



ADC-D3

Multi-tenant data isolation using AWS databases

Venkatesh Periyathambi

Principal Solutions Architect, Database
AWS

Agenda

- SaaS Fundamentals
- Multi-tenant deployment models
- Design Challenges
- Decision Matrix
- Design Considerations
- Takeaways
- Call to action


AWS Databases Landscape



Amazon Aurora




Amazon DocumentDB
(with MongoDB compatibility)




Amazon DynamoDB



Amazon MemoryDB
for Redis




Amazon Keyspaces
(for Apache Cassandra)




AWS Database Migration
Service (AWS DMS)



Amazon Neptune



Amazon Quantum Ledger
Database (Amazon QLDB)




Amazon Relational Database
Service (Amazon RDS)



Amazon RDS on
VMware



Amazon Timestream



Amazon ElastiCache

SaaS Fundamentals



What is SaaS

Software-as-a-Service (SaaS) is a business and software **delivery model that enables organizations to offer their solution in a **low-friction**, service-centric approach.**

Core SaaS architecture concepts

Data partitioning

How is data organized and stored in a multi-tenant environment?

Tenant isolation

How does your architecture ensure that one tenant can't access the resources of another tenant?

SaaS Identity

How is a user identity associated with a tenant identity and how is that context shared across your architecture?

Onboarding

How are new tenants added to your system in a frictionless manner?

Tiering

How do we offer different experiences to tenants at different price points?

Metering and billing

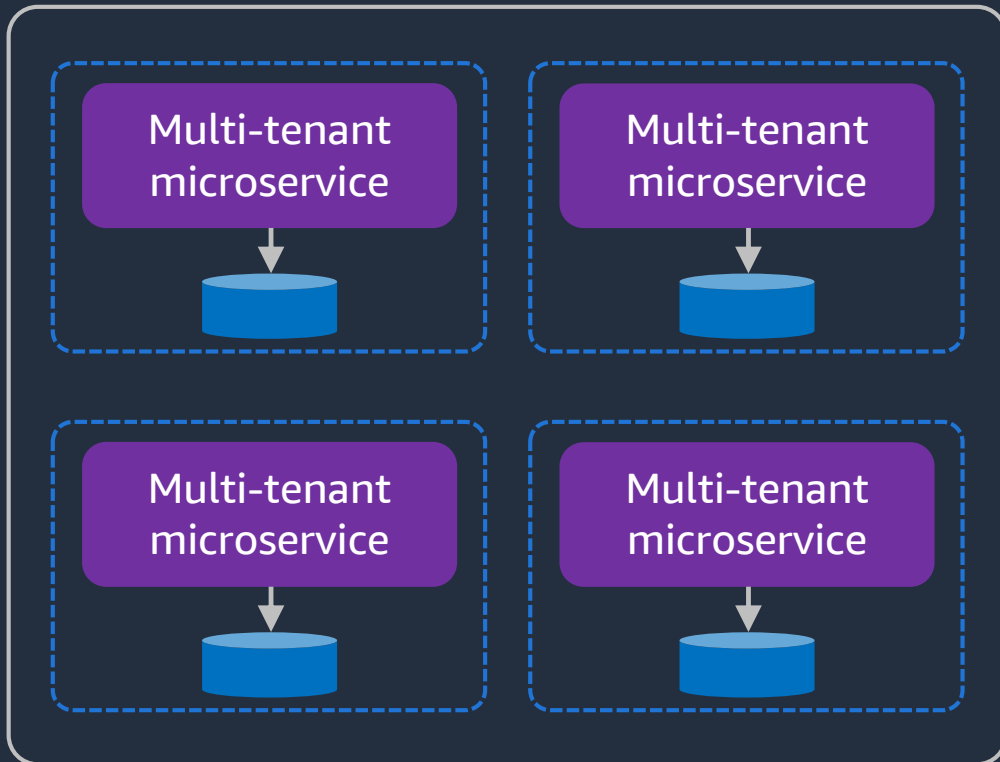
How do you instrument your application to meter tenant activity and generate a bill?

Two halves of SaaS

Application plane

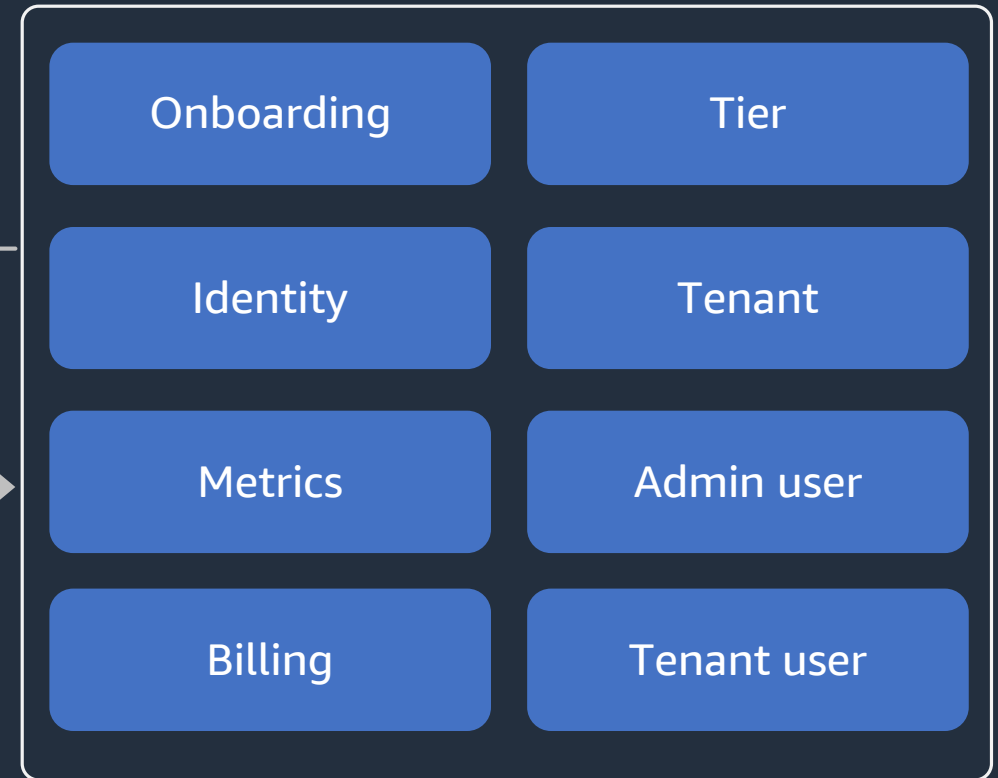


Web tier



Control plane

Admin console

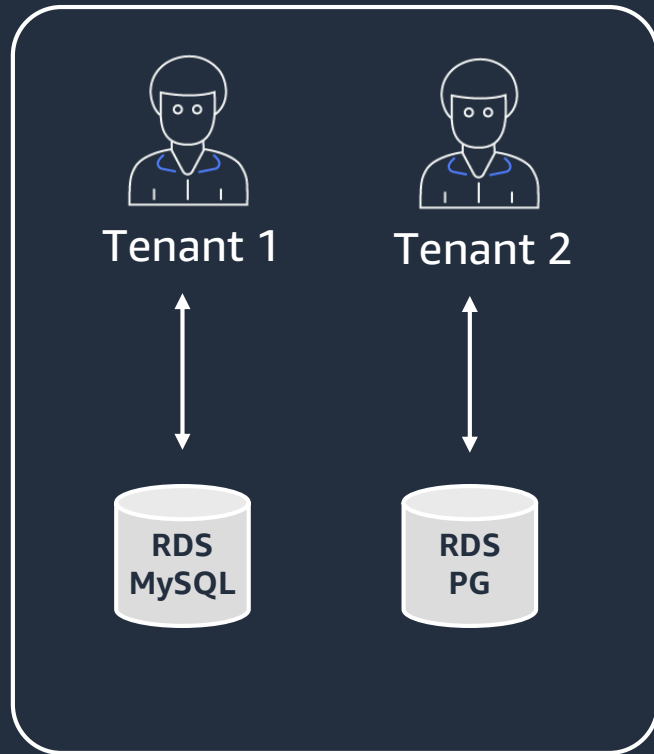


Multi-Tenant Deployment Models

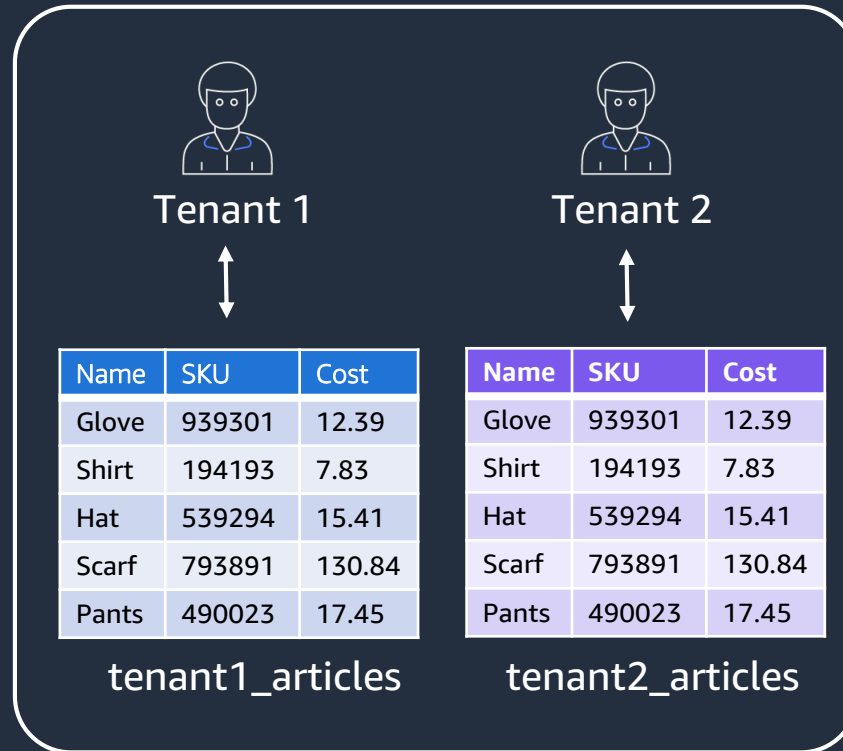


Multi-tenant Deployment Models

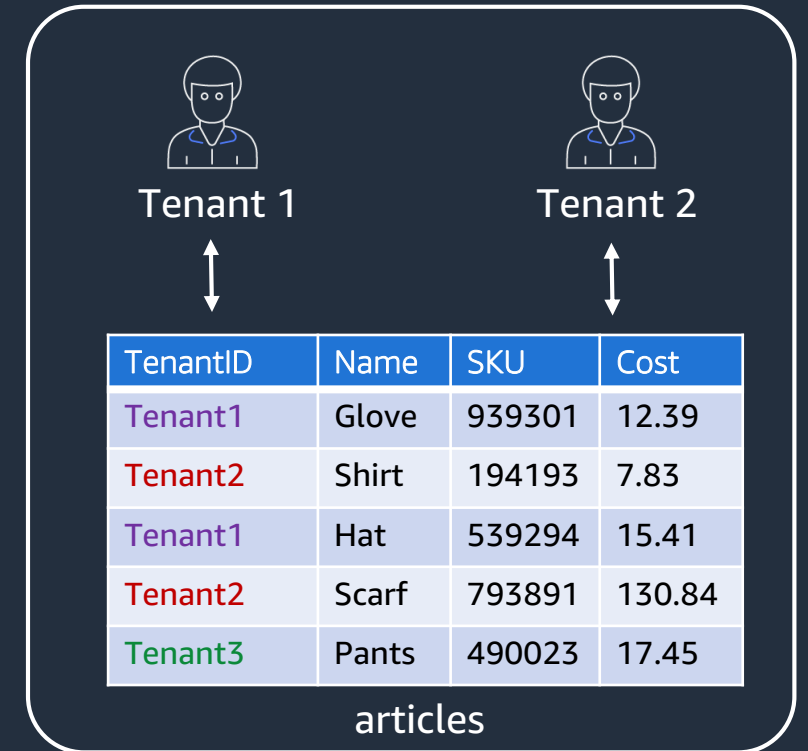
Silo model



Bridge model



Pool model



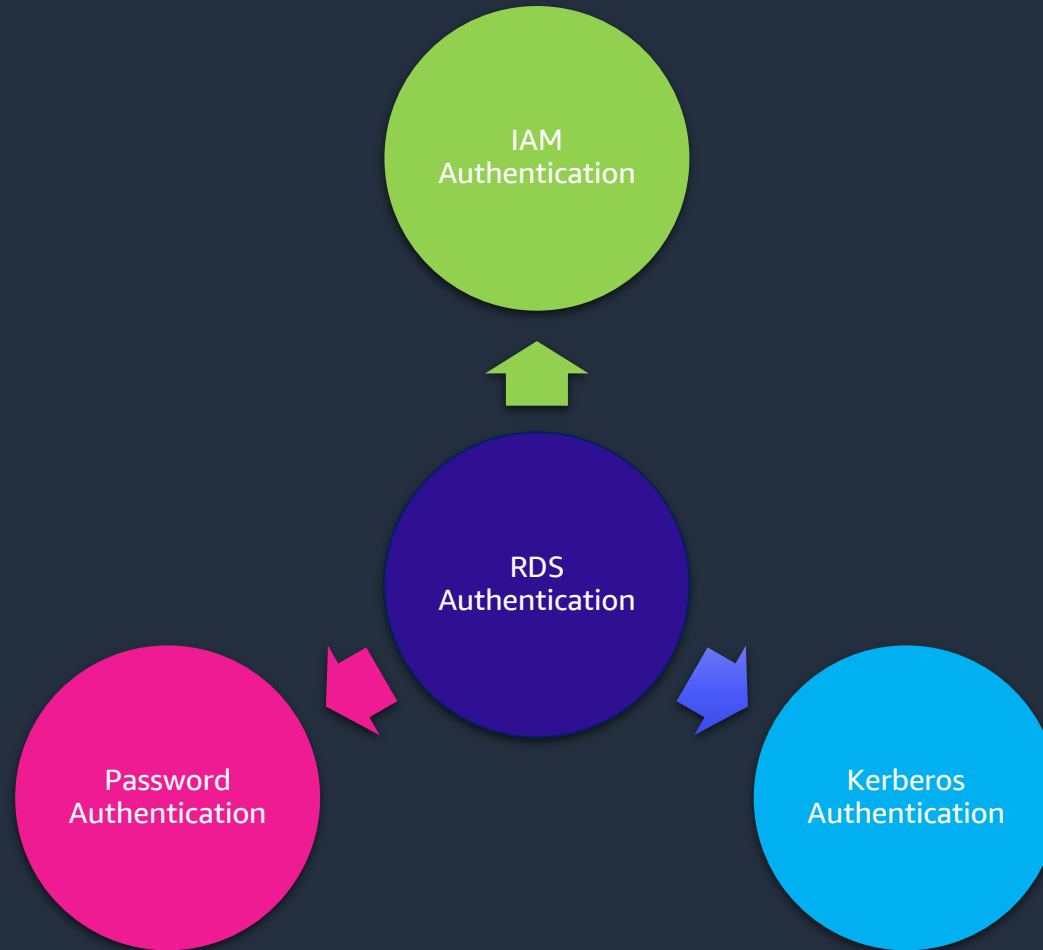
Max. isolation

Max. Resource Sharing

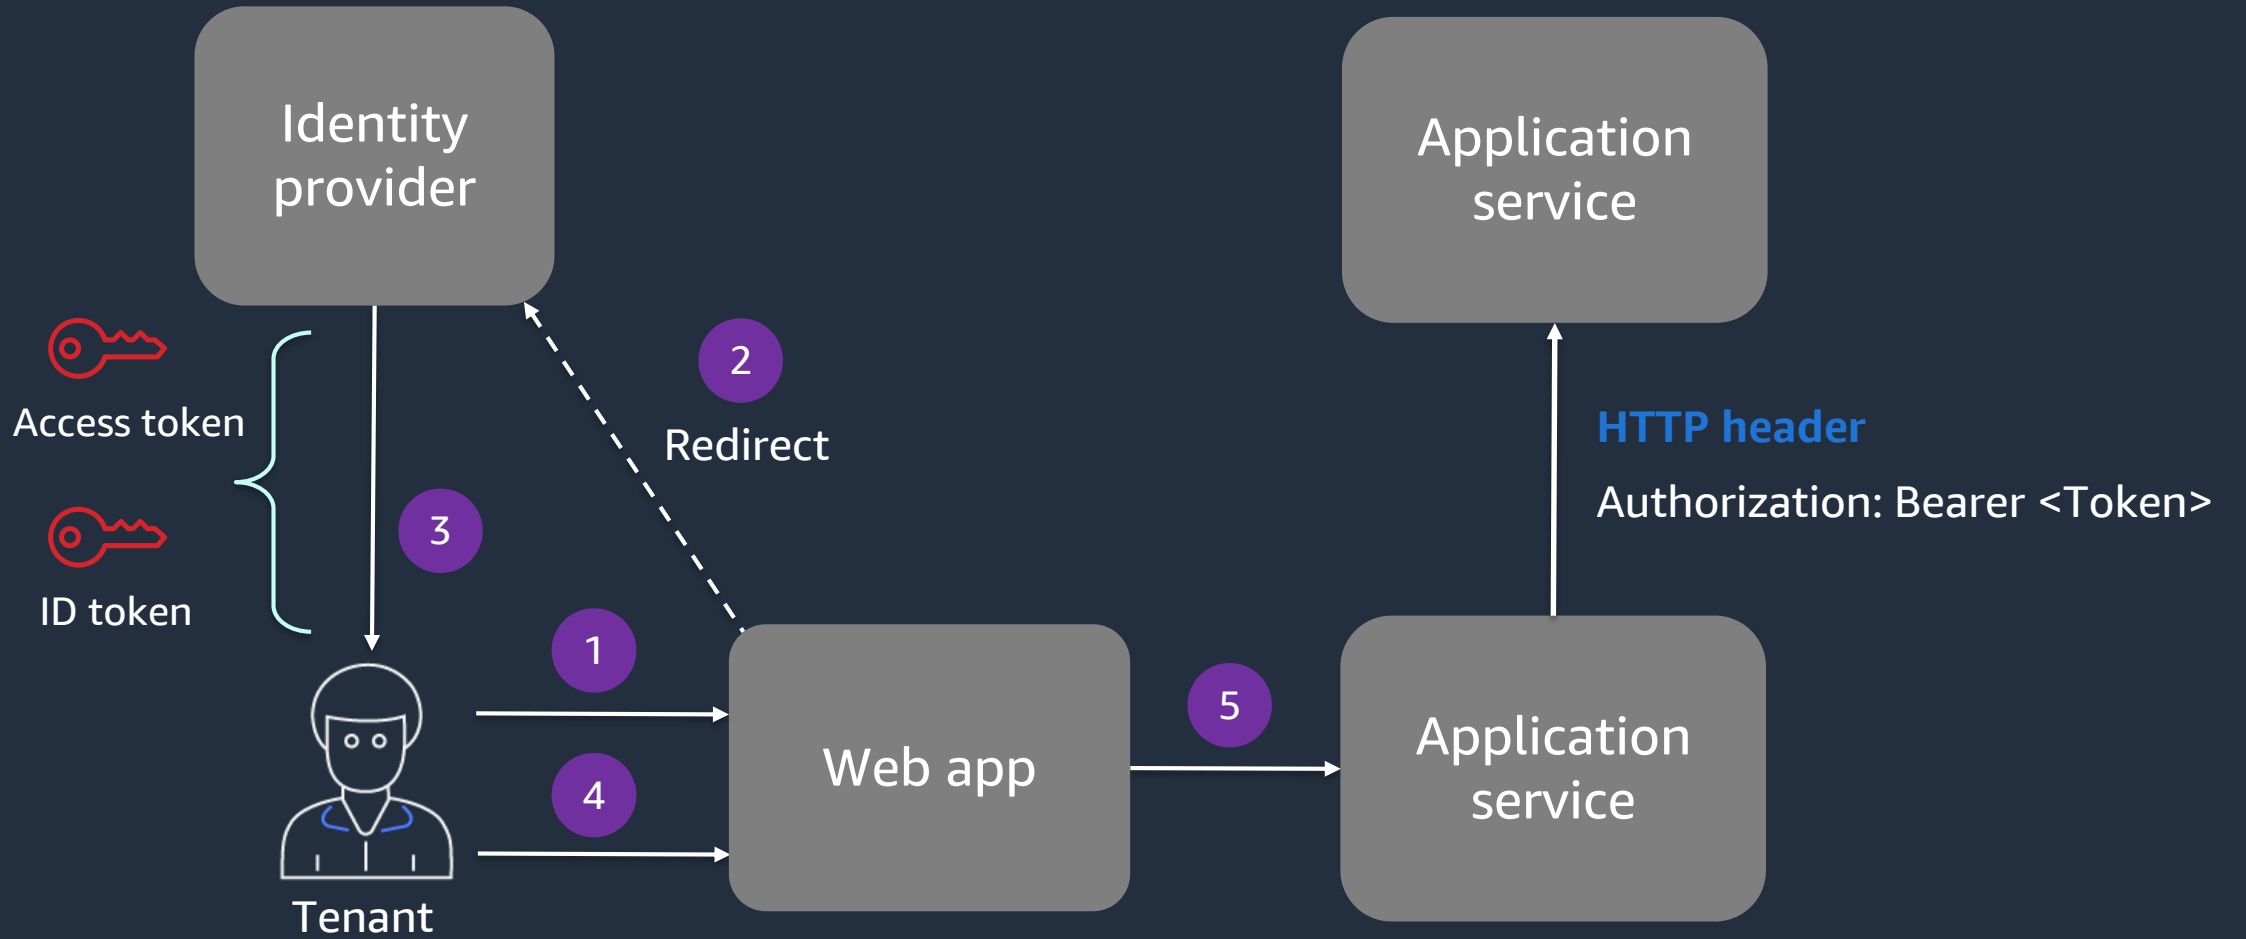
Multi-tenancy Authentication



Amazon RDS Authentication Methods



Tenant-aware identity

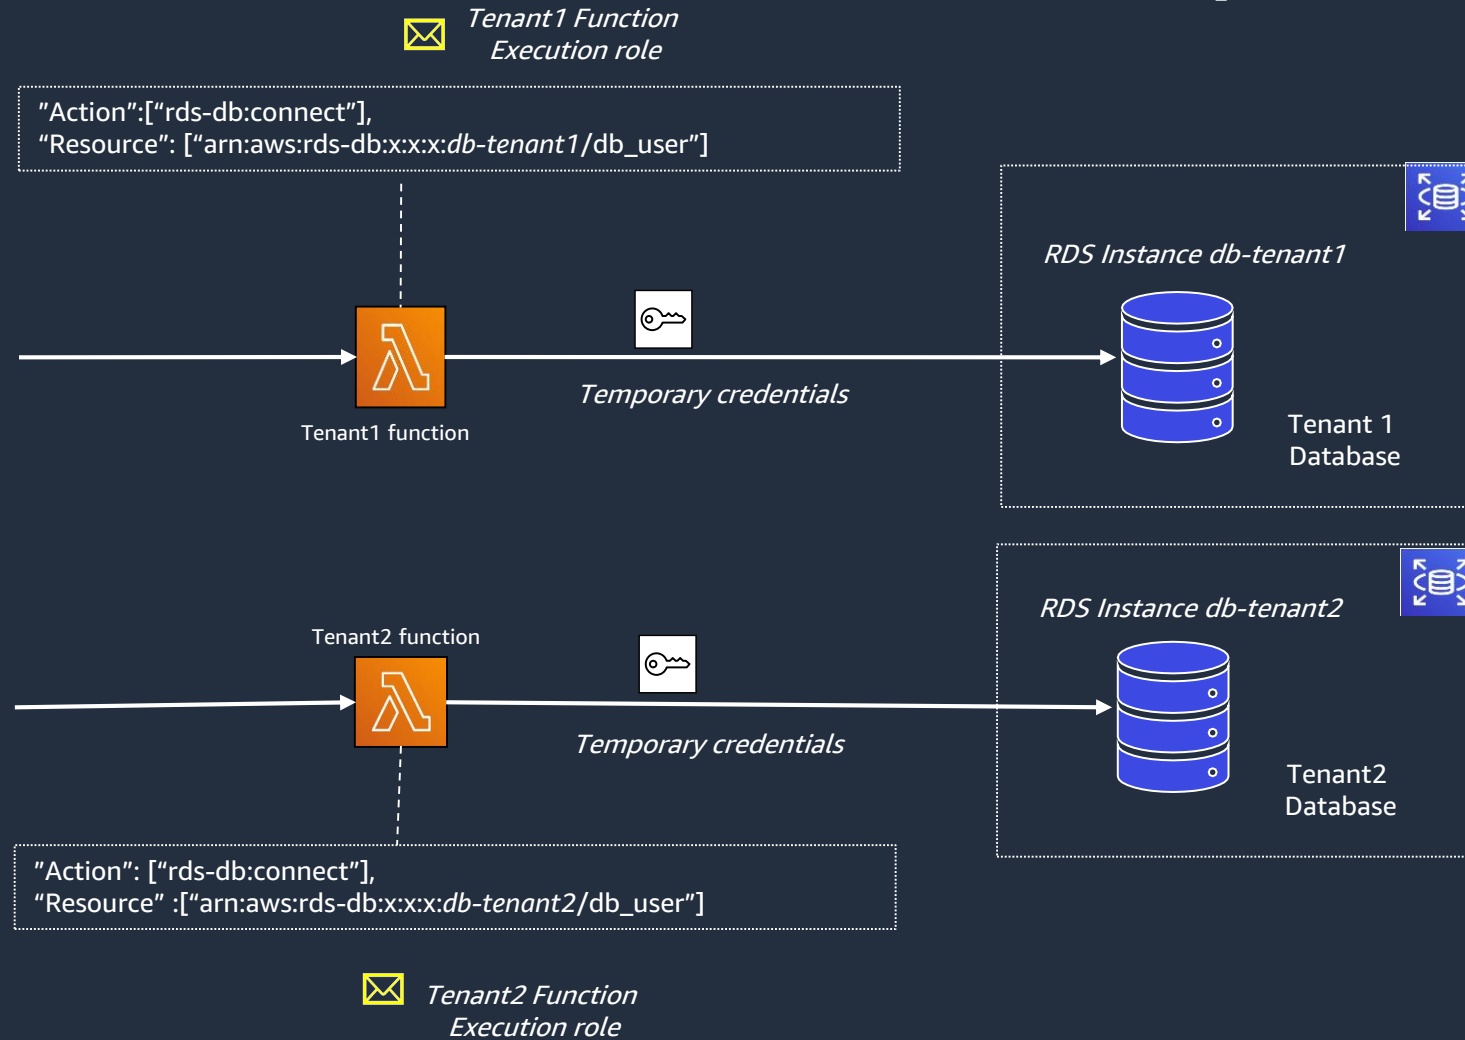


Design Challenges #1

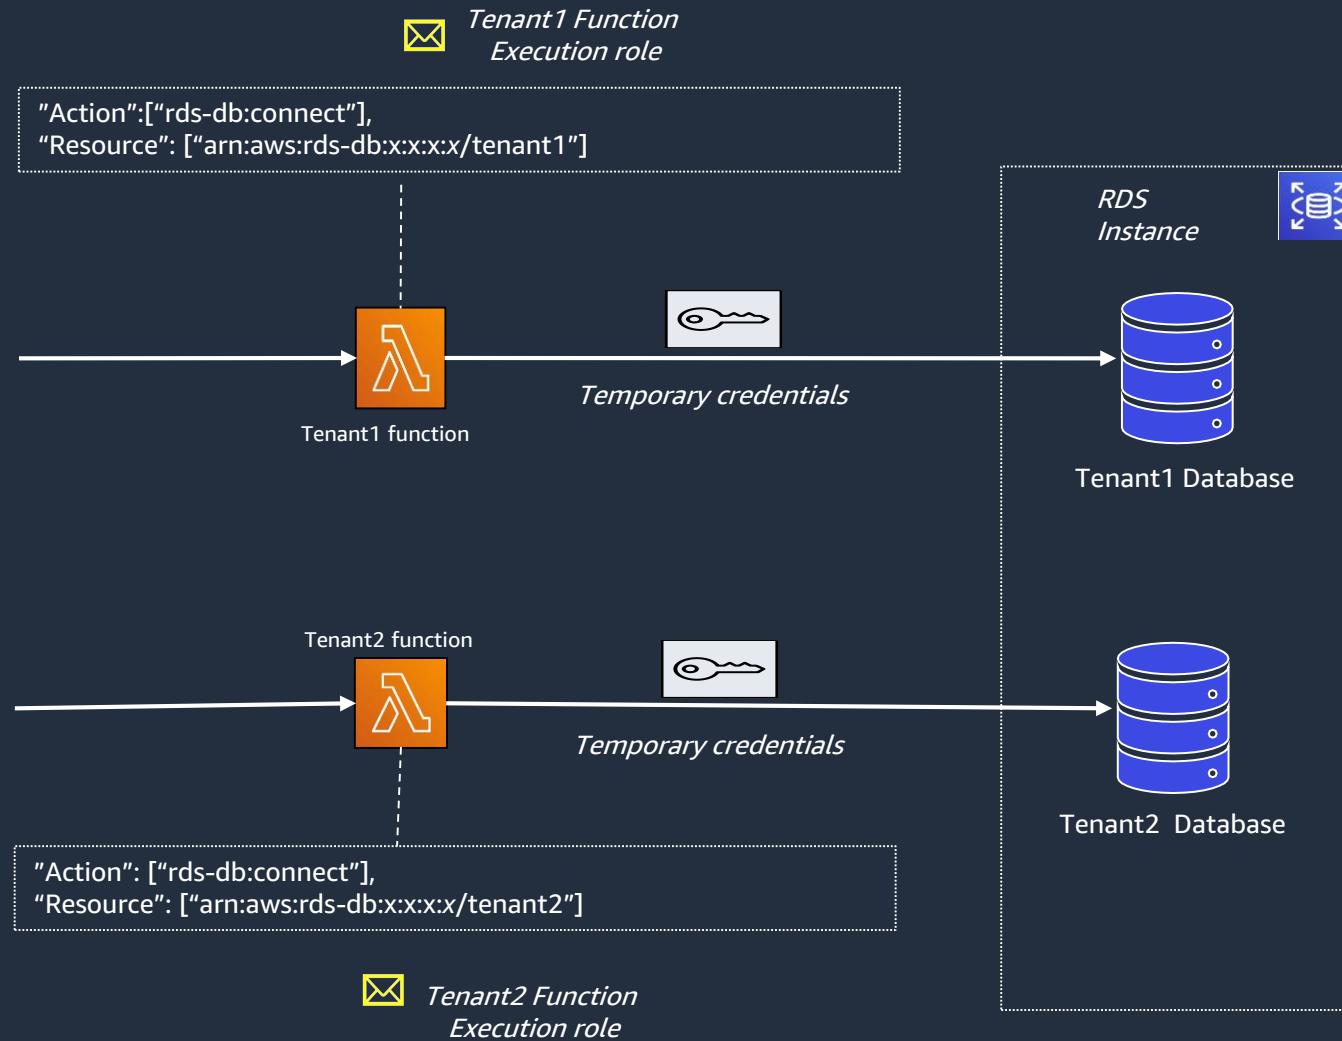
Data Access Patterns



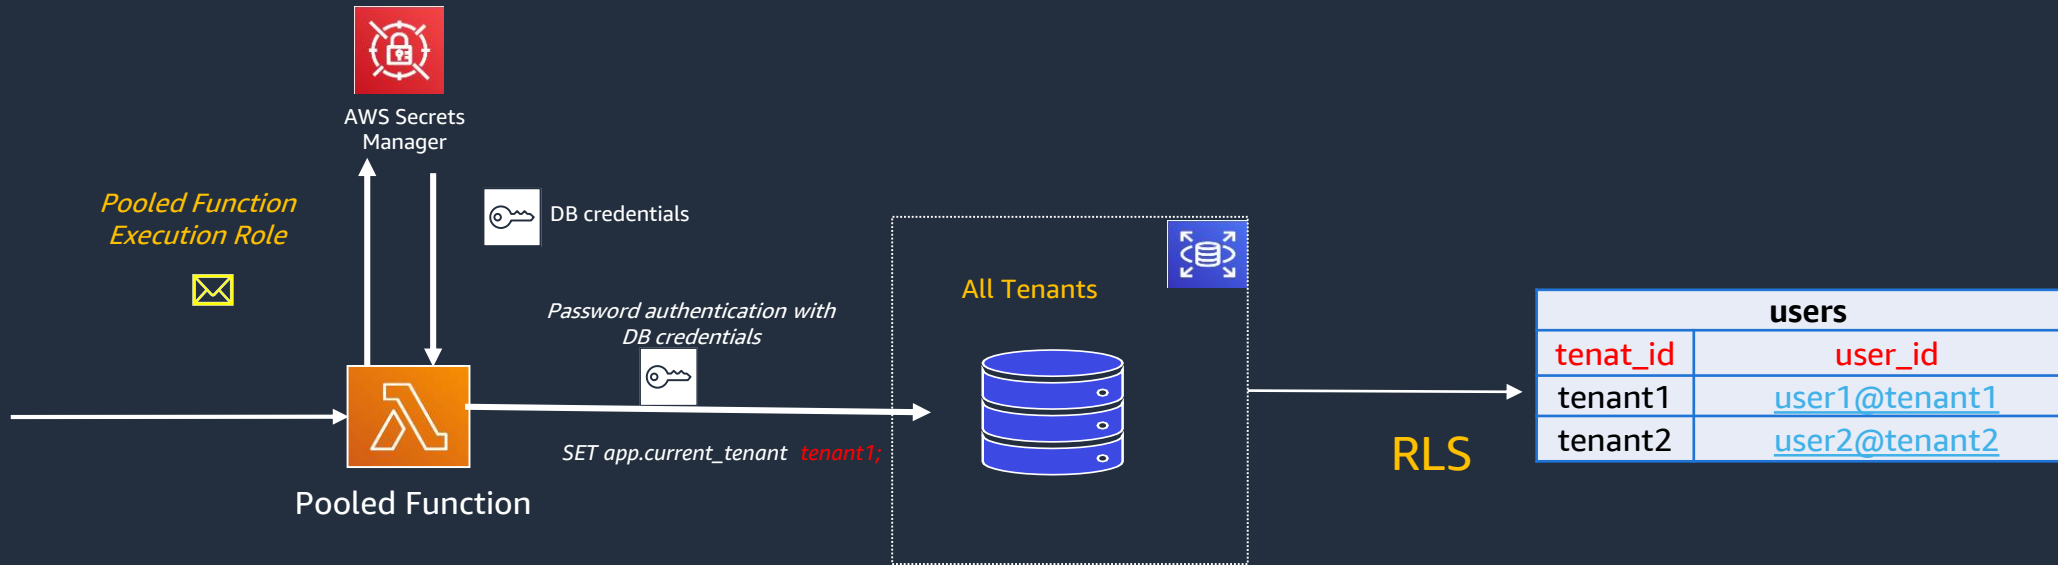
Silo database isolation with siloed compute



Bridge database isolation with siloed compute



Pool database Isolation with pooled compute



Pool with RLS and Amazon Aurora PostgreSQL

Initialize RLS

```
-- Turn on RLS
ALTER TABLE tenant ENABLE ROW LEVEL SECURITY;

-- Scope read/write by tenant
CREATE POLICY tenant_isolation_policy ON tenant
USING (tenant_id::TEXT = current_user);
```

Query with RLS

```
-- No tenant context required
rls_multi_tenant=> SELECT * FROM tenant;

-- Attempt to force tenant id
rls_multi_tenant=> SELECT * FROM tenant WHERE tenant_id
= 'tenant1'
```

FK	SKU	Name
Tenant1	93529-94	Black T-shirt
Tenant2	24411-01	Blue hoodie
Tenant1	76235-92	Wool socks
Tenant3	95419-37	Green polo
Tenant2	88314-99	White hat
Tenant1	24598-72	Tennis shoes

Design Challenges #2

Sizing considerations

The sizing challenge




Tenant 1



Tenant 2

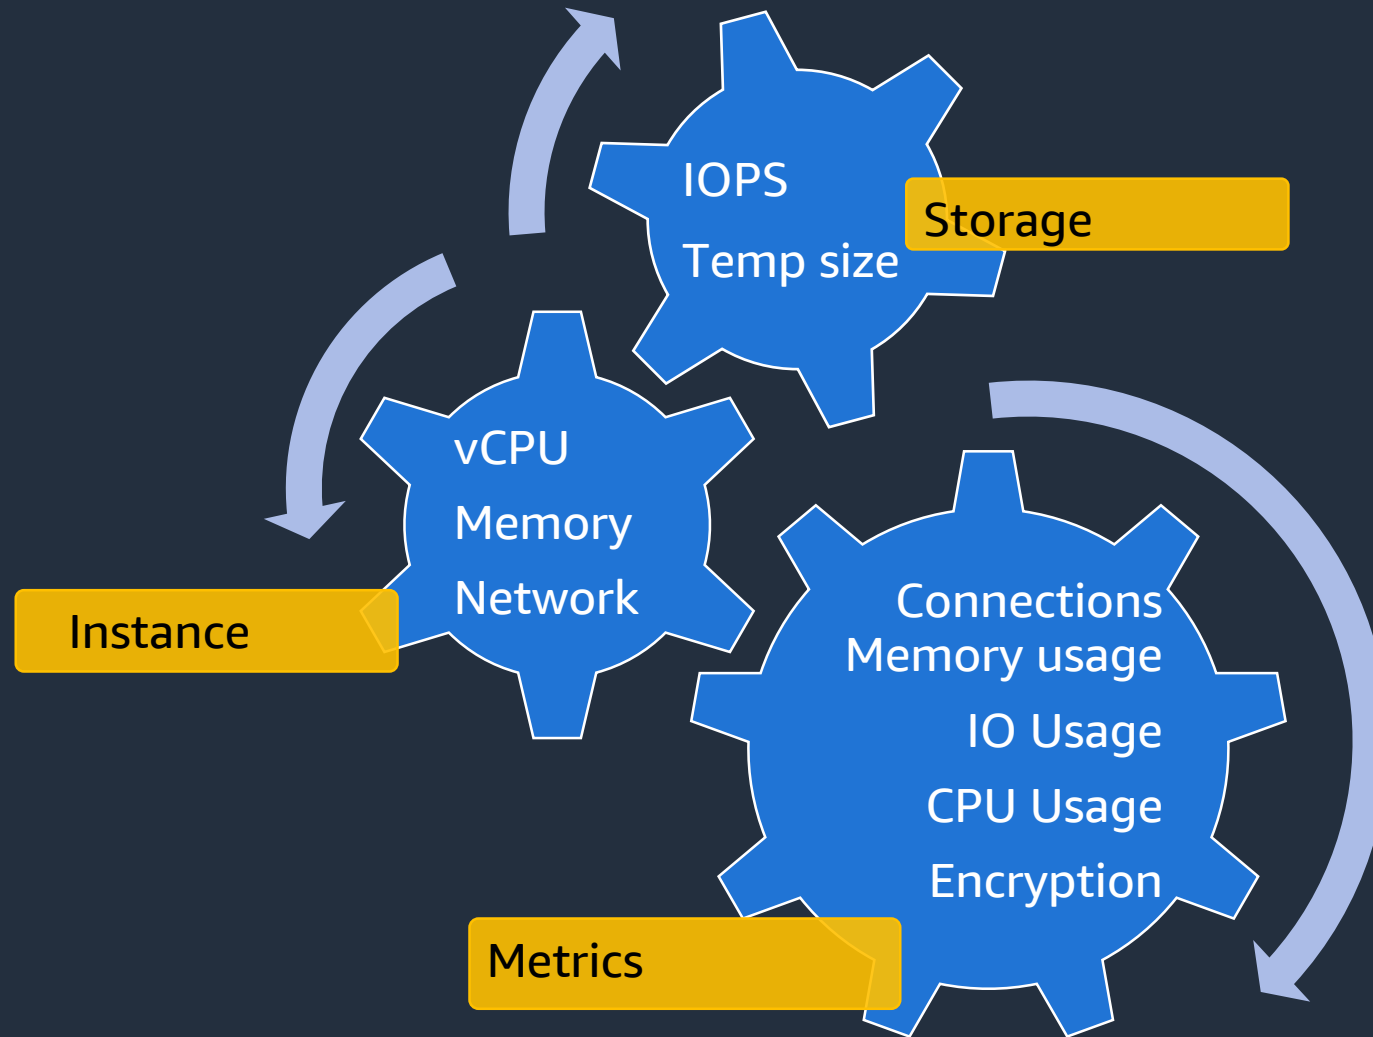


Tenant 3



- How do you accommodate different size tenants (noisy neighbor)?
- How do you prevent over-provisioning?
- How do you optimize based on actual consumption?
- How will you support tiering and SLAs?

Tenant bin packing



Design Challenges #3

Database engine



IAM granularity can be a factor

Course-grained IAM control




Amazon RDS



Amazon Elasticsearch Service

Effect: "Allow",
Action: rds-db:connect
Resources:
- arn:aws:rds-db:db-instance-id/db-user

Fine-grained IAM control



Amazon DynamoDB

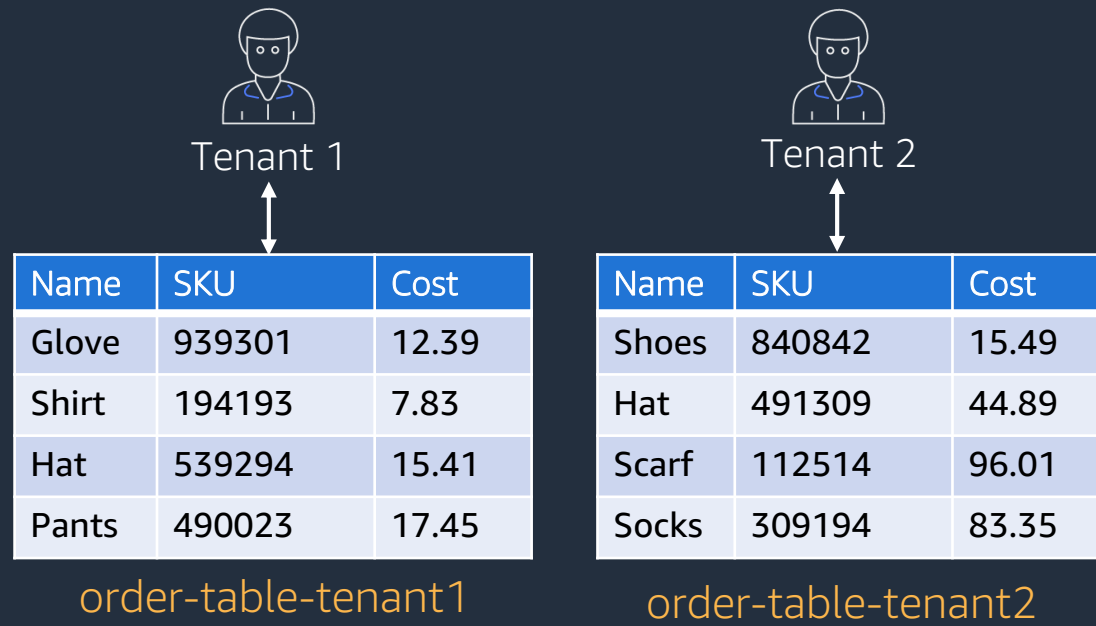


Amazon Simple Storage Service (S3)

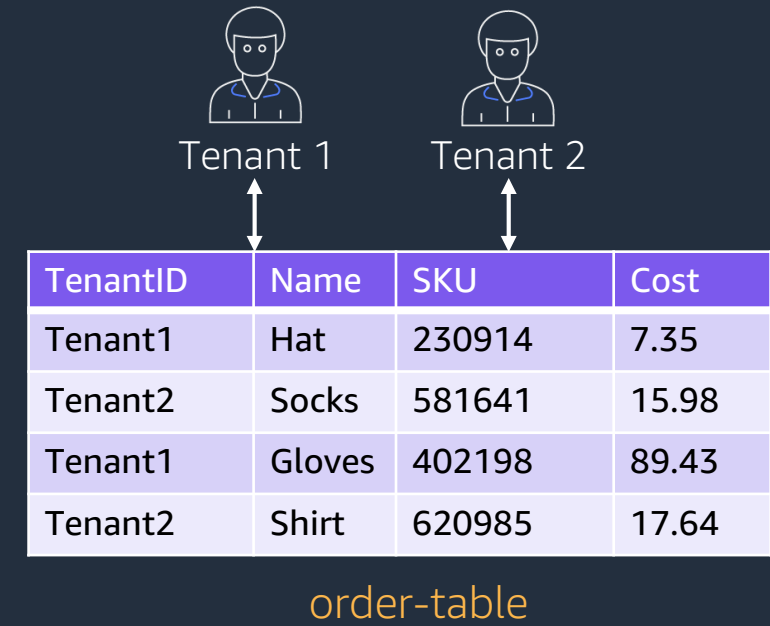
Effect: "Allow",
Action: dynamodb:GetItem
Resources:
- arn:aws:dynamodb:table/table

Data isolation with Amazon DynamoDB

Table per tenant (silo)




Shared tenant table (pool)

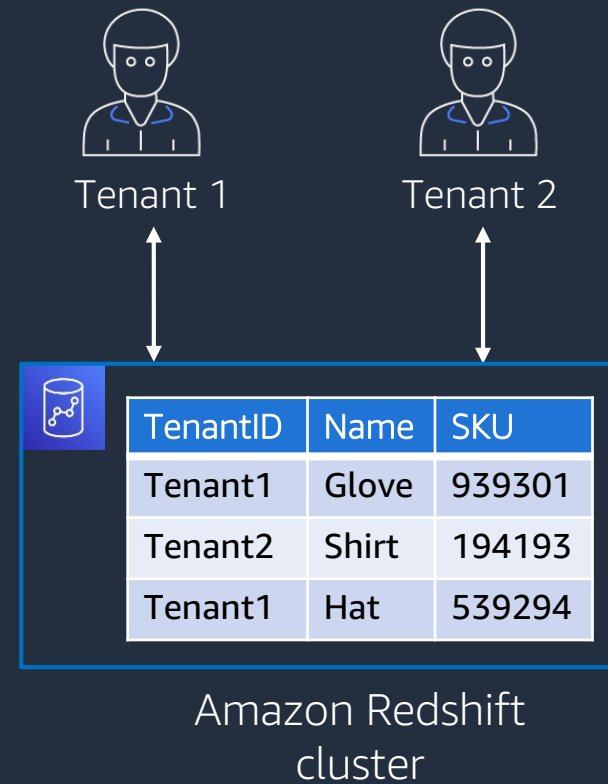


Data partitioning with Redshift

Cluster per tenant (silos)



Shared cluster for all tenants (pool)



A different strategy for each service



Amazon RDS




Amazon DynamoDB




Amazon Elasticsearch
Service



Amazon Simple
Storage Service (S3)



Amazon Timestream

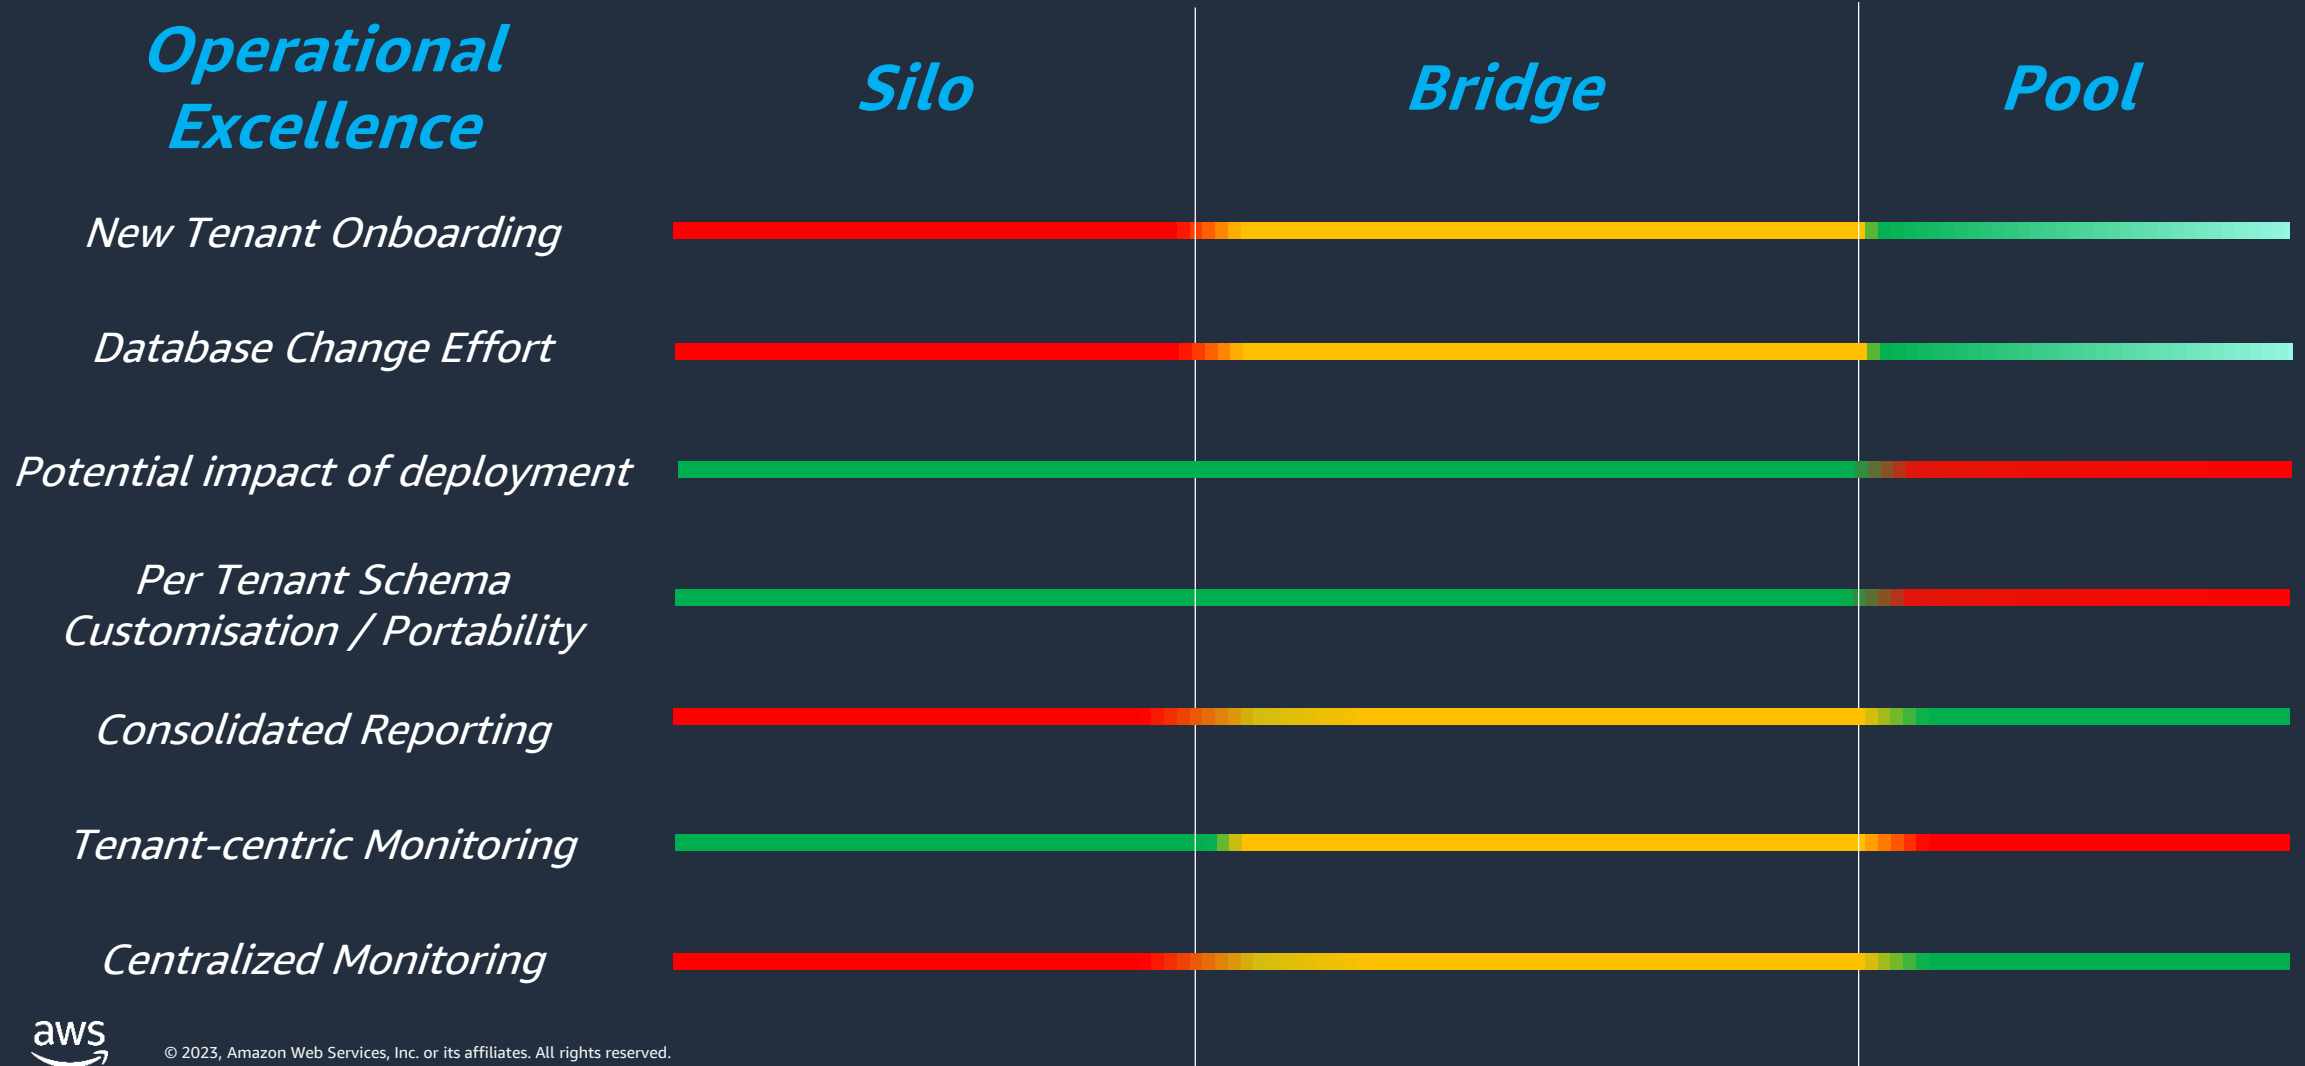


Amazon Redshift

Decision Matrix




Decision Matrix – Operational Excellence



Decision Matrix – Reliability



Decision Matrix – Performance Efficiency



Decision Matrix – Cost Optimization



Silo Model – Single vs Multi-tenant

Why people choose this model

- Compliance alignment
- Partitioned environments
- No cross-tenant impacts
- Tenant-specific tuning
- Tenant level availability

Challenges

- Cost
- Agility compromised
- Management complexity
- Deployment challenges
- Analytics/metering aggregation

Bridge/Pool Model- – Single vs Multi-tenant

Why people choose this model

- Agility
- Cost optimization
- Centralized management
- Simplified deployment
- Analytics/metering aggregation

Challenges

- Agility
- Tenant Data Distribution.
- Resistance due to “shared” nature of Pool.
- Cross-tenant impacts
- Compliance challenges

Design recommendations



Performance Considerations

Evaluate Workload

Connection Pooling

Partitions

Optimize for
Reads and
Writes

Avoid
SubTransactions

Avoid long
running
Transactions

Isolate Noisy
Tenants

HA/DR consideration

Configure HA for Critical Workloads

Backup Strategy

Business Continuity Plan

Proactive Monitoring

Plan downtime for upgradesHA/DR consideration

Takeaways

- Design with isolation in mind
- Authentication and Authorization \neq Isolation
- Factor scale and account limits into your isolation strategy
- Use decision matrix for isolation strategy on RDS & Aurora
- Use sizing considerations for tenant bin packing
- Validate that your isolation model is working
- Isolation is fundamental to success in a multi-tenant environment

Workshop outline

1: Setup at an AWS Event

2: Splitting the monolith

3: Multi-tenant data isolation


4: Database Operations with Infrastructure as Code

5: Data Migration & Portability

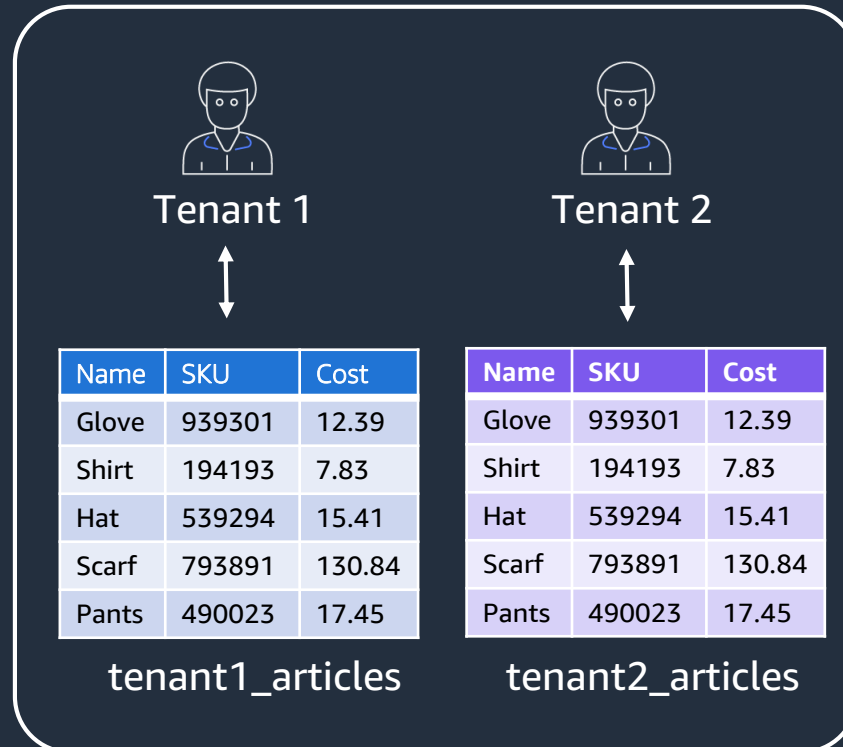
6: Cleanup

Multi-tenant data isolation Overview

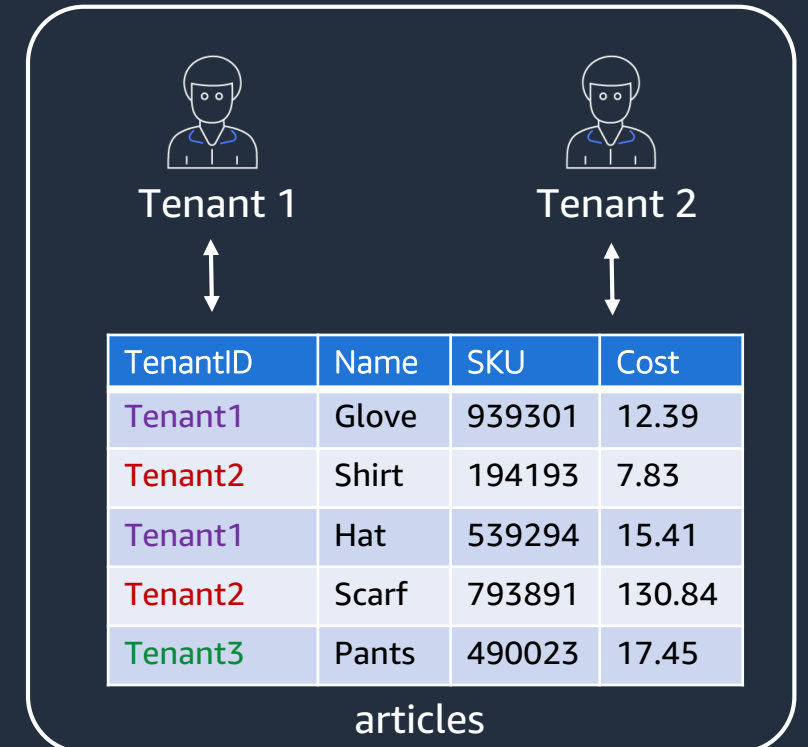
Silo model



Bridge model



Pool model

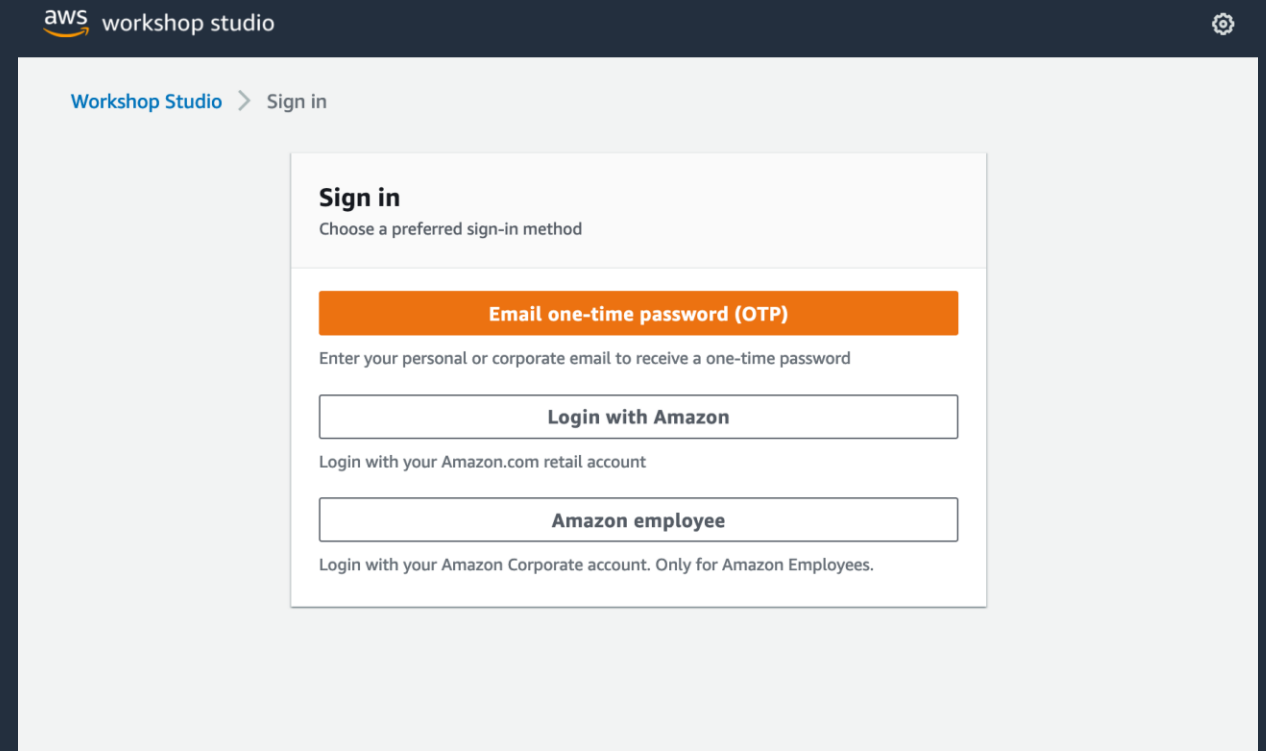


Max. isolation

Max. Resource Sharing

Step 1: Sign in via your preferred method

<https://catalog.workshops.aws/join>



Step 2: Enter event access code

Enter 12-digit event access code: **cf e8-07f6b4-48**

The screenshot shows the AWS Workshop Studio interface. At the top left, the 'aws workshop studio' logo is visible. In the top right corner, there are icons for settings and a user profile. The main content area is titled 'Workshop Studio > Join event'. On the left side, there is a progress indicator with 'Step 1 Enter event access code' (highlighted) and 'Step 2 Review and join'. The main heading is 'Enter event access code'. Below this, there is a section titled 'Event access code' with a sub-heading 'Event access code' and a description 'A 12 digit code that was given to you for this event'. A text input field is provided for the code. At the bottom right of the form, there are 'Cancel' and 'Next' buttons. The footer contains the copyright notice '© 2008 - 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.' and links for 'Privacy policy' and 'Terms of use'.

Step 3: Review terms and join event

aws workshop studio

Workshop Studio > Join event

Step 1
[Enter event access code](#)

Step 2
Review and join

Review and join

Event details

Name	Start time	Duration	Level
AWS General Immersion Day	9/23/2022 01:13 AM	12 hours	-

Description
AWS General Immersion Day

Terms and Conditions

Read and accept before joining the event

1. By using AWS Workshop Studio for the relevant event, you agree to the AWS Event Terms and Conditions and the AWS Acceptable Use Policy. You acknowledge and agree that are using an AWS-owned account that you can only access for the duration of the relevant event. If you find residual resources or materials in the AWS-owned account, you will make us aware and cease use of the account. AWS reserves the right to terminate the account and delete the contents at any time.
2. You will not: (a) process or run any operation on any data other than test data sets or lab-approved materials by AWS, and (b) copy, import, export or otherwise create derivate works of materials provided by AWS, including but not limited to, data sets.
3. AWS is under no obligation to enable the transmission of your materials through Event Engine and may, in its discretion, edit, block, refuse to post, or remove your materials at any time.
4. Your use of AWS Workshop Studio will comply with these terms and all applicable laws, and your access to AWS Workshop Studio will immediately and automatically terminate if you do not comply with any of these terms or conditions.

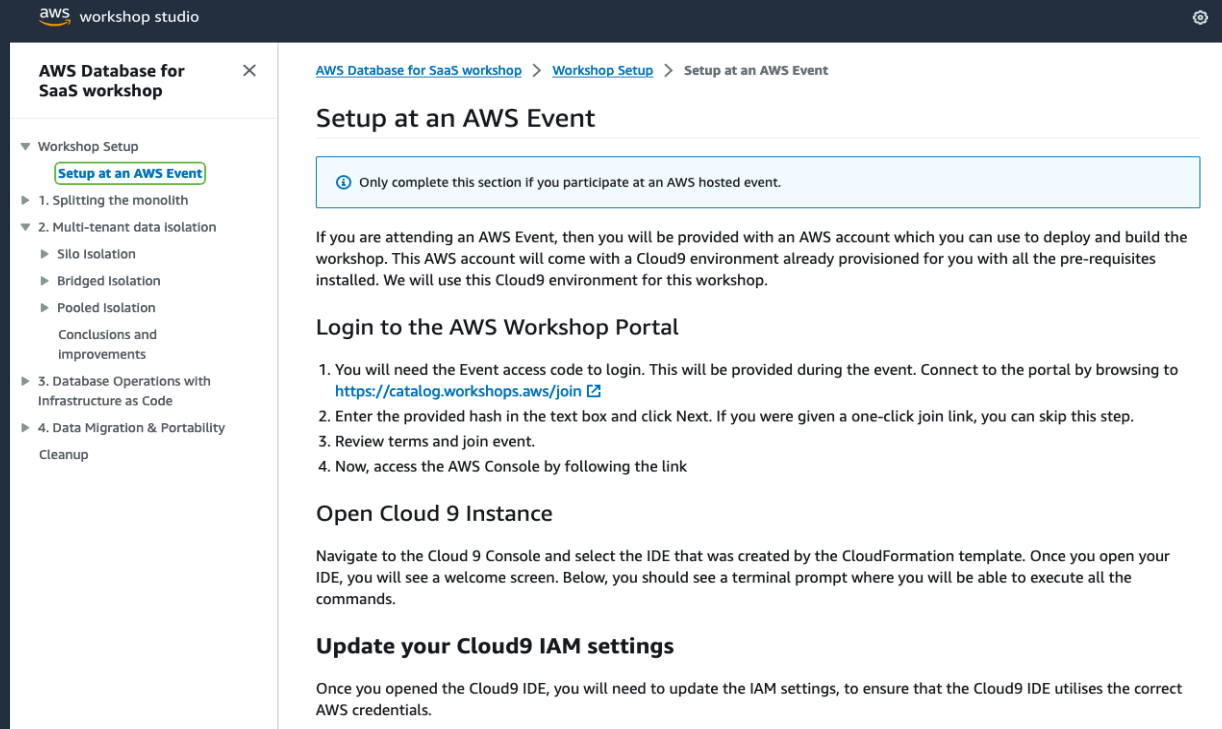
I agree with the Terms and Conditions

Cancel Previous **Join event**

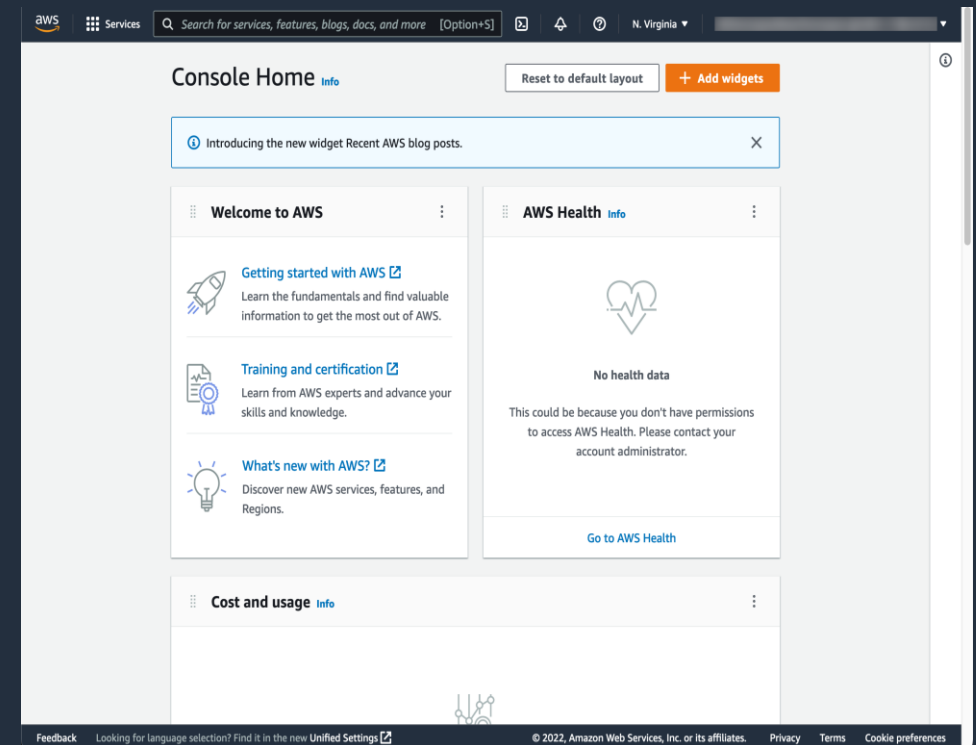
© 2008 - 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved. [Privacy policy](#) [Terms of use](#)

Step 4: Access AWS account

Access the AWS Management Console, or generate AWS CLI credentials as needed



The screenshot shows the 'AWS Database for SaaS workshop' documentation page. The left sidebar lists the workshop setup steps, with 'Setup at an AWS Event' highlighted. The main content area is titled 'Setup at an AWS Event' and includes a note: 'Only complete this section if you participate at an AWS hosted event.' Below this, it explains that an AWS account will be provided for the workshop. The section 'Login to the AWS Workshop Portal' lists four steps: 1. Obtain the Event access code, 2. Enter the hash or use a one-click link, 3. Review terms, and 4. Access the AWS Console. The 'Open Cloud 9 Instance' section instructs users to navigate to the Cloud 9 Console and select the IDE. The 'Update your Cloud9 IAM settings' section notes that IAM settings must be updated to use the correct AWS credentials.



The screenshot shows the AWS Management Console Home page. The top navigation bar includes the AWS logo, 'Services', a search bar, and the region 'N. Virginia'. The main content area features a 'Console Home' header with 'Reset to default layout' and '+ Add widgets' buttons. A notification banner at the top reads 'Introducing the new widget Recent AWS blog posts.' Below this, there are three widget panels: 'Welcome to AWS' with links for 'Getting started with AWS', 'Training and certification', and 'What's new with AWS?'; 'AWS Health' showing 'No health data' and a 'Go to AWS Health' button; and 'Cost and usage'.