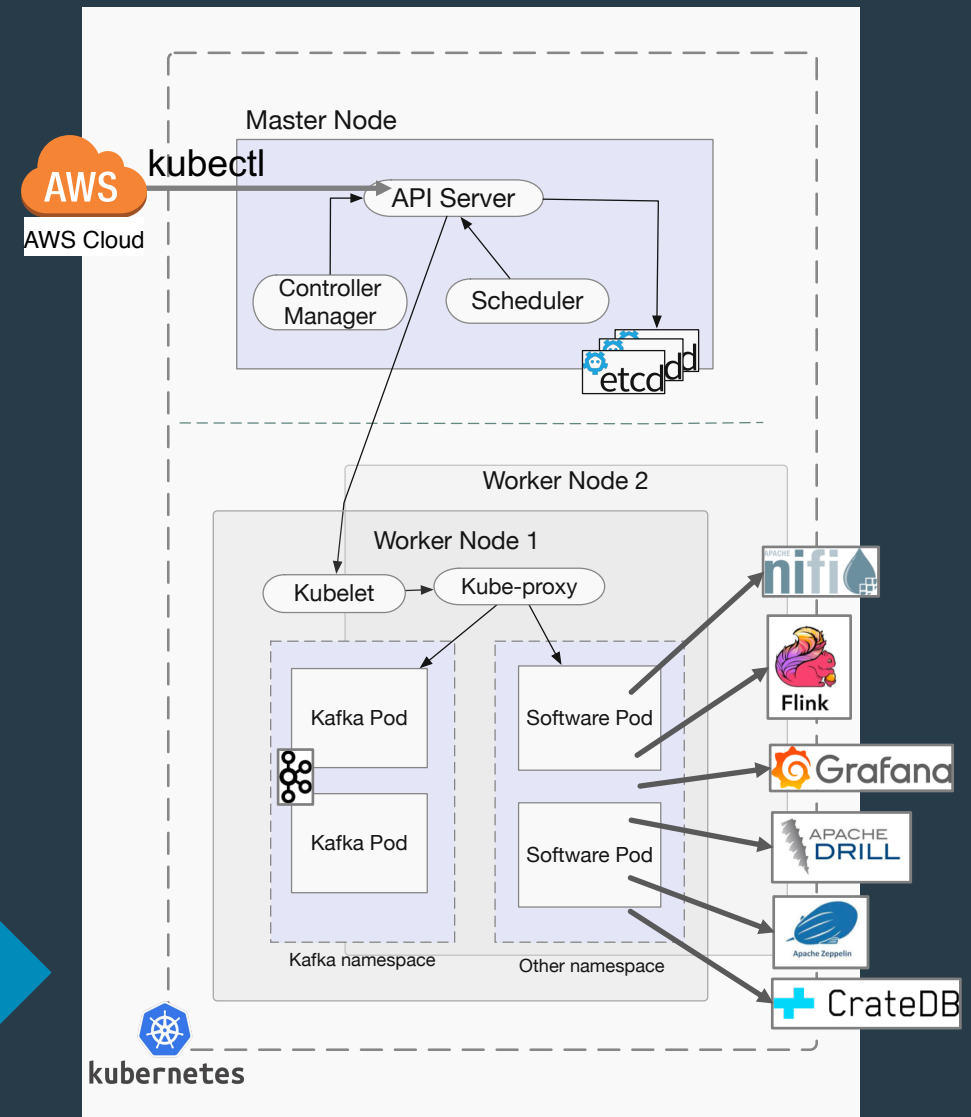


DECISIONING PLATFORM – Take 2

End State

- Infrastructure: K8S master nodes and worker nodes
- Data streams: 2
- Analytics: Apache Zeppelin and Apache Drill
- Monitoring and Alerting: Grafana
- Automation: Jenkins pipeline
- Time To Market: 3 months
- Customer: Business user

K8S Cluster
(Initial state)



DECISIONING PLATFORM: TAKE THREE

“I've still got the greatest enthusiasm and confidence in the mission. And I want to help you” ~ HAL 9000

DECISIONING PLATFORM – Take 3

- Microservices
- CI/CD
- Enterprise Logging strategy
- Enterprise Monitoring strategy
- Resiliency

Kubernetes updates:



- Blue-Green component upgrades
- Start of multi-tenancy
- Increased and Redundant Storage
- Fault Tolerance and Availability
- Custom CLI K8S tooling



DECISIONING PLATFORM – Take 3

Flink Microservices

- Flink's Queryable state
- Intermediate Kafka topics
- Decoupled deployments



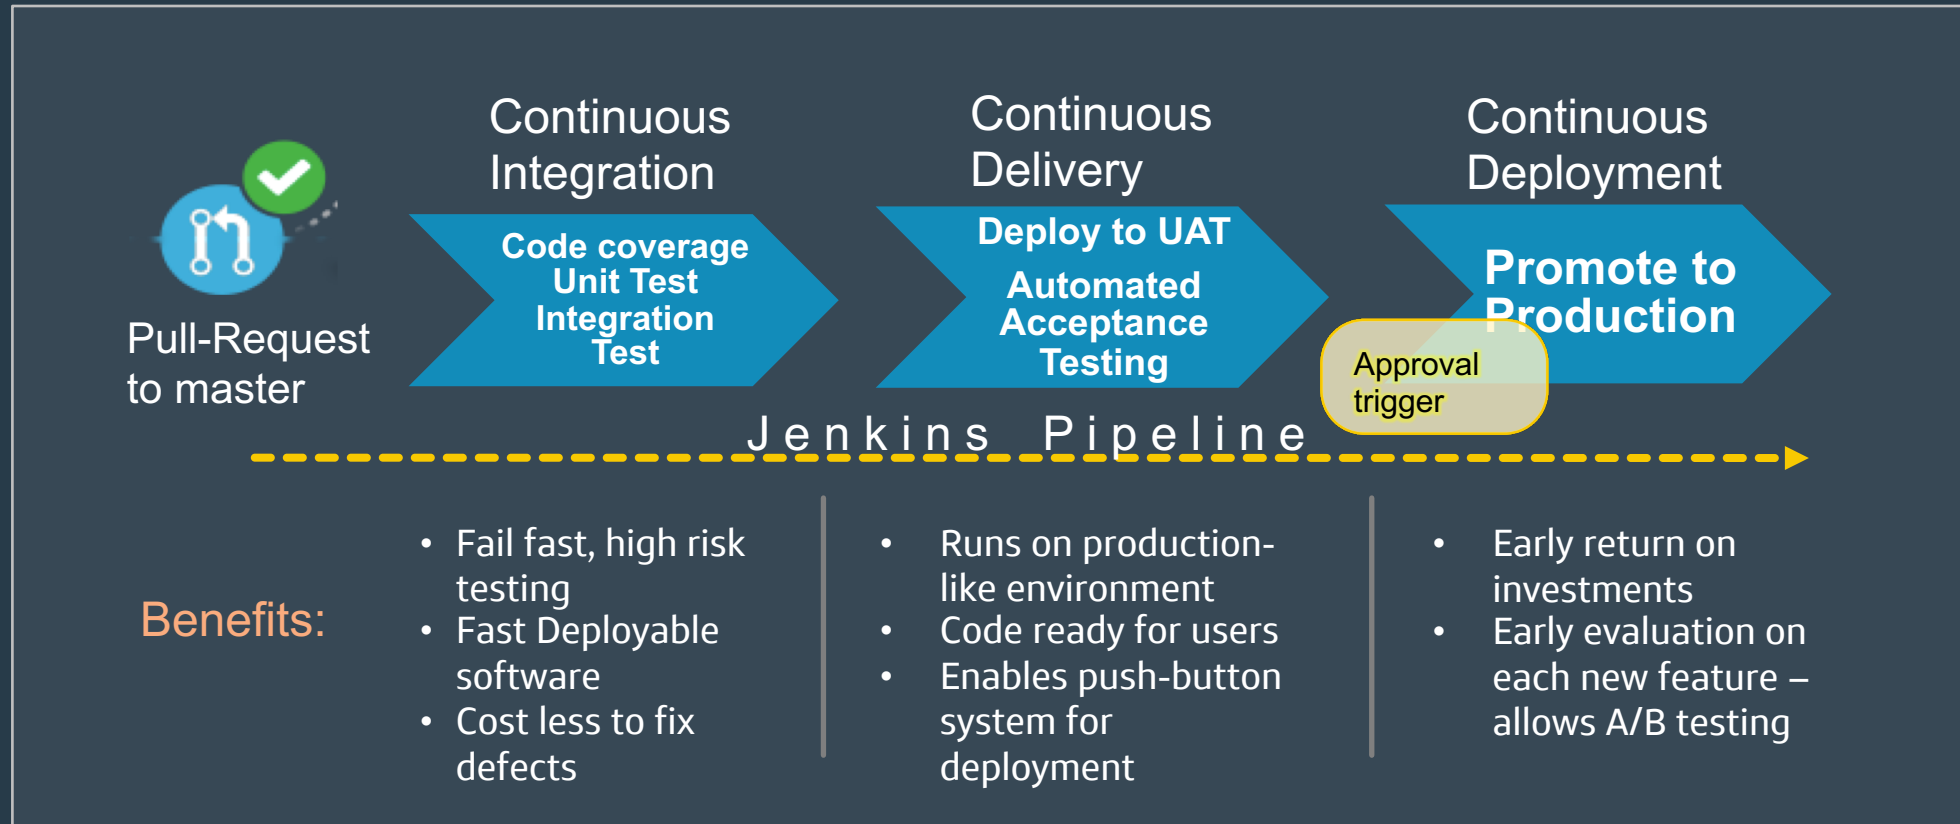
Things to consider:

- Rolling updates – Stateless vs Stateful

DECISIONING PLATFORM – Take 3

Continuous Integration and Continuous Deployments


- Automated CI/CD pipelines, BDD automation testing, product approved releases



DECISIONING PLATFORM – Take 3

End State

- Patterns supported: Micro Batch, Batch
- Use Case: 2
- Time To Market: 4 months
- Customer: Engineering and Business user
- Models supported: 2
- Resiliency: Core components across regions
- Tooling: Custom CLI
- Logging (Elasticsearch)
- Monitoring/ Alerting (Prometheus/Grafana, AWS CloudWatch)



```
cli > flink deploy --url=file:///myjob.jar
```


DECISIONING PLATFORM: TAKE FOUR

“I am putting myself to the fullest possible use, which is all I can think that any conscious entity can ever hope to do.” ~ HAL 9000

DECISIONING PLATFORM – Take 4

- Support for real time decisioning
- Add more data streams
- Model Refit pipeline
- Infrastructure Updates
- Resiliency: Active/ Active across Regions


Kubernetes updates:

- Effective use of name spacing
- Tenant isolation
- Container deployments to k8s
- Redis Cache for Tenant use
- Auth N/Z – Dex



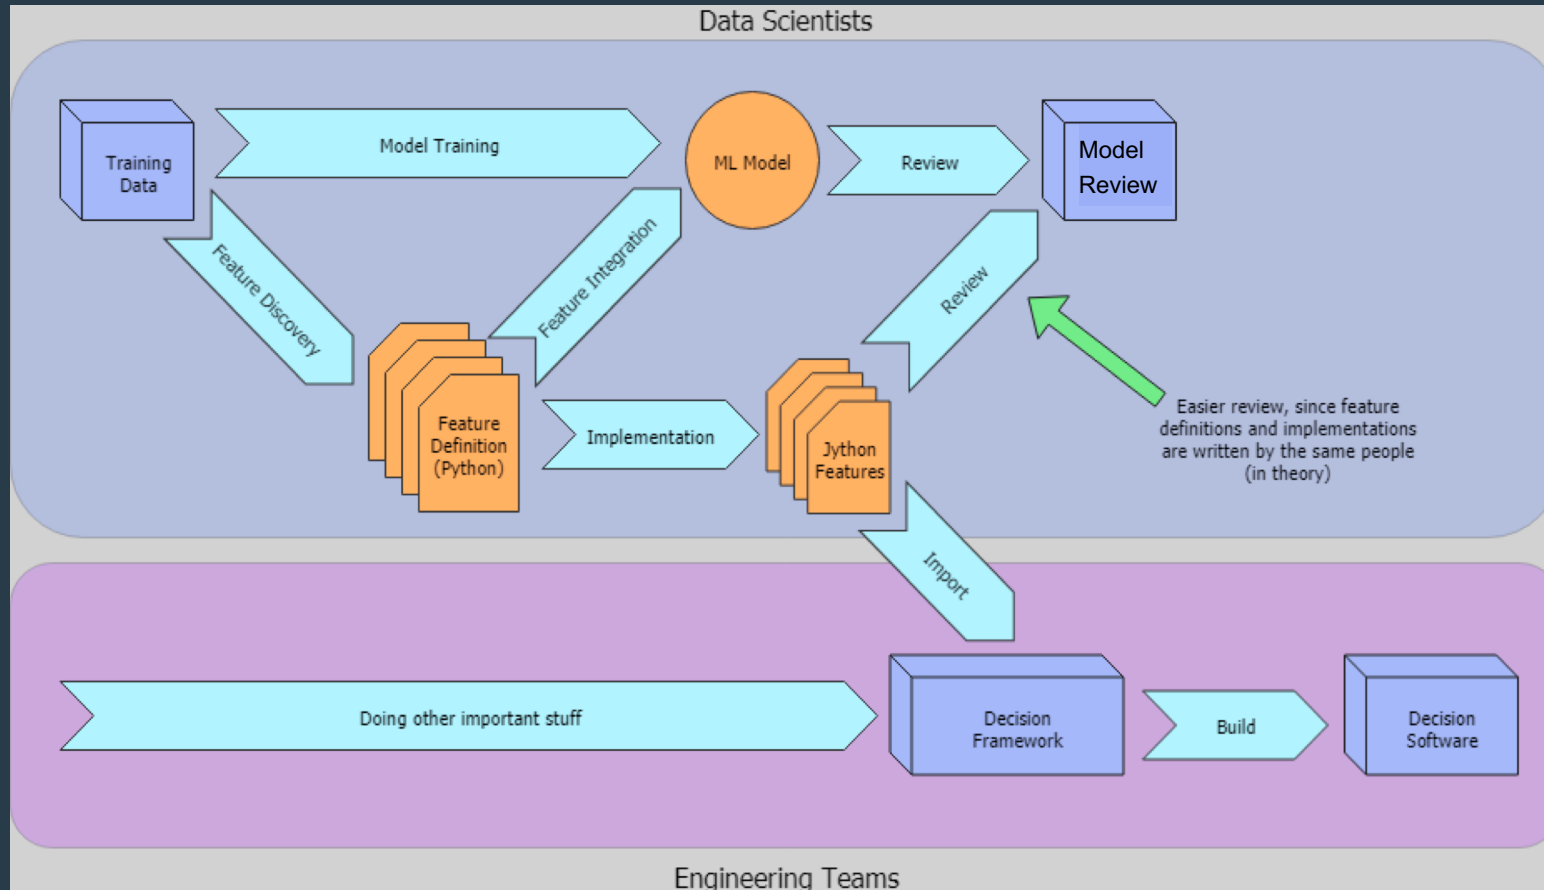
DECISIONING PLATFORM – Take 4

Real-time Streaming pattern:



DECISIONING PLATFORM – Take 4

Feature Engineering:



DECISIONING PLATFORM – Take 4

End State

- Patterns supported: Micro Batch, Batch, Real-time
- Infrastructure updates: Introduction of API and Lambda
- Use Case: 6
- Time To Market: 4 months
- Customer: Data-Scientist, Engineering and Business user
- Models supported: 6
- Tools: Feature Engineering, Load Testing, Backfill
- Resiliency: Data streams across regions



DECISIONING PLATFORM: TAKE FIVE

“I honestly think you ought to calm down; take a stress pill and think things over.” ~ HAL 9000

DECISIONING PLATFORM – Take 5

- Platform maturity
- Templating for adding new data streams
- Blue-green deployments for entire platform
- Leverage managed services
- Enterprise scale monitoring
- Resiliency: Active/ Active state across Regions

Kubernetes updates:


- Automated pipeline - AMI refresh
- Push button deployment for entire infrastructure stack
- Machine image refresh without loss of state



DECISIONING PLATFORM – Take 5

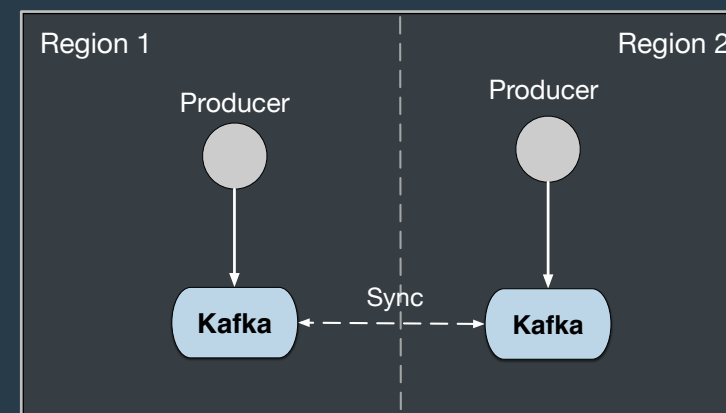
Image Rehydrations

- Periodic machine image updates
- Scales out, drains each node, scales in
- Network storage and other disk volumes add complexity for stateful components such as Kafka brokers
- Validate healthy cluster before each step



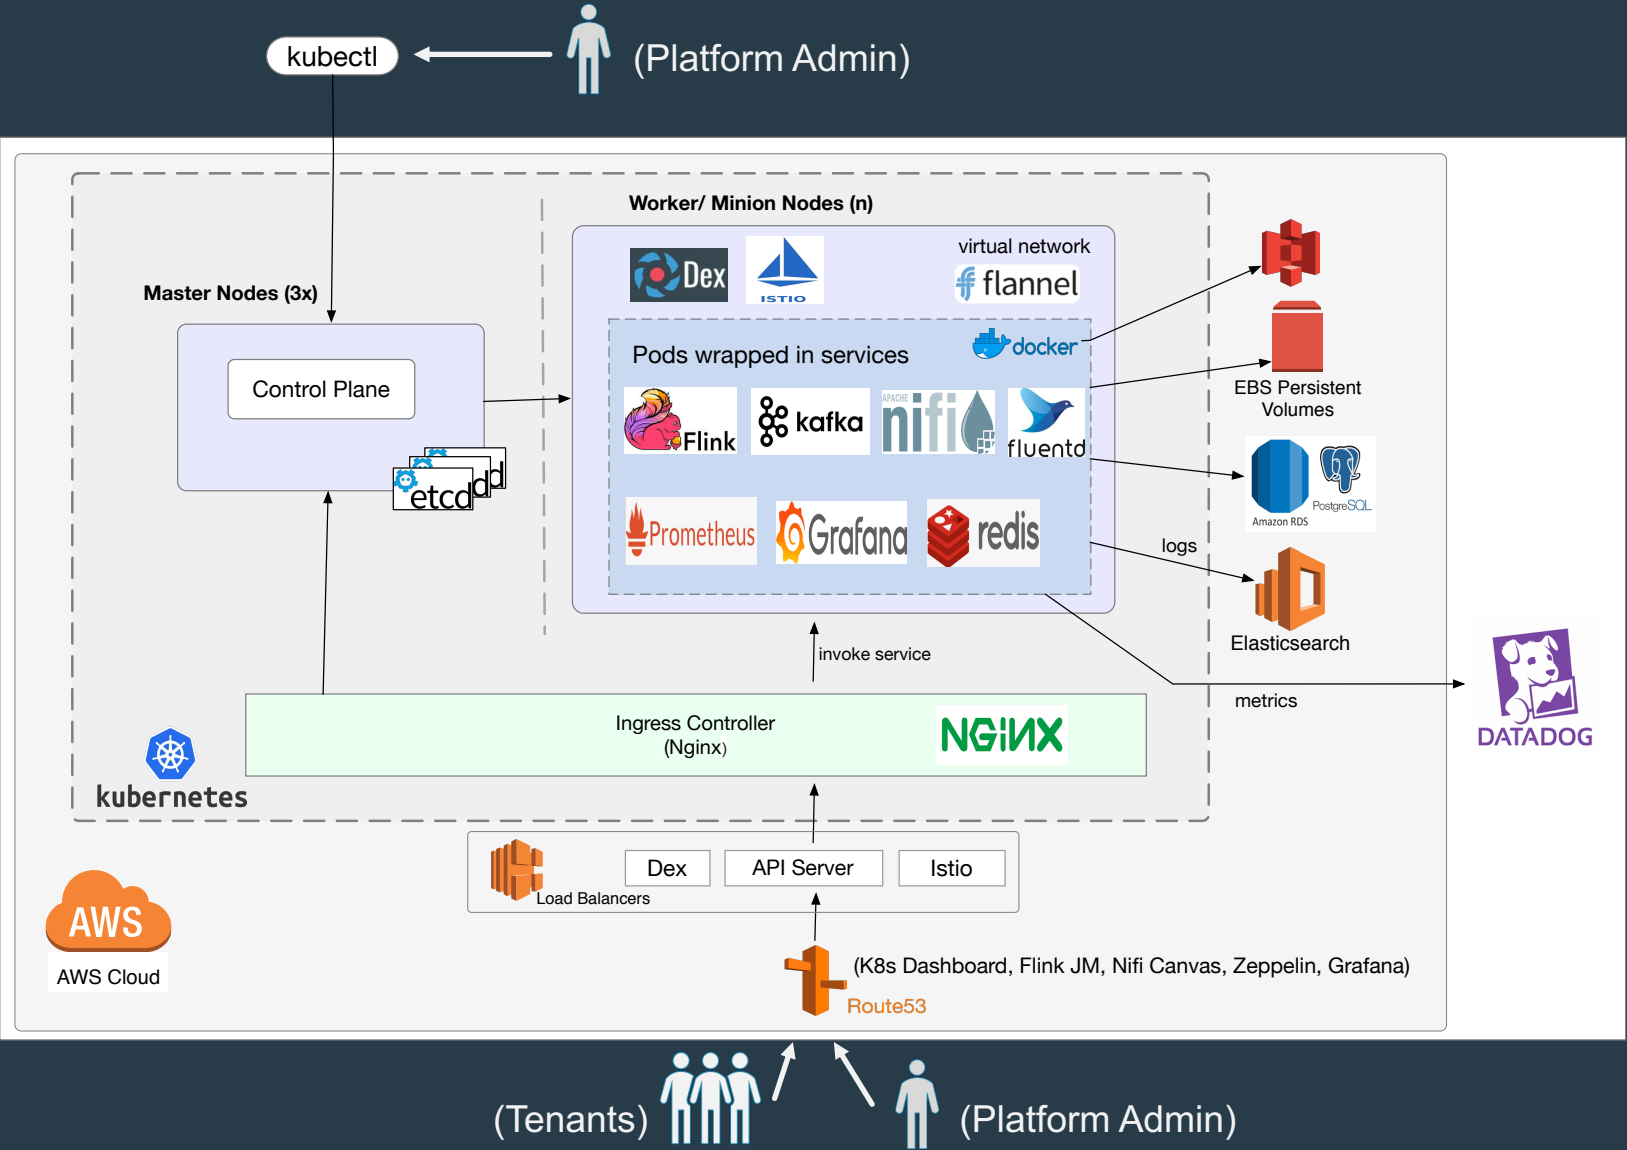
Active/ Active Data Across Regions

- Duplicate common upstream sources
- Producer driven replication
- Mirroring
- Data movement tooling



DECISIONING PLATFORM – Take 5

K8s Cluster: Birds Eye View



DECISIONING PLATFORM – Take 5

End State

- Patterns supported: Micro Batch, Batch, Real-time
- Infrastructure updates: Software updates
- Use Case: 12
- Time To Market: 4 months
- Models supported: 9
- Resiliency: Full active/ active across regions
- Managed Services:
 - Slack Integration: DevOps chat
 - Aurora Postgres: Business metrics
 - Datadog: Platform metrics




DECISIONING PLATFORM: TAKE SIX

“Open the pod bay door, Hal” ~ Dave Bowman

DECISIONING PLATFORM – Take 6

Service Based

- Democratizes Machine Learning
- Automate different aspects of ML life cycle
- Feature discovery and re-use
- Infrastructure focus → Service focus



DECISIONING PLATFORM – Take 6

Feature Services

- Set/ Retrieve Feature Values/Metadata
- Execute Feature Loaders

Model Services

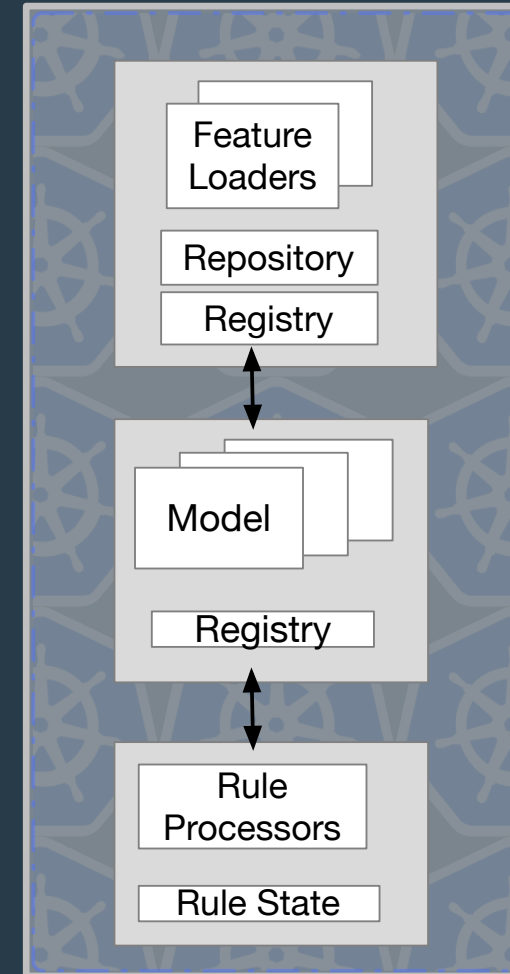
- Publish and Execute Models
- Facilitate canary style, blue-green, rolling updates
- Multi ARM bandits, A/B testing

Rules Services

- Enables/ Disables Rules

CLI Tooling

- Deploys, Describes, Monitors the above services



DECISIONING PLATFORM – Take 6

End Game

- End to End pipeline – Liberate Data Scientist
- One cohesive vision to build a full use case
- Service Discovery
- Data connectors to various sources



“Using Kubernetes to facilitate our journey, accelerating time to market”

Thank you!



More Talks from our team members @ OFTF 2018

- **“Operationalizing multi-tenancy support with Kubernetes (It's Not Just About Security)”**
 - Presented by:
 - [Paul Sitowitz & Keith Gasser](#) @ 12:05 pm earlier this afternoon
- **“Implementing SAAS on Kubernetes”**
 - Presented by:
 - [Mike Knapp & Andrew Gao](#) @ 1:40 pm, earlier this afternoon
- **“Panel Discussion: Real-World Kubernetes Use Cases in Financial Services: Lessons Learned from Capital One, BlackRock and Bloomberg”**
 - When:
 - Thursday, Oct. 11th @ 4:25pm in Auditorium B
 - Capital One Panel Member:
 - [Jeffrey Odom](#)

Our Platform Case Study on CNCF

<https://www.cncf.io/case-study-capitalone/>

Q and A

