

Cloud Agnostic Multi-Tenant SaaS Applications - Challenges & Solutions

Abhay Dutt Paroha

Software Team Leader, SLB


Conf42



A word about me

- 3+ years managing engineering teams working on real-time upstream operational data ingestion and delivery
- 14 years of experience in various aspects of the oilfield (operational and technical)

*Disclaimer: This session is not about cloud infrastructure; it is about using cloud to build applications.


An illustration on the left side of the slide shows a person from behind, wearing a dark blue shirt, looking up at a large, stylized cloud. The cloud is rendered in shades of orange, yellow, and blue, with several white lines extending downwards from it, suggesting data or connectivity. The background is a gradient of blue and orange, with other smaller clouds scattered around.

What is Cloud Agnostic?

Applications, Services, Systems, Tools, and Workloads that are designed to be compatible with multiple cloud providers, rather than cloud-native being tied to a specific cloud platform

Why *Cloud*?

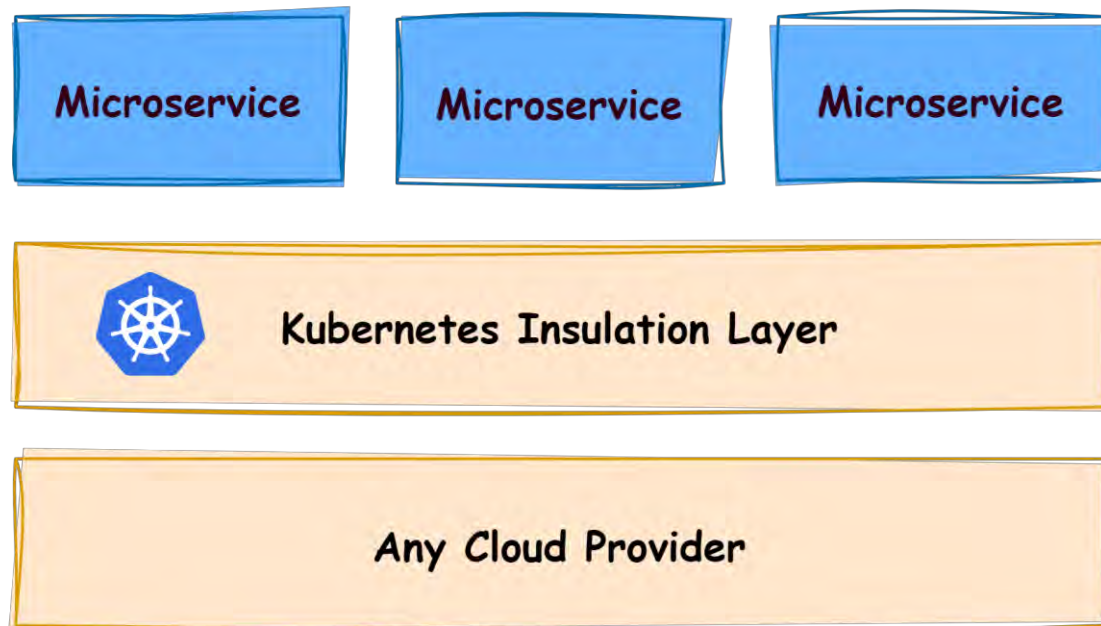
By 2028, cloud computing will shift from being a technology disruptor to becoming a necessary component for maintaining business competitiveness - *Gartner*



Cloud agnostic pros and cons

- Pros
 - Avoid the risk of vendor lock-in
 - Performance
 - Flexibility
 - Resilience
- Cons
 - Challenging implementations
 - Time to market

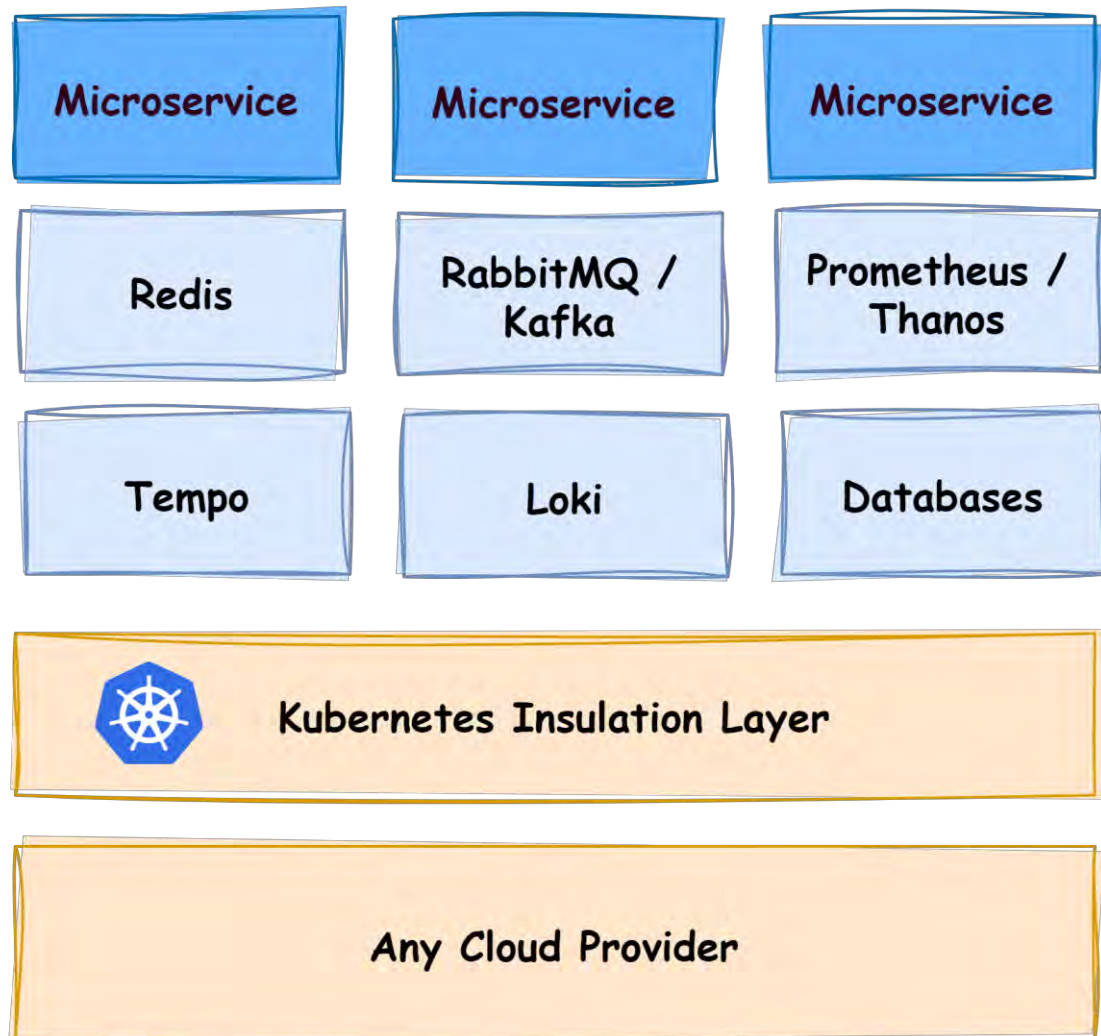
How to design cloud agnostic architecture?



Use Kubernetes and you are done!

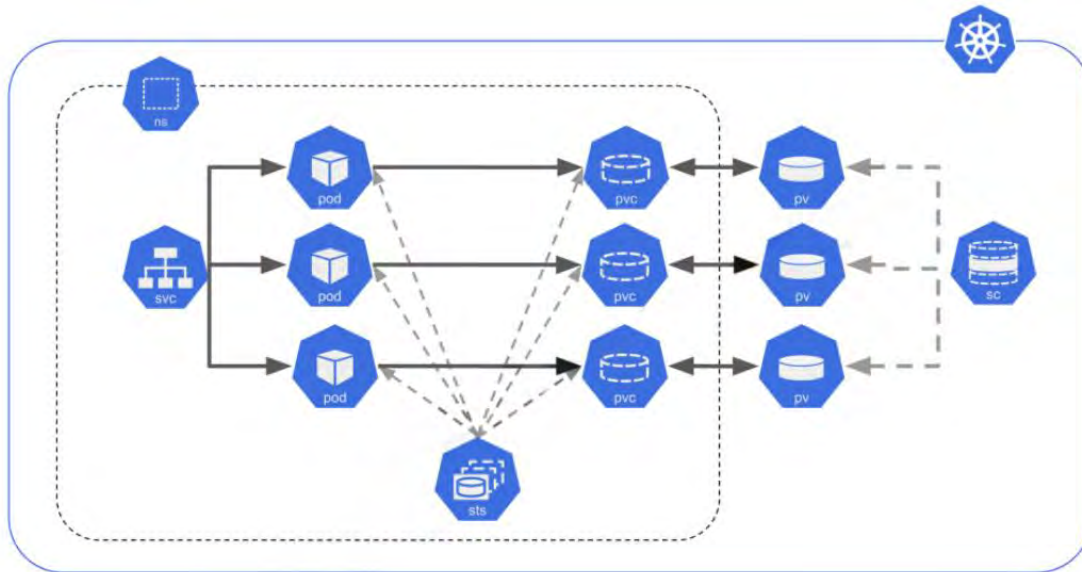
- Build services as containerized workloads, our friend Docker.
- Deploy containers to the Kubernetes product offered by the cloud vendor. (AKS, EKS, GKE)
- New capability -> New Container
- Switch providers as long as Kubernetes is available

New capability -> New Container



- Need Messaging? RabbitMQ or Kafka in containers
- Need Cache? Redis High Availability cluster in containers
- Need a RDBMS? PostgreSQL in containers
- Need an Object Storage like AWS S3? Minio in containers
- Need Monitoring? ELK (Elasticsearch, Logstash, Kibana) stack in containers

Run a stateful application on Kubernetes

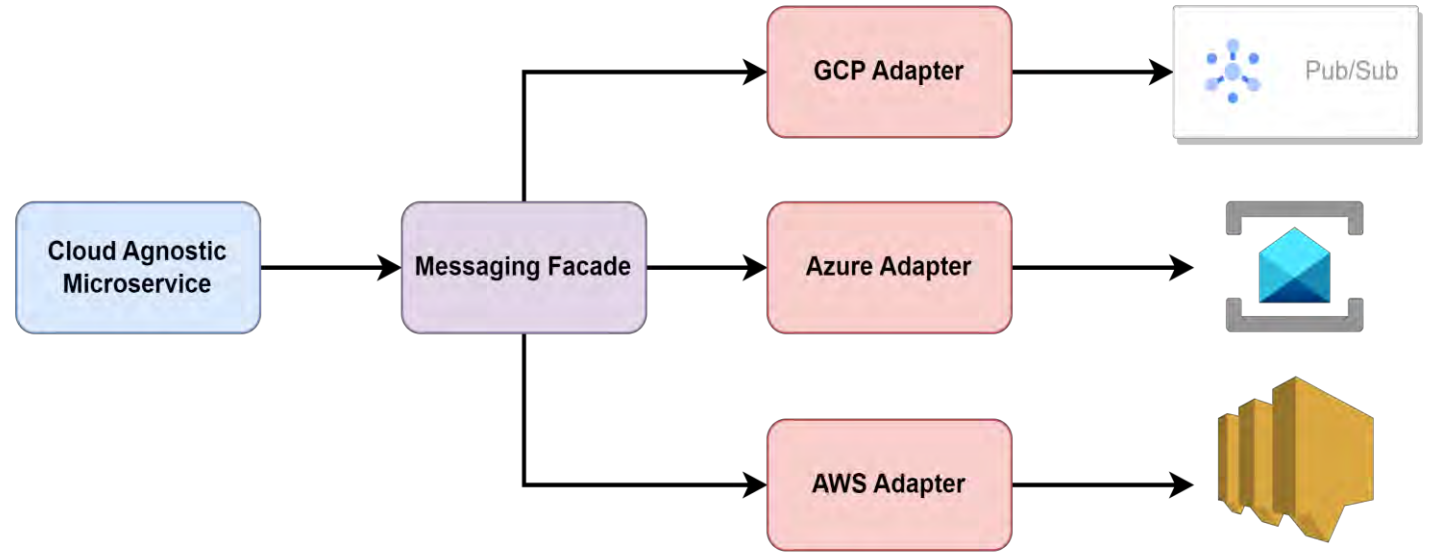


- StatefulSets
 - Startup, Scale-up/down, Rolling Upgrades, Termination - Ordered operations with ordinal index
 - Unique network ID/name to maintain affinity
 - Persistent storage disk linked to the ordinal index – network, local, cloud
 - Headless service (don't forget network policies for security)

Other challenges

- Cloud capabilities
 - Different cloud providers -> different data centers, data residency, GDPR
 - Failover, resilience, and latency -> depend on the location of data centers
- Networking
 - Unlike AWS and Azure, GCP's Virtual Private Cloud resources are not tied to any specific region. It's a global resource.
- Data egress cost -> Ingress is free, but egress can be expensive
- Infrastructure as a code

Loosely coupled architecture

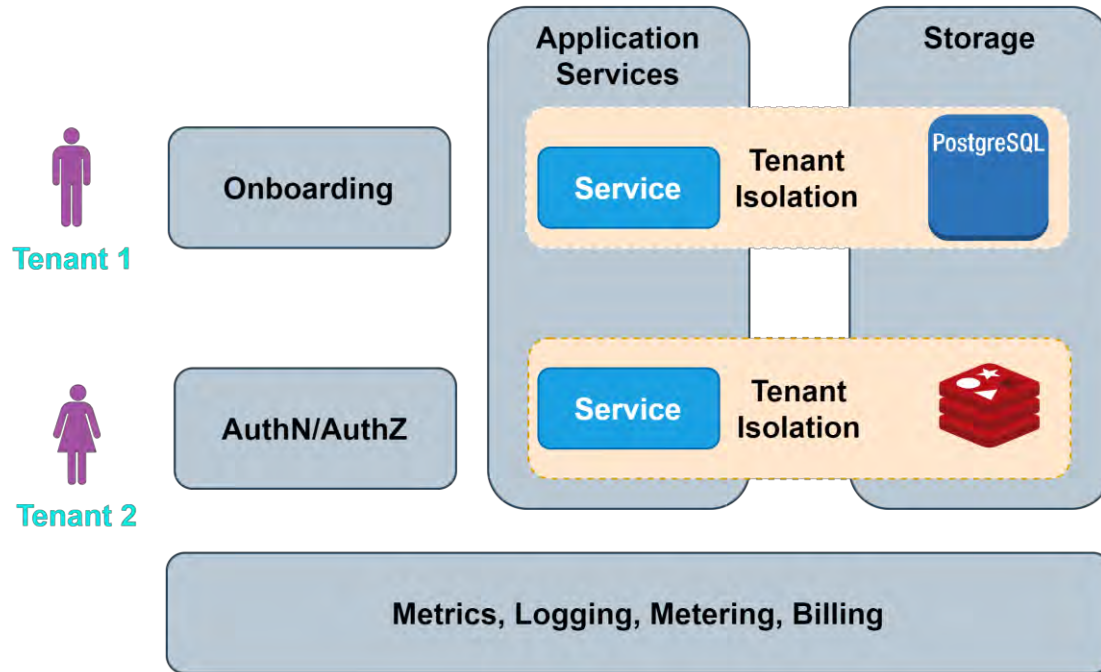


A facade is a structural design pattern that provides a simplified interface to a library, a framework, or any other complex set of classes.

Strategic lock-in

- Identify the areas where lock-in must be kept to a minimum.
- Only use products, that have corresponding counterparts on the other platforms.
 - Databases
 - SQL DB -> GCP Cloud SQL, Azure database for PostgreSQL
 - Runtime -> GCP GKE, AWS EKS
 - Serverless -> Knative
 - Timeseries database -> GCP BigTable, AWS DynamoDB

Fundamentals of SaaS



Tenant aware operations dashboard

- Onboarding
- AuthN/AuthZ
- Data partitioning
- Tenant isolation
- Metering and billing
- Tenant aware operation models

Multi-tenant impact



Frontend

Authentication
Routing
Feature Flags



API Gateway

Authorization
Throttling
Caching



Business logic

Metrics
Logging
Metering



Data persistence

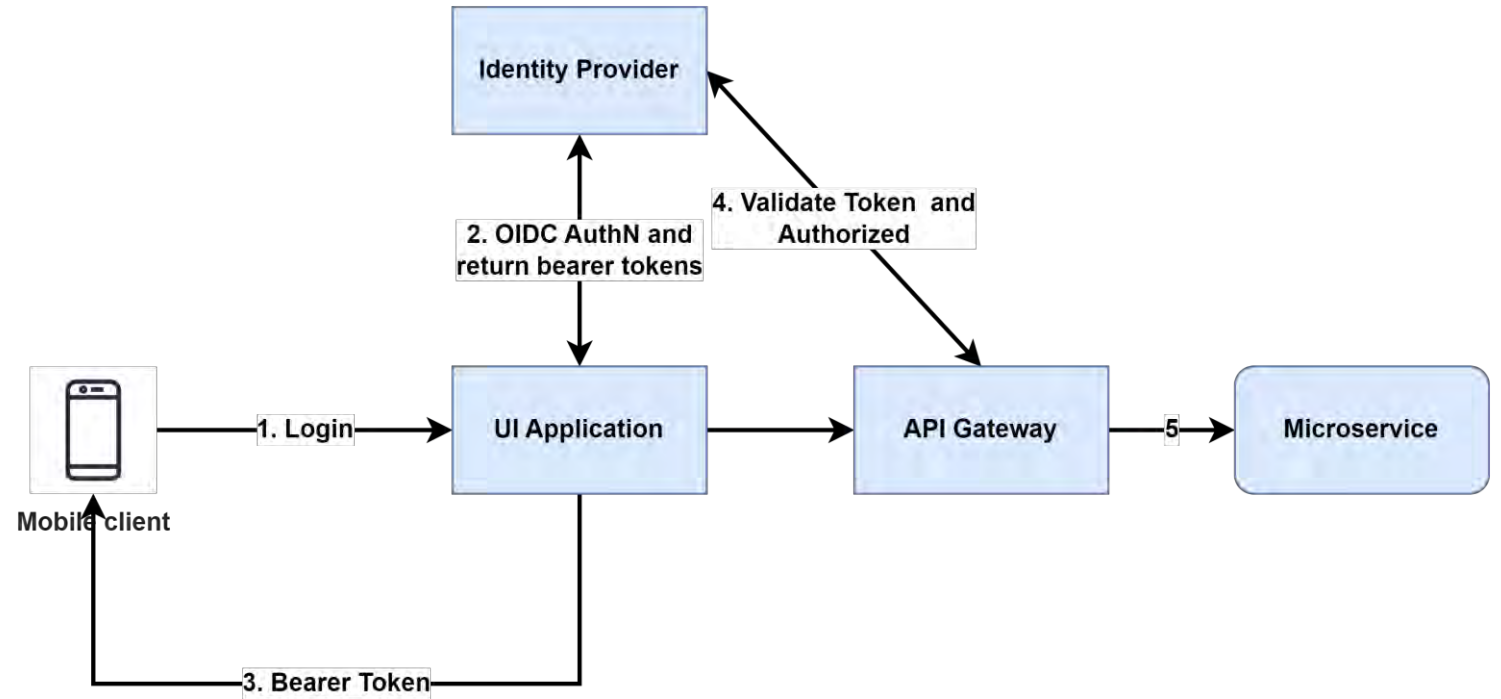
Data Access, Partitioning
and Isolation
Backup/restore



Infrastructure

Provisioning
Isolation
Maintenance
Tenant lifecycle

A normal microservice flow



Non-SaaS microservice

```
# Create a DynamoDB client
dynamodb = boto3.resource('dynamodb', aws_access_key_id=aws_access_key_
                           aws_secret_access_key=aws_secret_access_key,
                           region_name=aws_region)

# Specify the table name
table_name = 'your_table_name'

# Define API endpoint to get data
@app.route('/get-data', methods=['GET'])
def get_data():
    try:
        # Get key from request query parameter
        key_to_retrieve = request.args.get('key')

        # Get a reference to the table
        table = dynamodb.Table(table_name)

        # Use the get_item method to retrieve data
        response = table.get_item(
            Key={
                'your_primary_key_name': key_to_retrieve
            }
        )

        # Check if item exists
        item = response.get('Item')
        if item:
            return jsonify({"message": "Item retrieved successfully", "
        else:
            return jsonify({"message": f"No item found with key: {key_t

    except Exception as e:
        return jsonify({"message": f"Error: {str(e)}"}), 500
```

SaaS microservice

```
# Create a dynamodb client
dynamodb = boto3.resource('dynamodb', aws_access_key_id=aws_access_key,
                           aws_secret_access_key=aws_secret_access_key,
                           region_name=aws_region)

# Define a function to get the table based on the tenant ID
def get_table(tenant_id):
    table_name = f'{tenant_id}_your_table_name'
    return dynamodb.Table(table_name)

# Define API endpoint to get data
@app.route('/get-data', methods=['GET'])
def get_data():
    try:
        # Get tenant ID from request headers or other methods
        tenant_id = request.headers.get('X-Tenant-ID')

        if not tenant_id:
            return jsonify({"message": "Tenant ID not provided in header"})

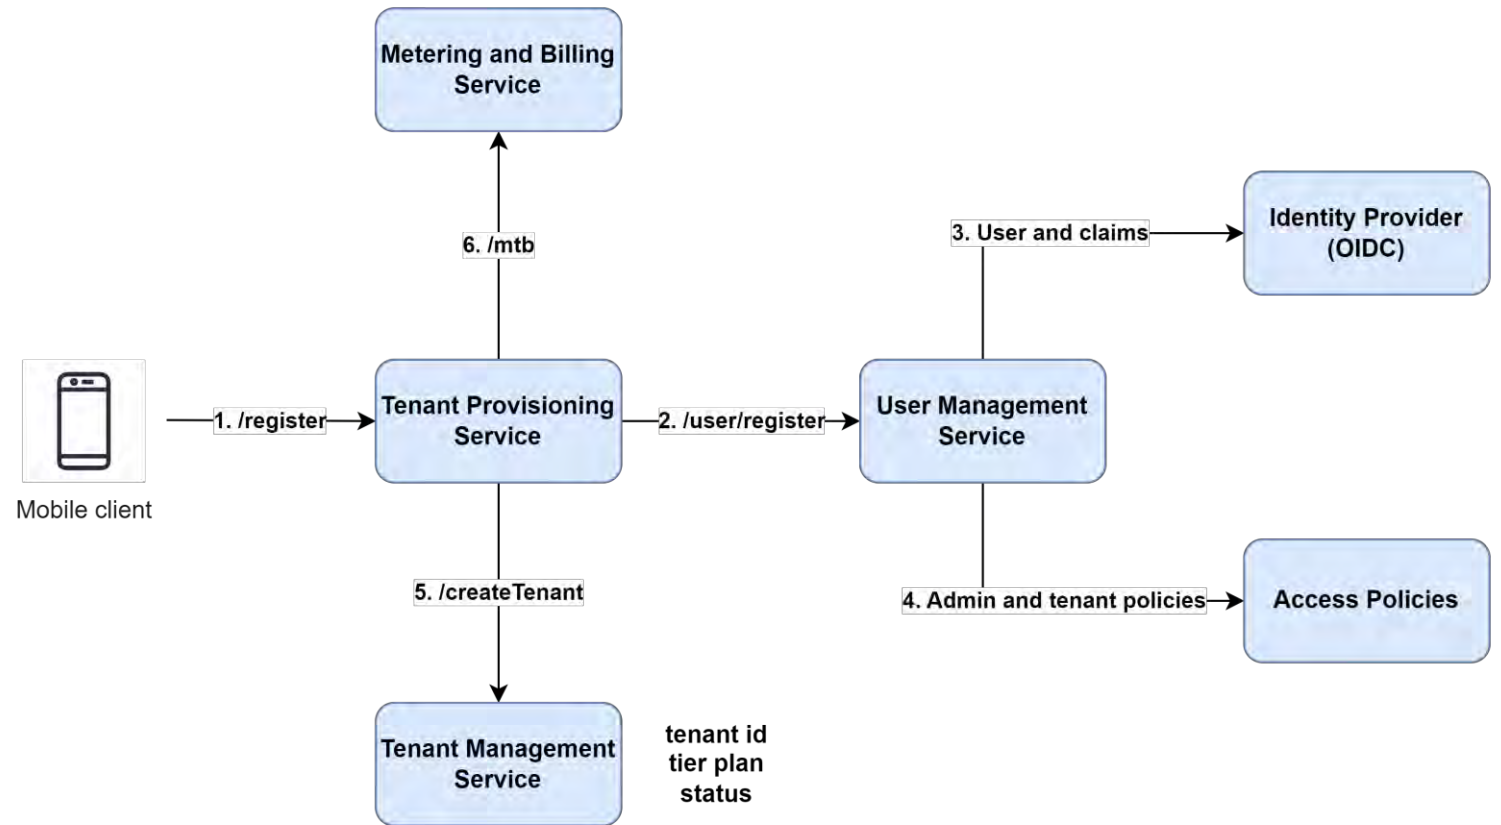
        # Get key from request query parameter
        key_to_retrieve = request.args.get('key')

        # Get a reference to the table for the specific tenant
        table = get_table(tenant_id)

        # Use the get_item method to retrieve data
        response = table.get_item(
            Key={
                'your_primary_key_name': key_to_retrieve
            }
        )

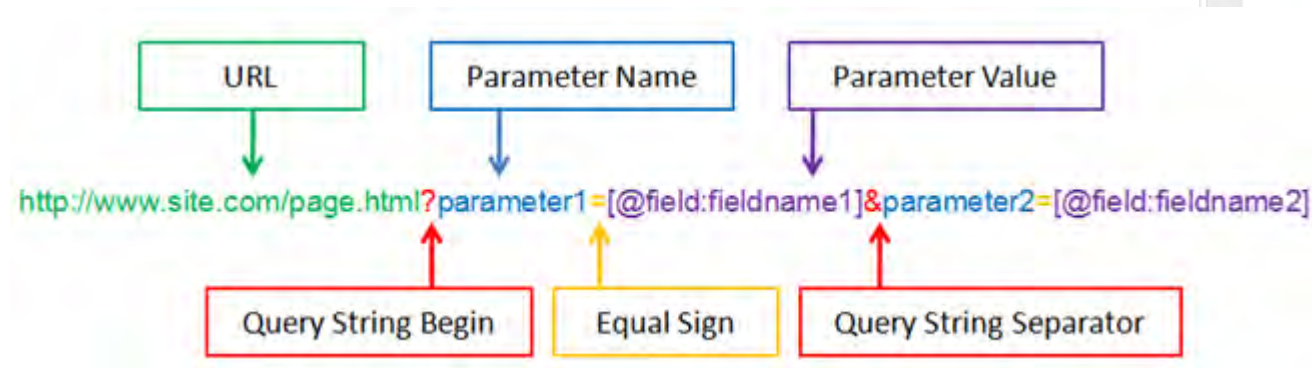
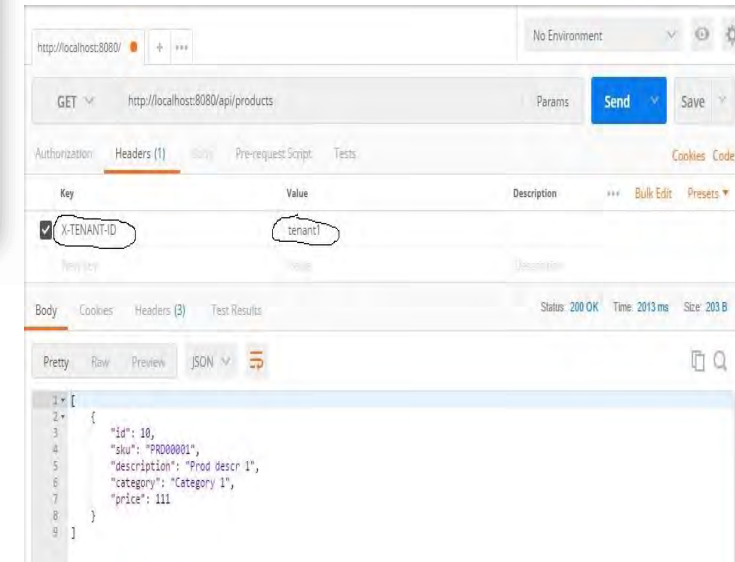
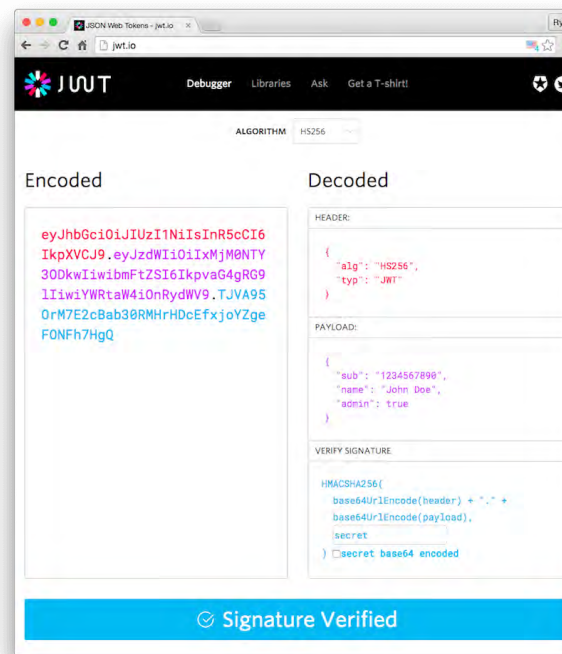
        # Check if item exists
```


Provision a tenant



Acquiring Tenant Context

- JWT token (JSON web token)
- In the URL as a query string parameter
- Request Header
- Separate microservice



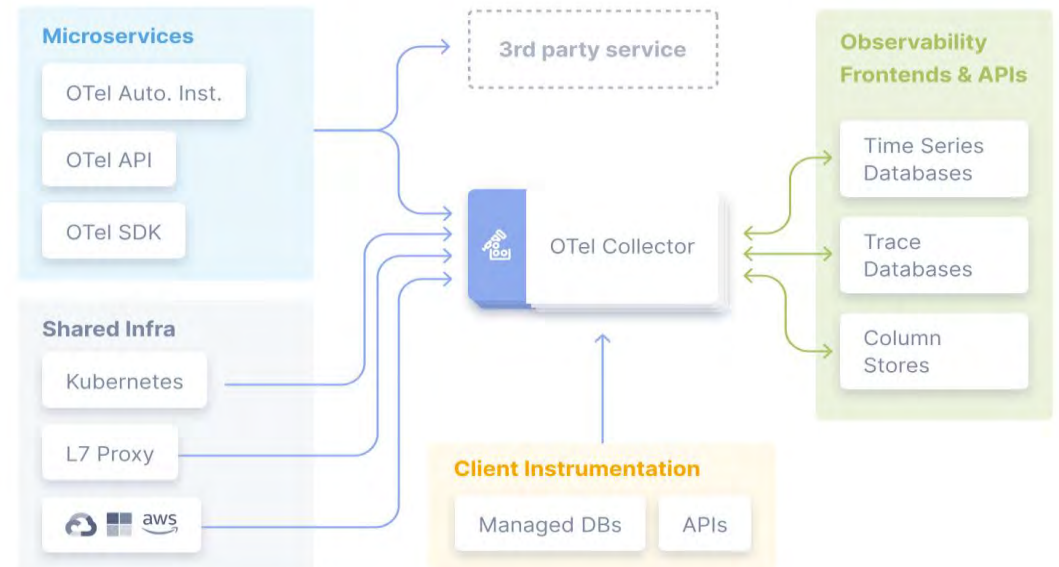
Common libraries

- Get tenant ID from JWT token
- Structured logging
- Metrics

```
auth_header = request.headers.get('Authorization')
bearer_token = re.split(r"^[B|b]earer +", auth_header)[1]
header, payload, signature = bearer_token.split('.')
claims = base64.b64decode(payload).decode('UTF-8')
tenant_id = claims['custom:tenant_id']
```

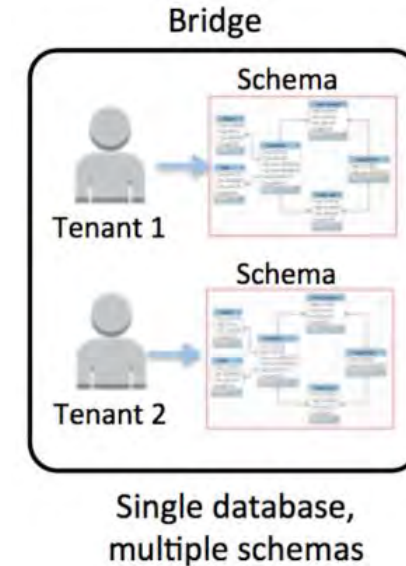
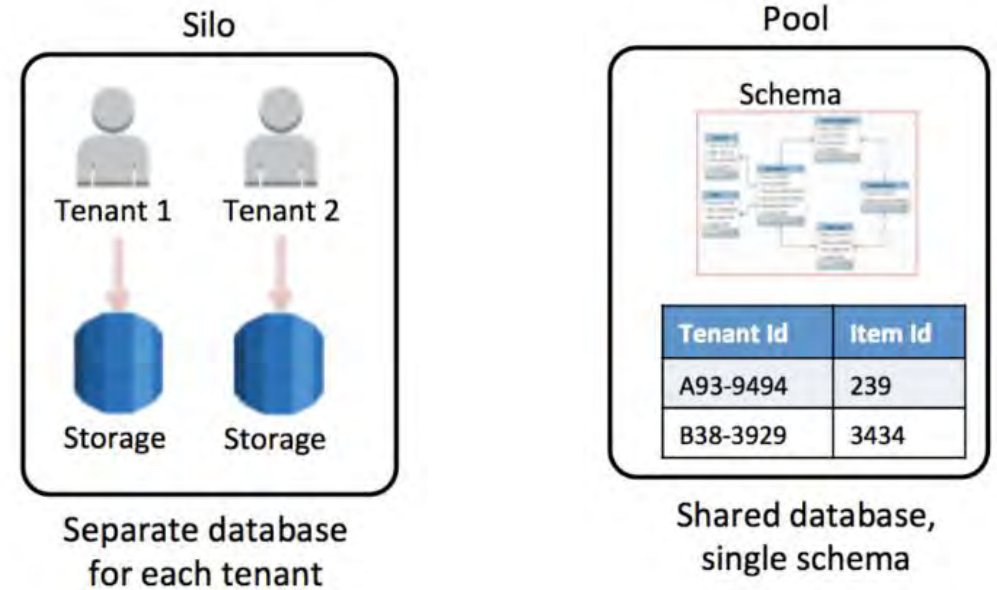
```
{
  "message": "Computing 10001 for client_id 1 shipment_id 2",
  "payload": {
    "client_id": 1,
    "shipment_id": 2,
    "item_id": 10001
  },
  "metadata": {
    "code": {
      "file_url": "/a/url/code.py",
      "line_number": 186,
      "file_name": "code.py",

```



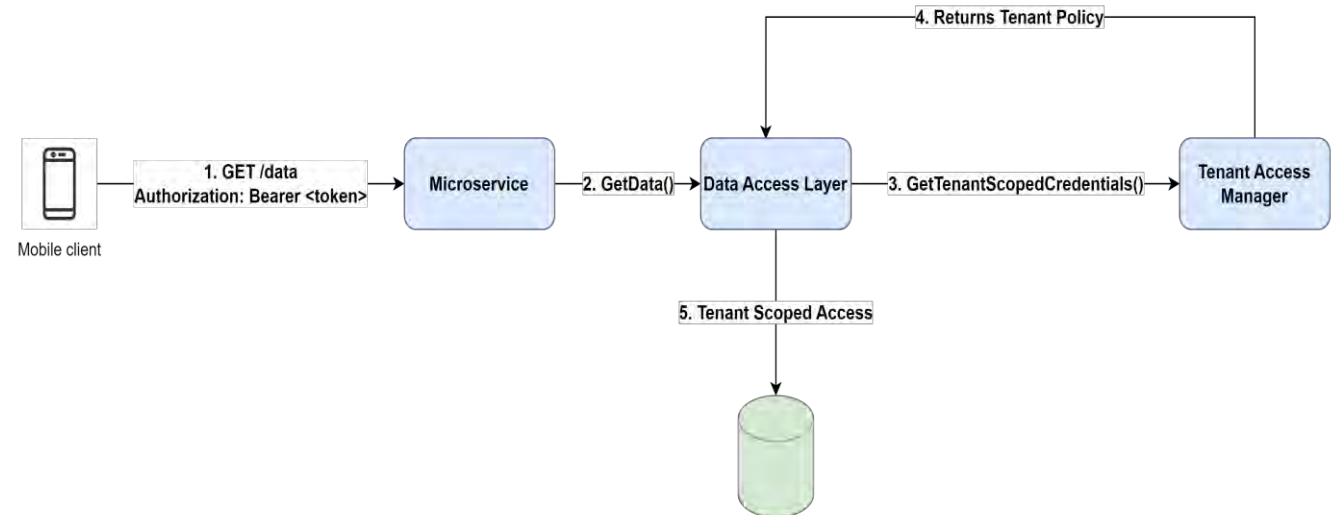
Data partitioning

- A microservice-based decision
 - Compliance & security
 - Performance
 - Data distribution
 - Noisy neighbor



Tenant Isolation

- Silo Model – every tenant gets their own environment
- Pool Model – isolation through runtime policies
 - Before touching any resource, get the policies assigned to a specific tenant



Take aways

- Strategic lock-in
- Loosely coupled architecture
- Tenant lifecycle
- Tenant context
- Data partition
- Tenant isolation



Thank You!
