

Database Synergy Day 2025: Swiss Oracle & PostgreSQL User Groups Unite

ORACLE



# Open Source services & PostgreSQL on OCI

Live Demo


**Piotr Kurzynoga**

Principal Solution Engineer – Open Source Data Platform





# ABOUT ME



Degree: Electronic and Computer Engineering



SAP KDE for Partner Developed Solutions in SAP Cloud for Customer & SAP Business ByDesign



Oracle Solution Engineer since October 2020  
Focus on: Data Management / Data Development



<https://www.linkedin.com/in/piotr-kurzynoga/>



<https://medium.com/@devpiotrek>

**SAILGP™**

ORACLE  
CloudWorld



**POUG**

# Agenda

## Presentation






















- Oracle and Open-Source
  - That can't be right...
- Managed or not?
  - Kubernetes vs DBaaS.
- Demo Architecture
  - The bigger picture.

## Demo

- Machine learning pipeline, streaming data into PostgreSQL with Apache Spark.
- Text2SQL Chatbot with PostgreSQL.

# Open Source & Oracle Cloud Infrastructure (OCI)

# OCI Managed Open Source Data and AI services
















Ingestion, Event Streaming	Big data and ETL	Build, train, deploy ML models	Generative AI and RAG	High-Speed Caching	Search, Analytics and Vectors	Combine OLAP and OLTP	Extensible SQL, Vector database	NoSQL Distributed Database
								
OCI Streaming with Apache Kafka	OCI Data Flow	OCI Data Science	OCI Generative AI (LLM, Agents)	OCI Cache	OCI Search with OpenSearch	Oracle Heatwave MySQL	OCI Database with PostgreSQL	OCI NoSQL Database
Fully managed, highly available service available in all OCI regions								
<b>Distributed event streaming</b> High-performance data pipelines, streaming analytics, data integration	<b>Serverless Spark at scale</b> Big data processing, ETL, interactive SQL and Spark streaming at cloud scale	<b>Machine Learning (ML) and MLOps</b> Use Python and open source tools to build, train and deploy ML models at scale	<b>Multiple LLMs with dedicated AI clusters</b> Choice of models, content moderation, flexible fine tuning with dedicated AI clusters	<b>Flexible memory configuration</b> Enable sub-millisecond application response time	<b>Hybrid search, observability and security</b> Zero down time scaling and high availability	<b>In-memory query accelerator</b> Improve query performance, and get real-time analytics without ETL	<b>Optimized pay as you grow storage</b> High performance, scale with pay-per-usage and achieve zero RPO & HA	<b>Optimized low pay-per-use serverless NoSQL table</b> Low-latency access to global data replicated to multiple OCI regions
			  	 				



# Leverage a wide range of existing technology investments with OCI

## Use OCI managed open-source services

Simplifies the migration of existing applications










 <b>kubernetes</b> MANAGED CONTAINERS	 <b>docker</b> CONTAINERS	 <b>OpenTofu</b> INFRASTRUCTURE AS CODE	 AUTONOMOUS LINUX
 DATABASE	 BIG DATA	 ANALYTICS ENGINE	 CACHING DATABASE
 DATA COLLECTION	 EVENT DELIVERY	 STREAMING	 DATABASE
 SERVERLESS PLATFORM	 SEARCH AND ANALYTICS	 INTERFACE DEFINITION	

## Use the technologies, dev tools, and integrations you're familiar with

Extends existing investments

 Linux OS Red Hat Oracle Linux Ubuntu CentOS Debian SUSE	 WINDOWS SERVER OS	 VIRTUAL ENVIRONMENT
 PROGRAMMING LANGUAGES	 MACHINE LEARNING FRAMEWORKS	 APPLICATION PLATFORM
 PROGRAMMING LANGUAGES	 MACHINE LEARNING FRAMEWORKS	
 DATABASE	 APPLICATION FRAMEWORK	 NOSQL DATABASE
 APPLICATION FRAMEWORK	 BLOCKCHAIN	 AI PLATFORM

Simplifies application build, deploy, and manage with skills you already have


 DEV-OPS	 AUTOMATION
 DEV-SEC-OPS	 AUTOMATION
 CONTAINER MANAGEMENT	 PACKAGE MANAGEMENT
 INFRASTRUCTURE AS CODE	 TEAM COLLABORATION
	 RISK MANAGEMENT





# Managed or Not?

# Get more with OCI Database with PostgreSQL



## Fully Managed:

- Automated deployment and patching
- Reader endpoints
- Metrics/Logging



## Scalable:

- Scales vertically with minimal downtime
- Scales horizontally with zero downtime
- Up to 750K IOPS\*
- Up to 64 OCPU's & 1024GB RAM



## Storage:

- Automatic shrink & expand
- Cross-region backup



## Highly Available:

- SLA is 99.99%
- Data distributed across multiple Fault Domains (single AD)



## Secure:

- End-to-end encryption
- Automatic security patching

**No vendor owned extensions!**

\*(Over 16 OCPU's)



# Manageability vs Flexibility – Managed services and Kubernetes





## DBaaS PostgreSQL

- Easy deploy & connect
- Separate scaling mechanism
- Security & Patching
- Cloud based

## K8s hosted PostgreSQL

- Configuration required
- In K8 wholesome cluster scaling
- OS Level Access & maintenance
- Environment agnostic

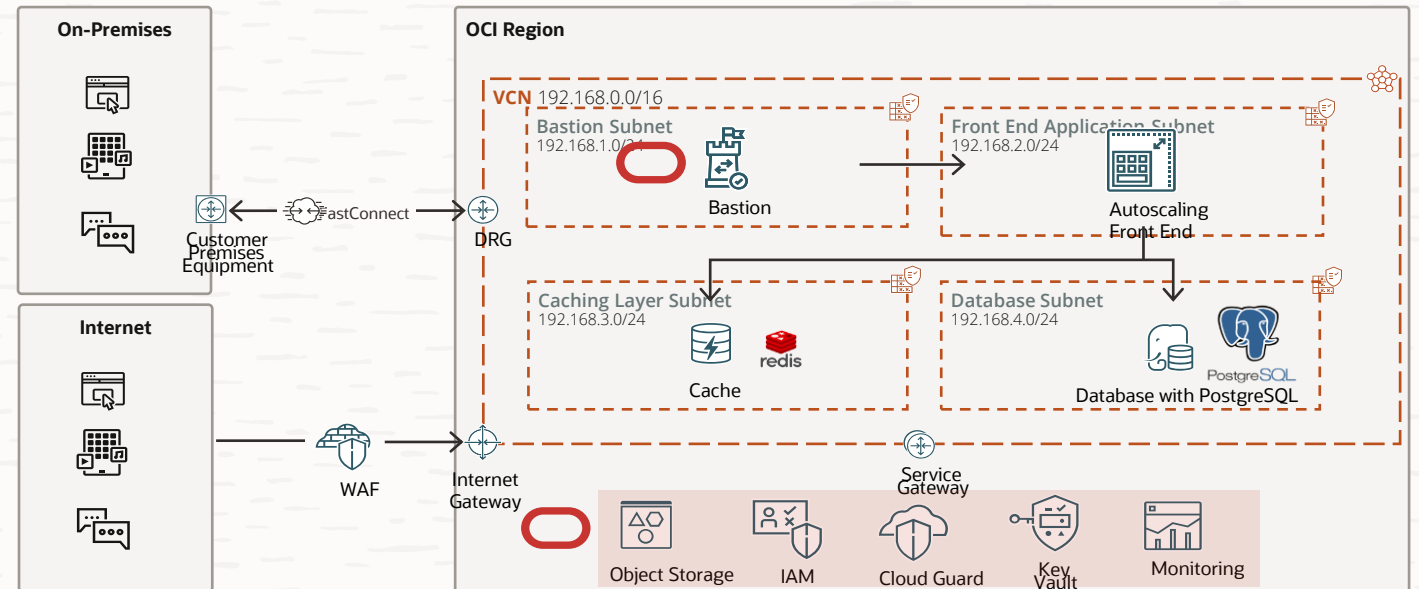


Ease of management

Flexibility

# Build and secure your web application

- Well-suited for applications requiring fast data access and complex transactions, such as e-commerce platforms
- With secure bastion to gain restricted and time-limited access to cloud resources without public-facing endpoints

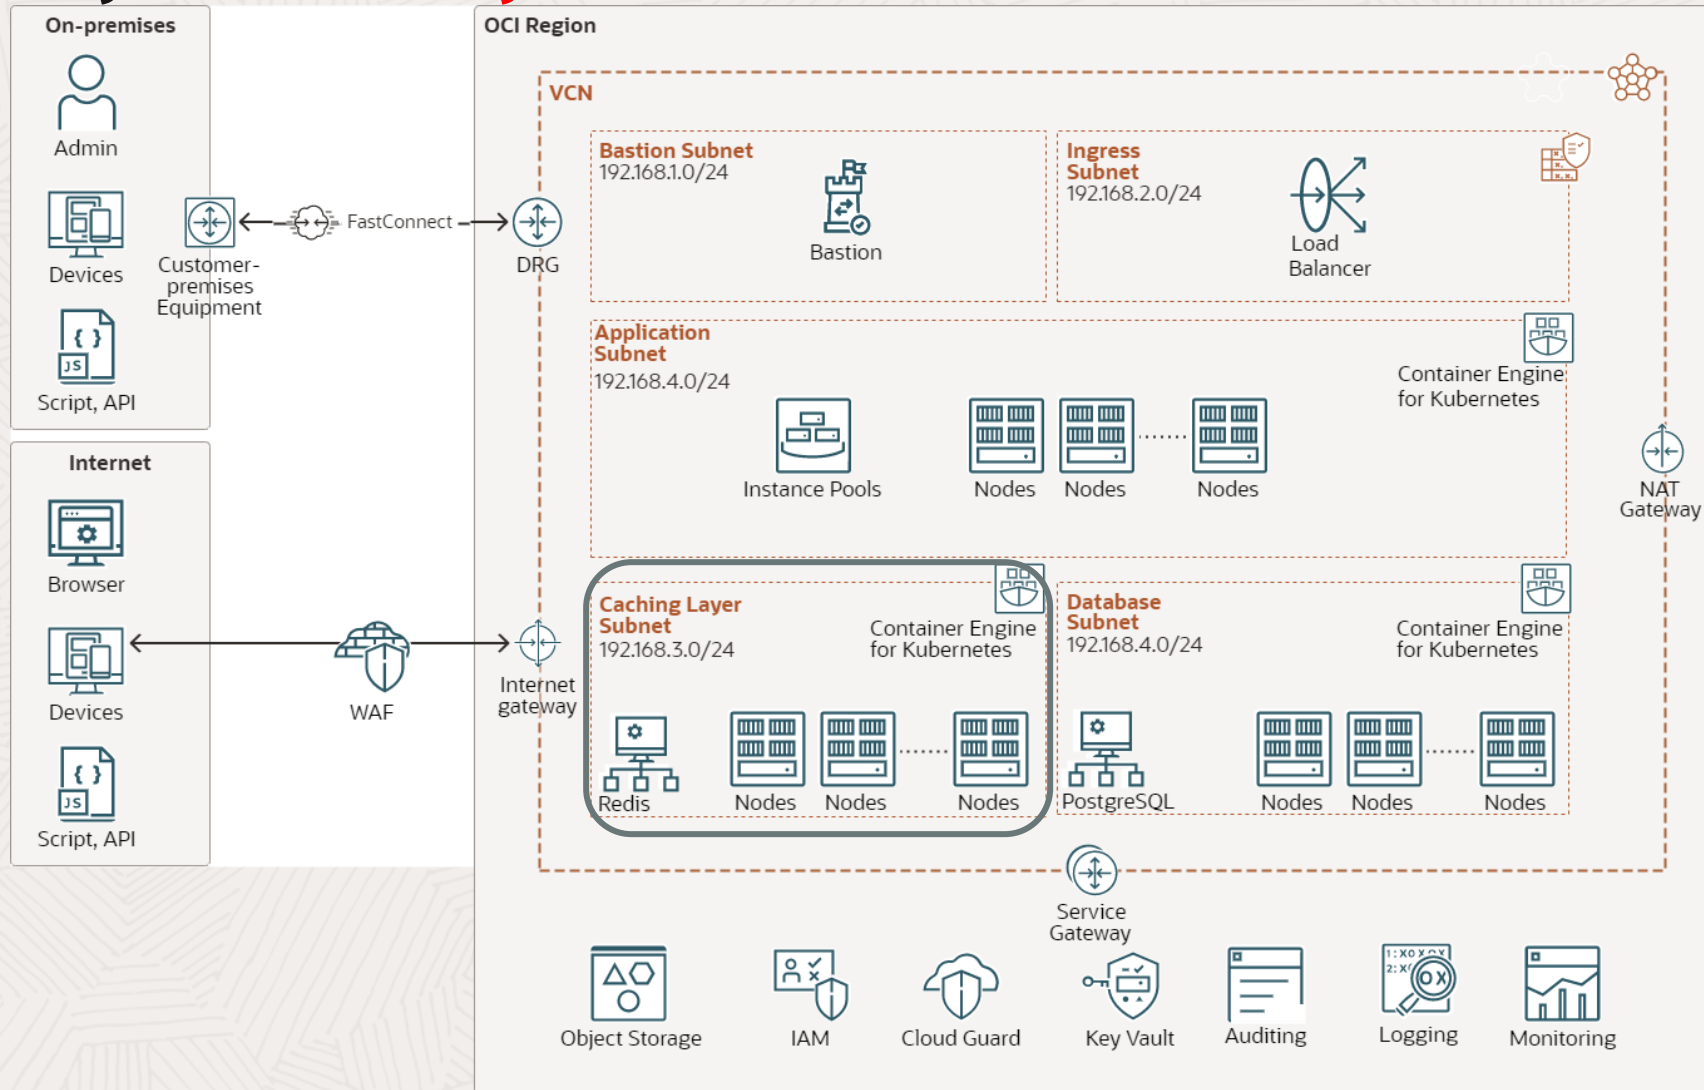


Full

Split

Minimal

# Manageability vs Flexibility

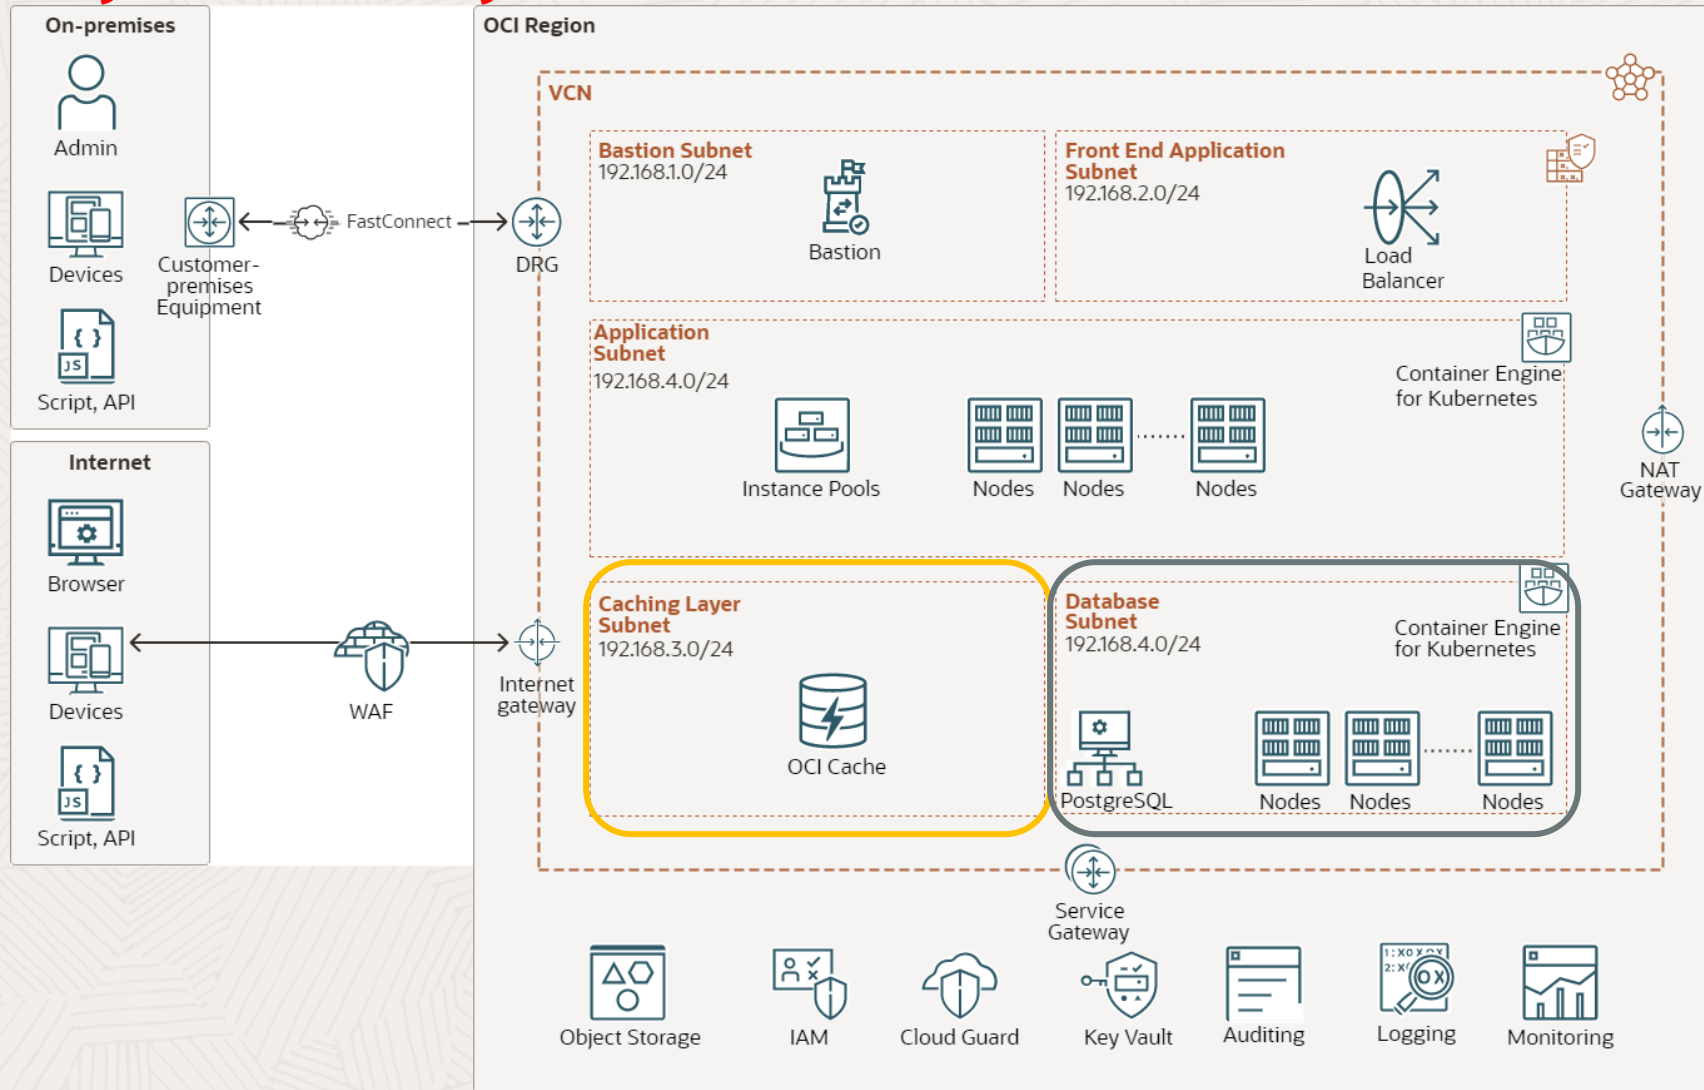


Full

Split

Minimal

# Manageability & Flexibility

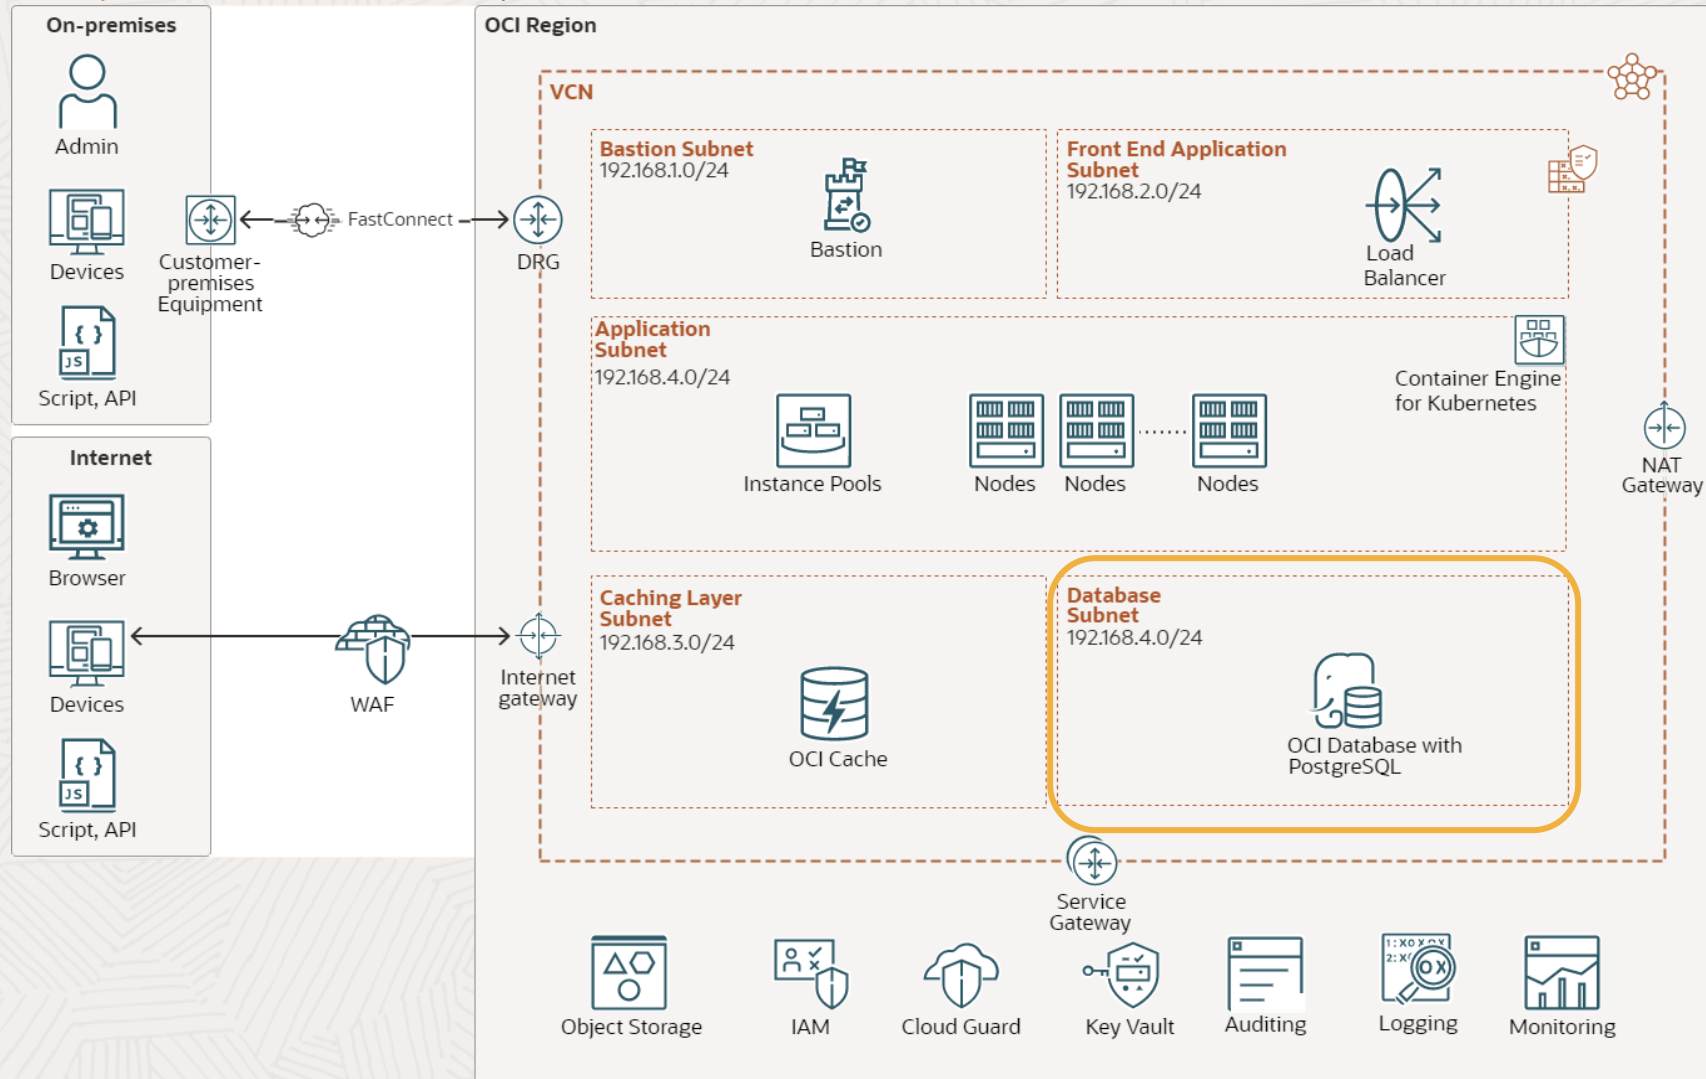


Full

Split

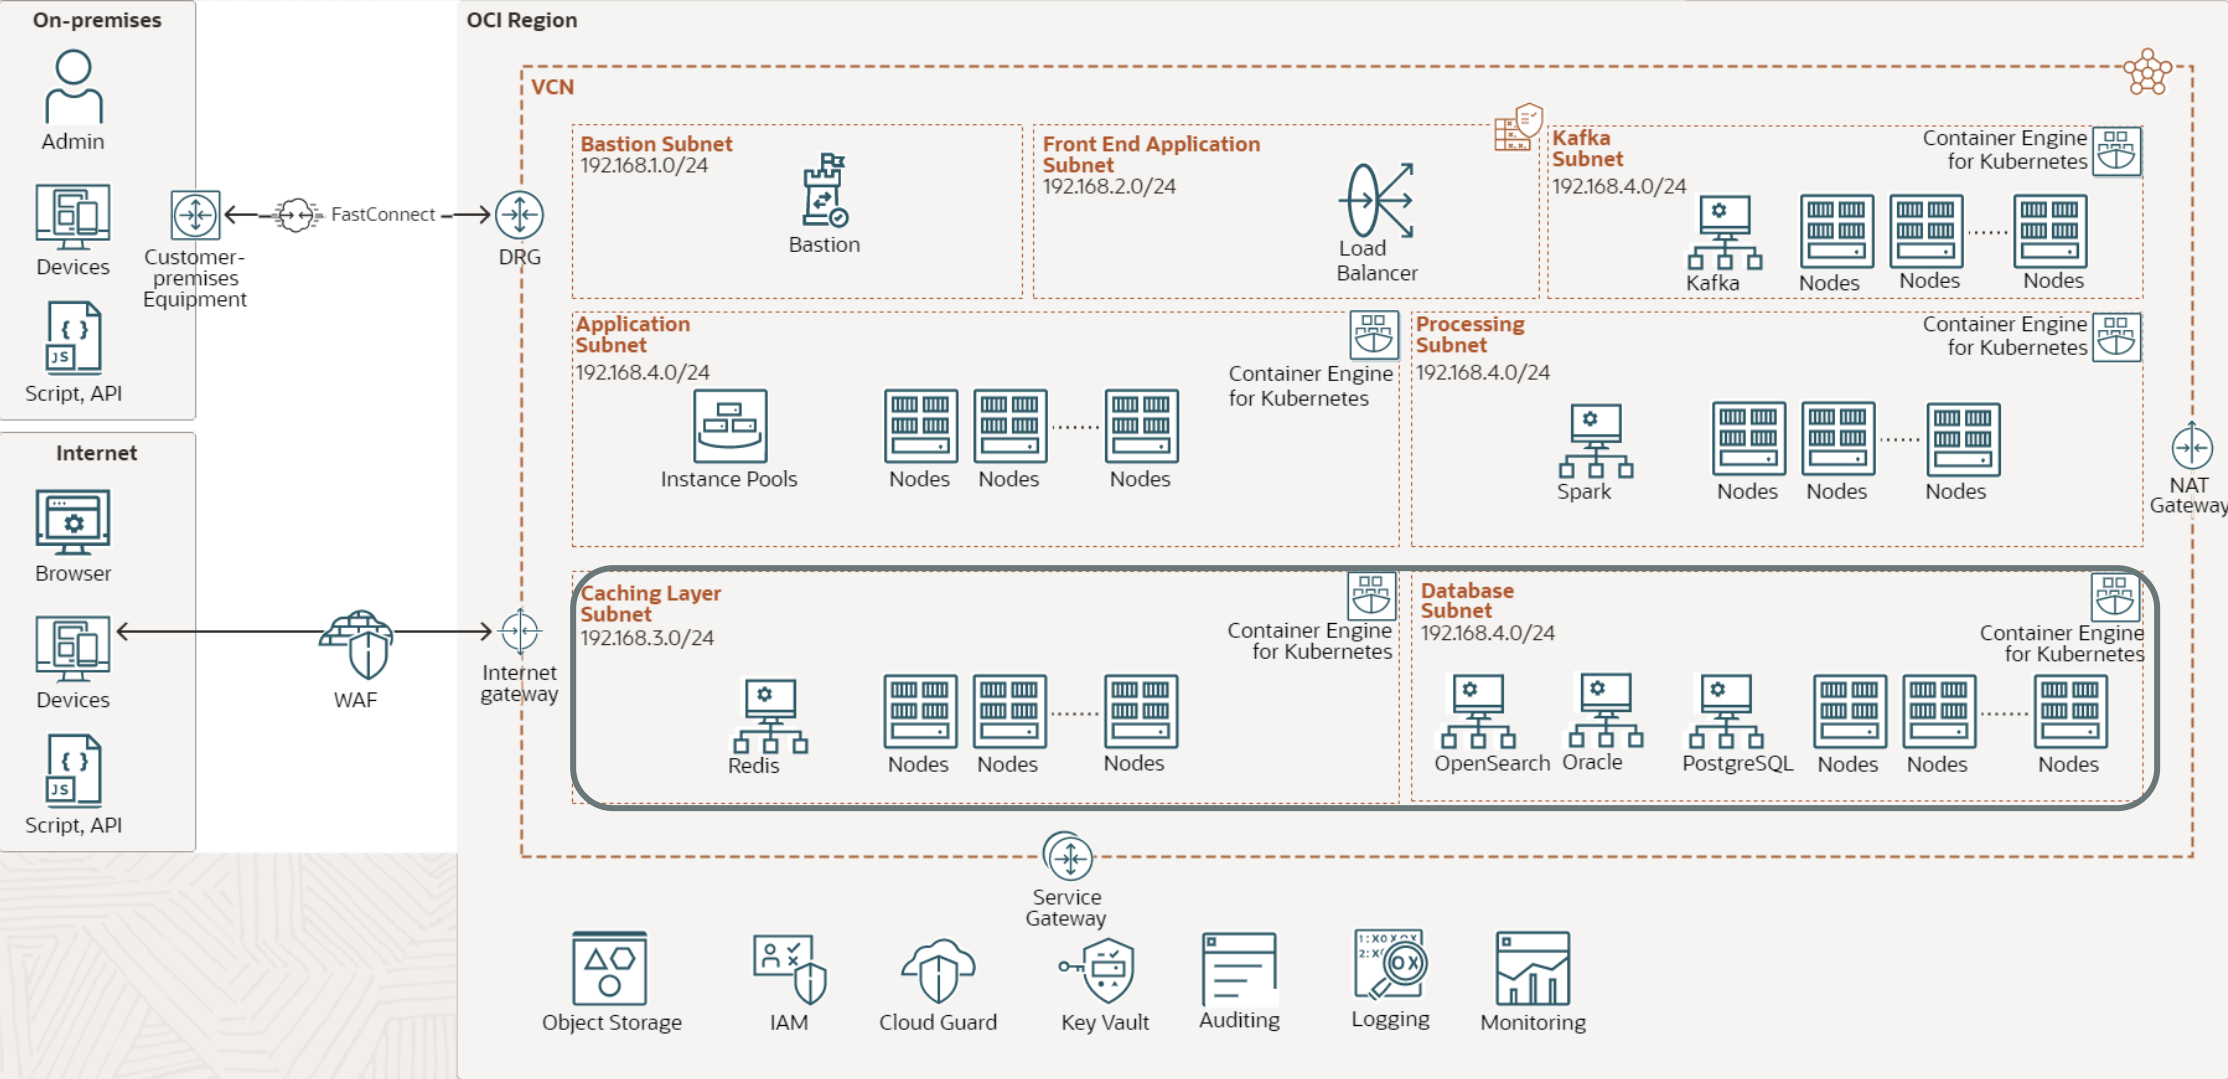
Minimal

# Manageability vs Flexibility

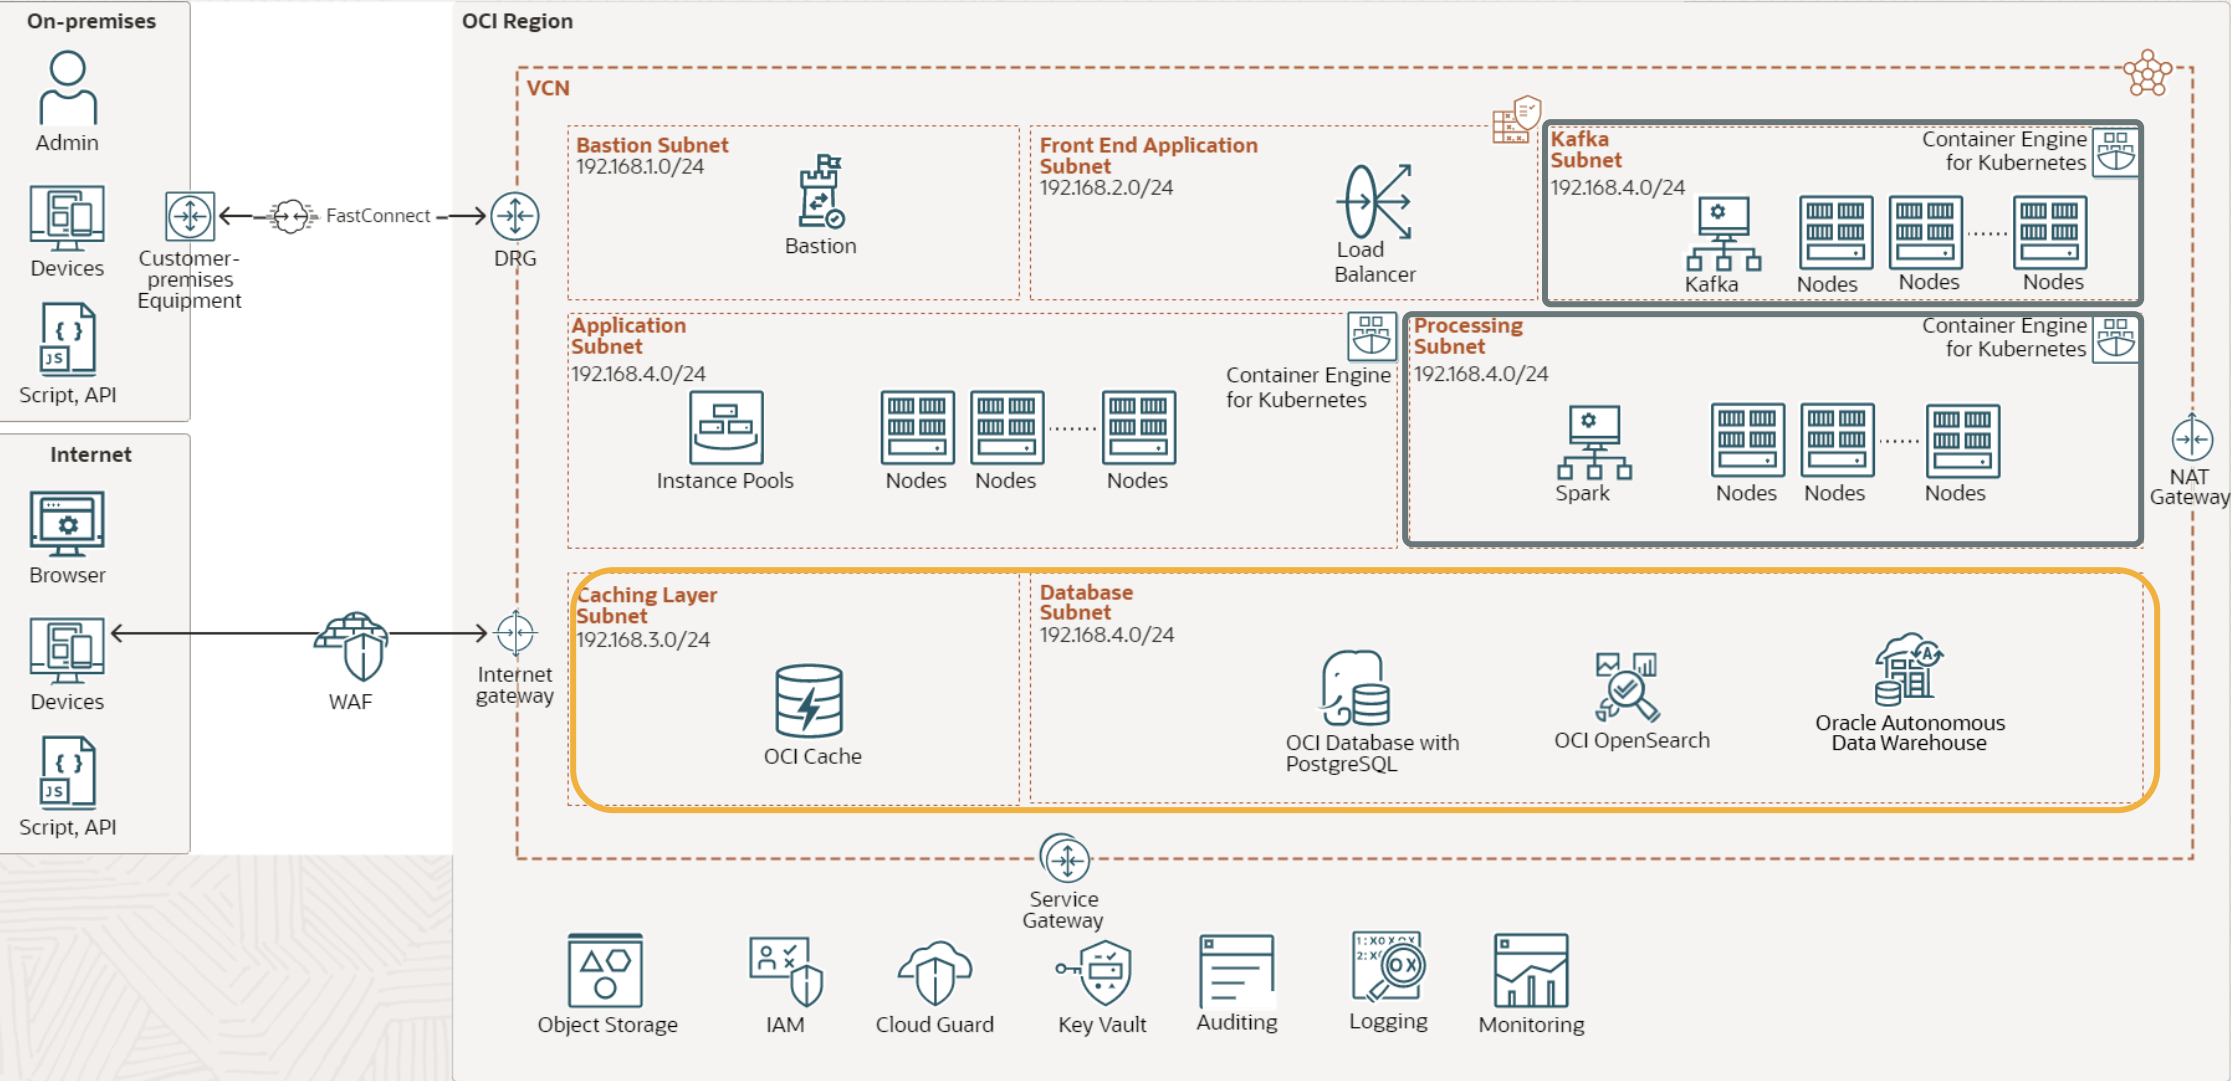




# Manageability vs Flexibility

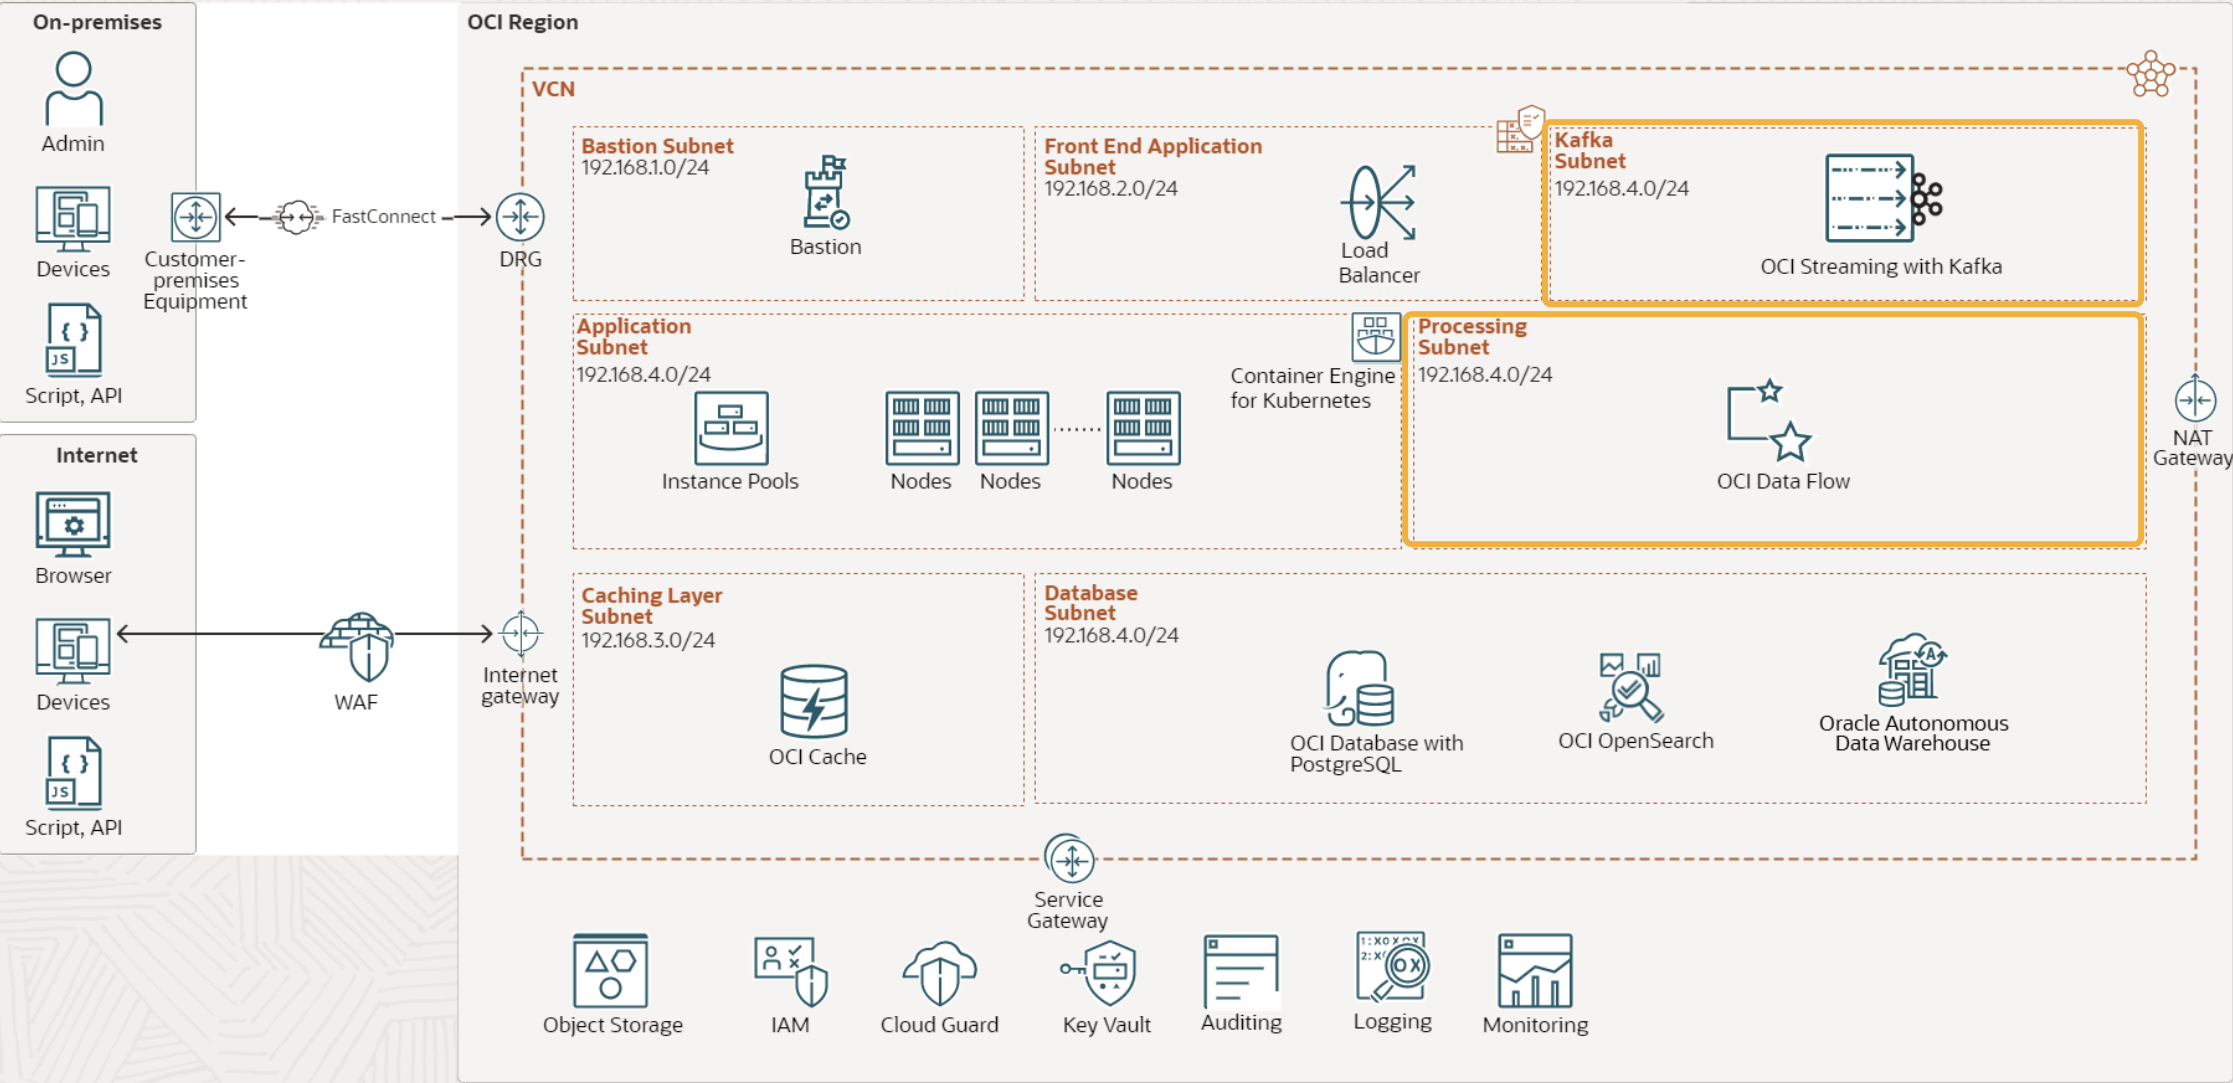


# Manageability & Flexibility



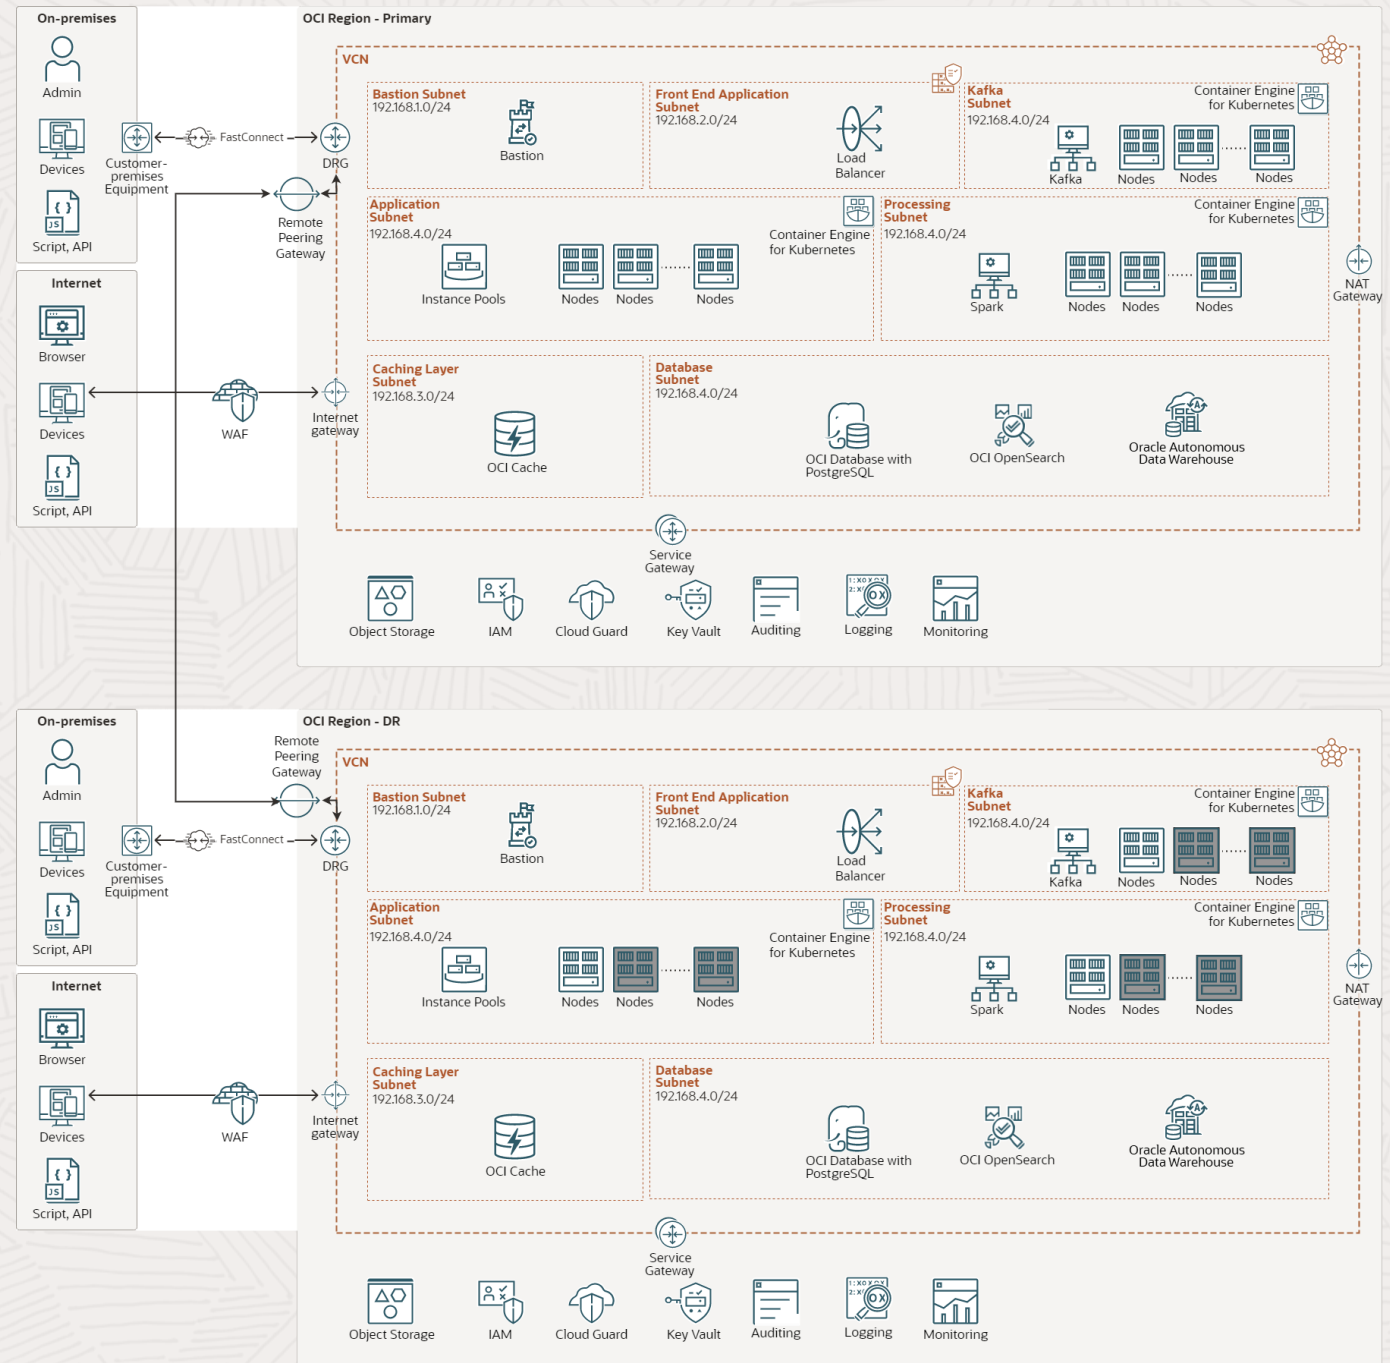


# Manageability vs Flexibility

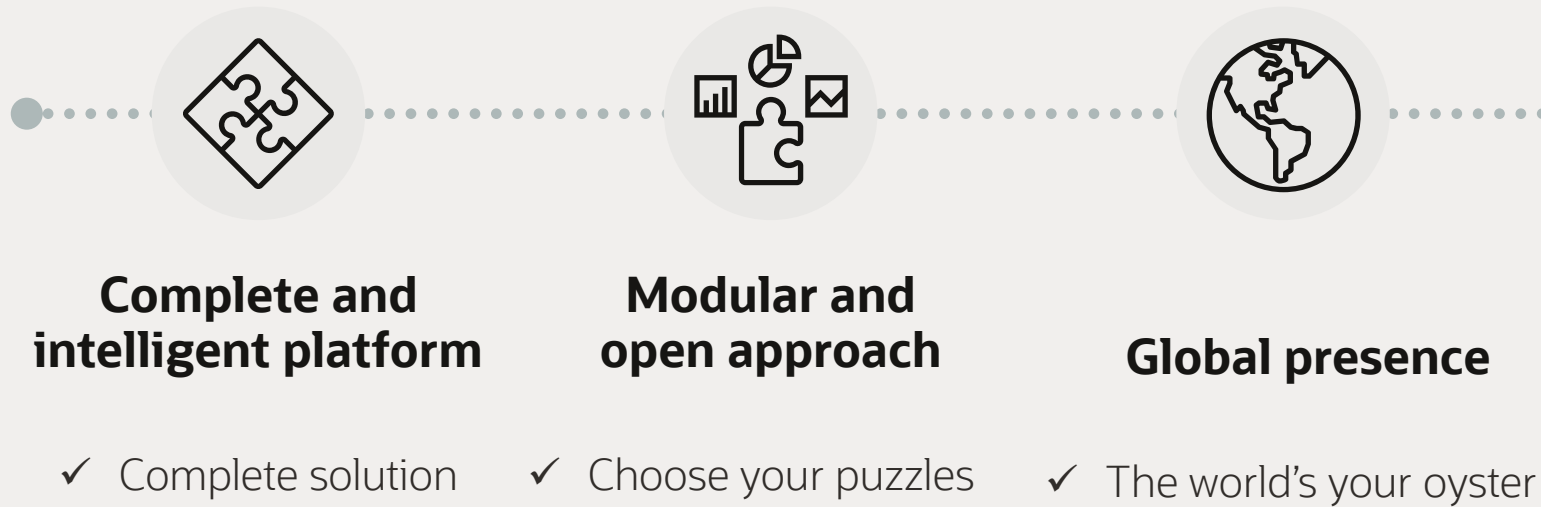


# What about DR?

- Maintaining pipelines
- Cross-region replication
- Globally distributed applications



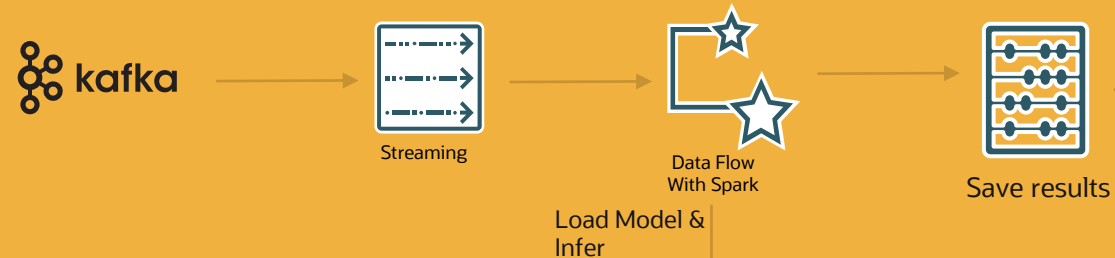
# Ultimately...



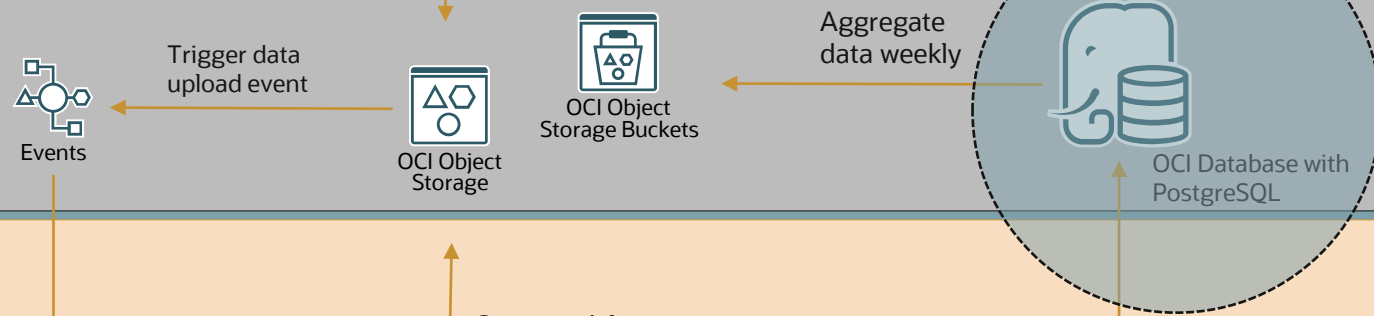
# Demo Time

# Solution Overview

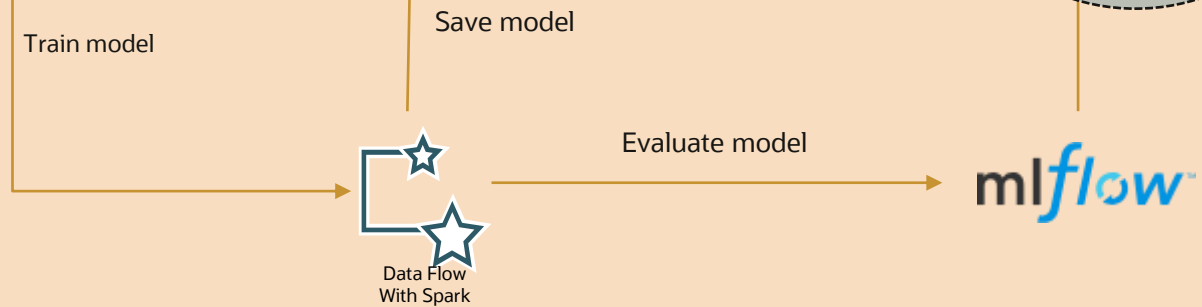
## Inference Pipeline - Streaming



## Common Zone



## Training Pipeline - Batch





Default >

receptive-slug-50

OverviewModel metricsSystem metricsTracesArtifacts

Description

No description

Details

Created at	03/17/2025, 02:23:37 PM
Created by	dataflow
Experiment ID	0
Status	Finished
Run ID	82c34a7cc8864113884efda07bdaa1c5
Duration	26.6s
Datasets used	—
Tags	<a href="#">Add tags</a>
Source	ML_Train.py
Logged models	spark
Registered models	mpg_estimator v1
Registered prompts	—

Parameters (2)

Search parameters	
Parameter	Value
labelCol	mpg
numTrees	50

mlflow

Tables

alembic_version	24K
datasets	32K
experiment_tags	16K
experiments	48K
input_tags	16K
inputs	24K
latest_metrics	16K
metrics	16K
model_version_tags	16K
model_versions	16K
params	24K
registered_model_aliases	16K
registered_model_tags	16K
registered_models	16K

Database Navigator

Enter a part of object name here

System Info	
Roles	
zabbix	
mlflow	
Schemas	
public	
mlflow	
Tables	
alembic_version	24K
datasets	32K
experiment_tags	16K
experiments	48K
input_tags	16K
inputs	24K
latest_metrics	40K
metrics	40K
Columns	
Constraints	
Foreign Keys	


metrics

Enter a SQL expression to filter results (use Ctrl+Space)

	key	value	timestamp	run_uuid	step	is_nan
1	rmse	4.3634829814	1,742,217,829,832	82c34a7cc8864113884efda07bdaa1c5	0	[ ]
2	r2	0.7341970187	1,742,217,829,940	82c34a7cc8864113884efda07bdaa1c5	0	[ ]

# Chatbot – End User Experience

- Ask questions in natural language.
- Fetch data from relational tables – no embedding required.
- Preset questions for quicker interaction.
- Show SQL (Optional) – Preview generated SQL.



## OCI SQL Agent Chat for Postgres

Enter your query:

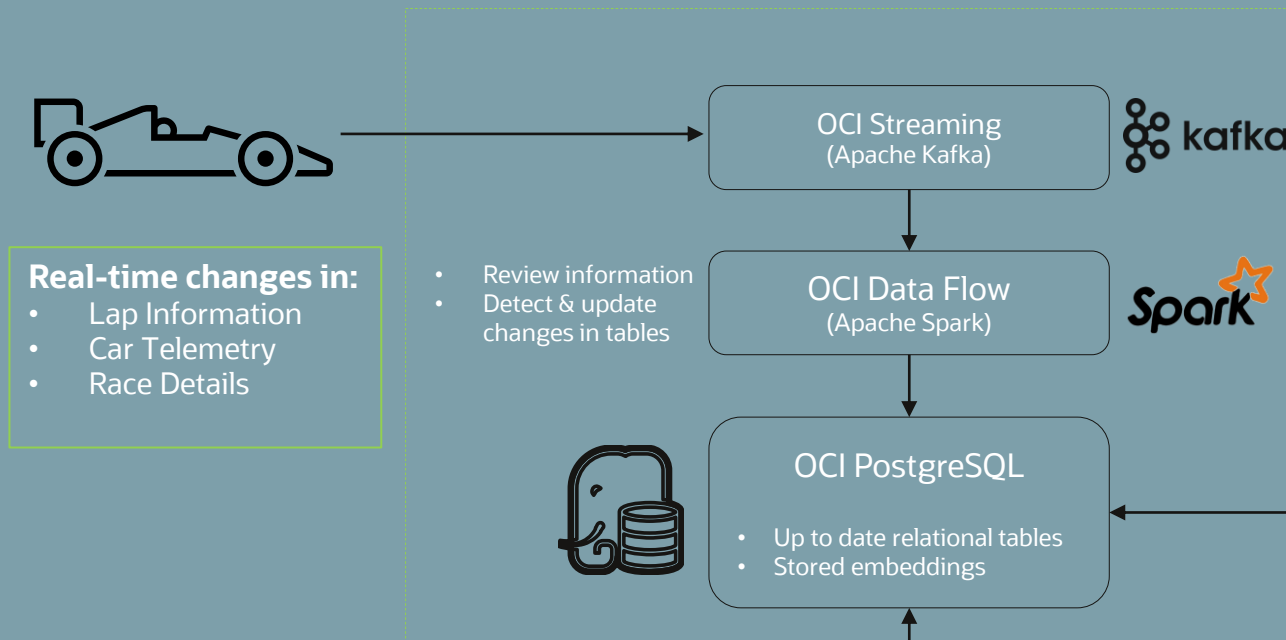
show me the prediction and normal consumption for ford torino

Run Query

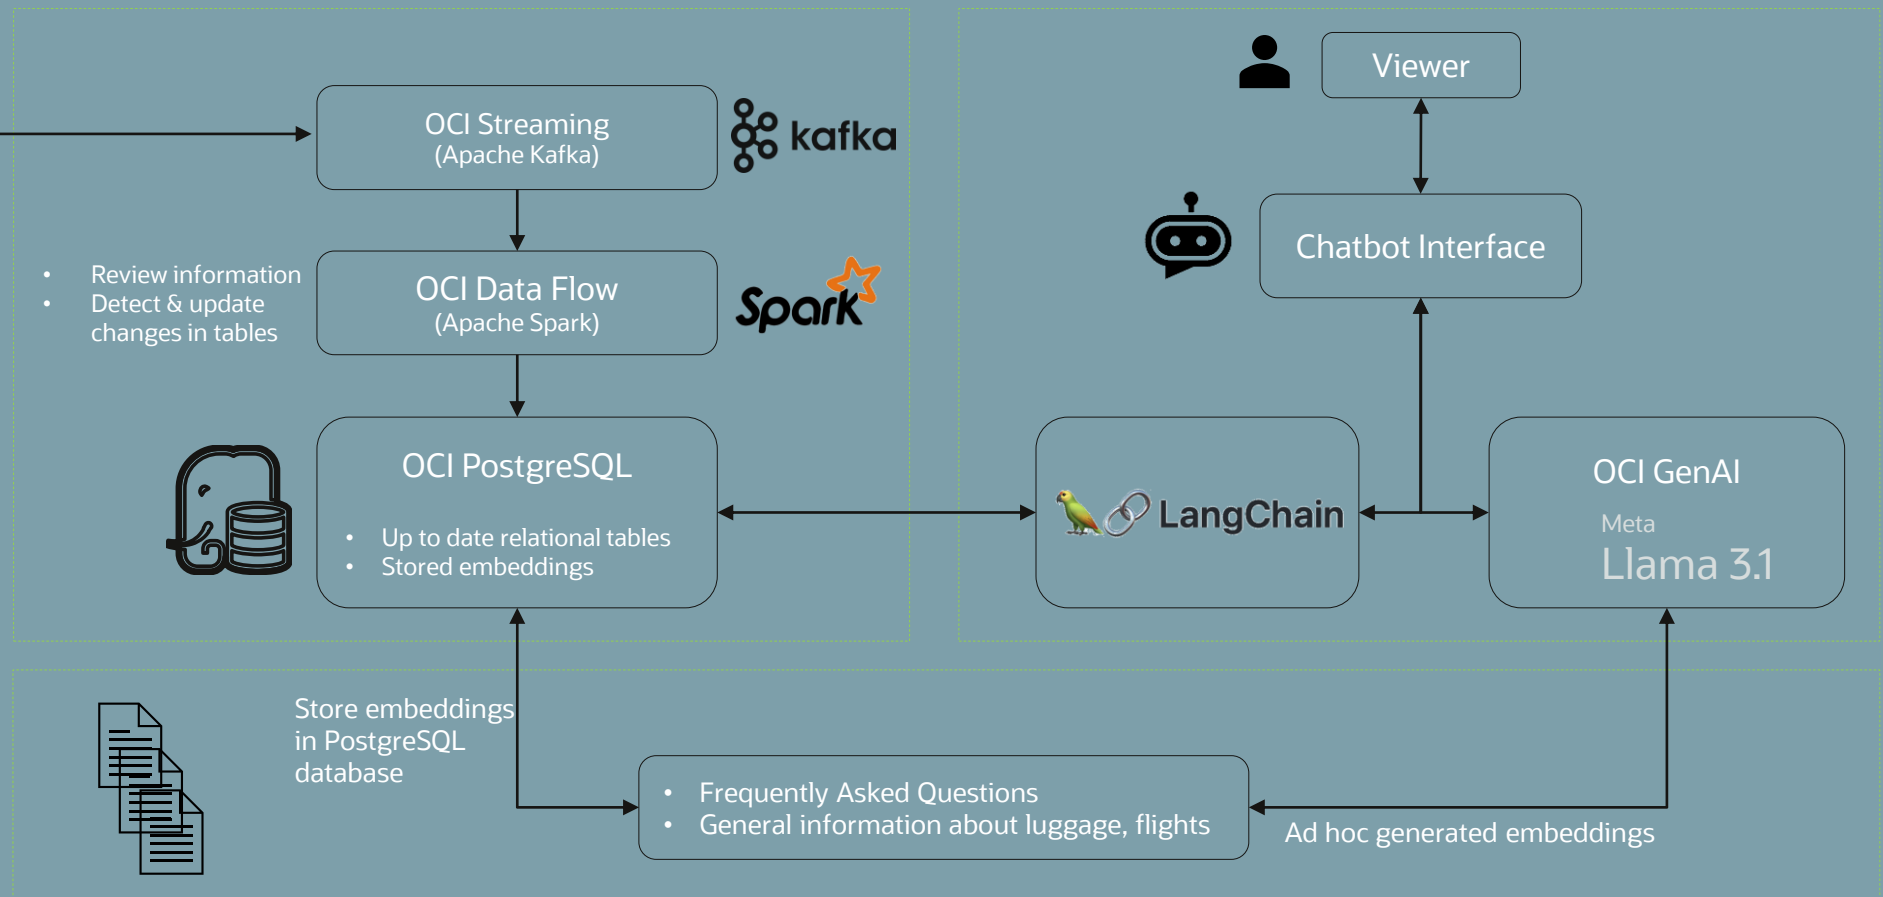


# Real-time Retrieval Augmented Generation in F1!

## Real-time Information Processing



## Retrieval Augmented Generation



# Useful Resources

# Find out more

## Examples and documentation

[Open-Source Data Platform Examples on Github](#)

Create your own project today with our examples!

[OCI GenAI Reference Architecture](#)

Discover the magic behind SQL Agent.

[Open Source Projects @ Oracle](#)

See how Oracle is contributing to Open-Source projects.

[Managed Open Source Services on OCI](#)

Checkout how Oracle Cloud Infrastructure can empower your projects with our managed services.



# Thank you

---

**Piotr Kurzynoga**

**[piotr.kurzynoga@oracle.com](mailto:piotr.kurzynoga@oracle.com)**

