



# Integrating Amazon RDS for MySQL with Amazon Redshift, AWS Glue and Amazon QuickSight for Analytics and Reporting

**Anuja Malik**

Sr. Database Specialist SA

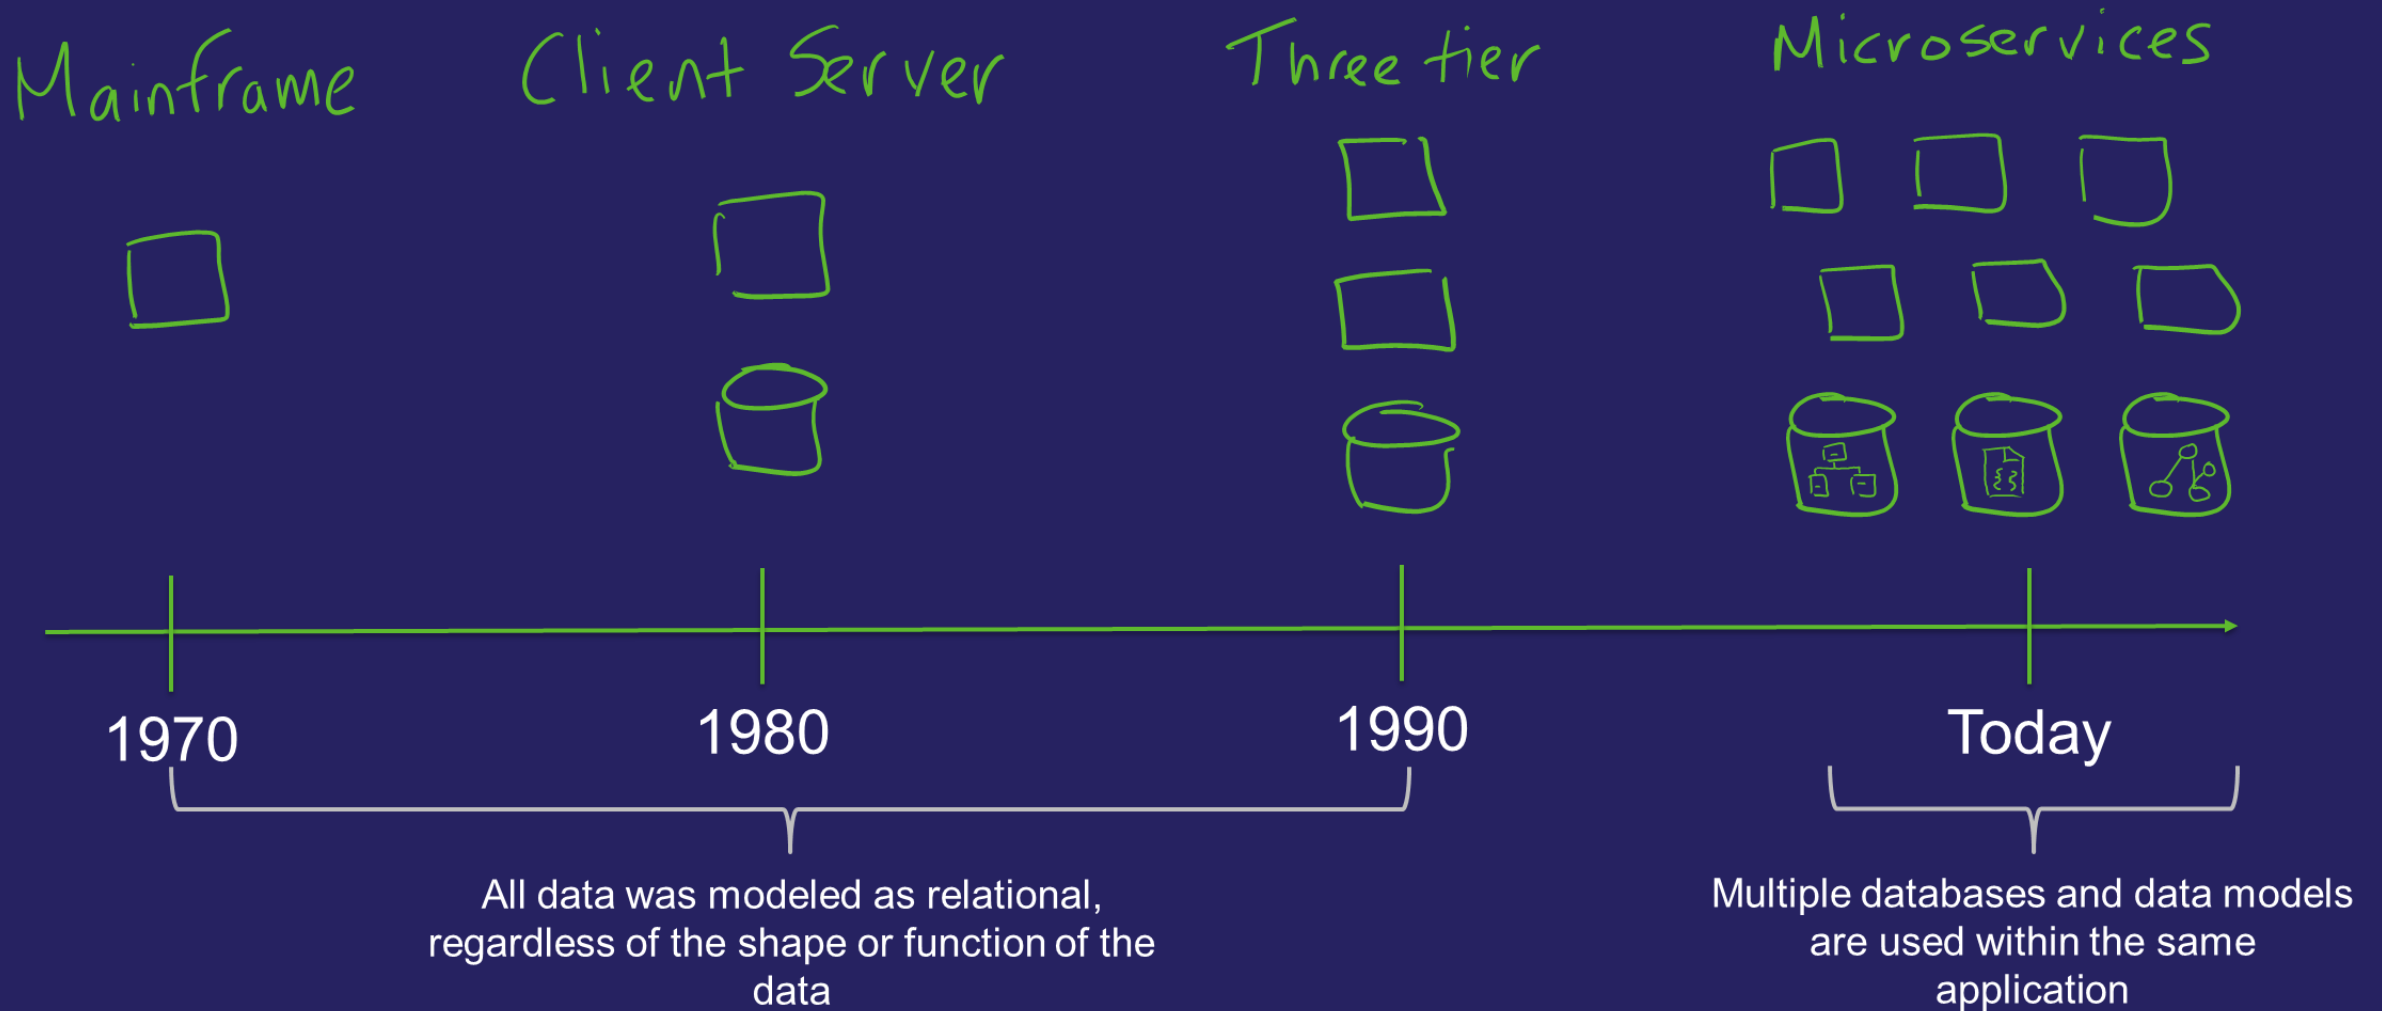
**Phil Bates**

Sr. Analytics Specialist SA

# A Modern Data Strategy



# App architectures and patterns have evolved



# A modern data strategy

Data at any scale

The best price performance


Seamless data access

Unified governance

AI and ML to solve business challenges



# Decoupling your data



# Decoupling data into each service improves resilience




# Purpose-built databases

*What are the data needs of each modular service?*



Massive data volume?



High-speed rendering?



Data warehousing?

**One size does not fit all**

# AWS delivers the right database for the right job



## Relational

Referential integrity, ACID transactions, schema-on-write




## Key value

High-throughput, low-latency reads and writes; endless scale



## Document

Store documents and quickly access querying on any attribute



## In memory

Query by key with microsecond latency



## Graph

Quickly and easily create and navigate relationships between data



## Time series

Collect, store, and process data sequenced by time



## Ledger

Complete, immutable, and verifiable history of all changes to application data



## Wide column

Scalable, highly available, and managed Apache Cassandra-compatible service

### AWS service(s)



RDS



Aurora Redshift



DynamoDB



DocumentDB



ElastiCache Memory DB



Neptune



Timestream



QLDB



Keyspaces Managed Cassandra

### Common use cases

Lift-and-shift, ERP, CRM, finance

Real-time bidding, shopping cart, social, product catalog, customer preferences

Content management, personalization, mobile

Leaderboards, real-time analytics, caching

Fraud detection, social networking, recommendation engine

IoT applications, event tracking

Systems of record, supply chain, health care, registrations, financial

Build low-latency applications, leverage open source, migrate Cassandra to the cloud



# Amazon Relational Database Service (Amazon RDS)

Managed relational database service with a choice of 7 popular database



PostgreSQL-Compatible Edition



MySQL-Compatible Edition



**Easy to administer**



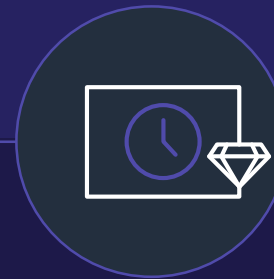
Easily deploy and maintain hardware, OS, and DB software, with built-in monitoring

**Secure and compliant**




Data encryption at rest and in transit, with industry compliance and assurance programs

**Available and durable**



Automatic Multi-AZ data replication, with automated backup, snapshots, and failover

**Performant and scalable**



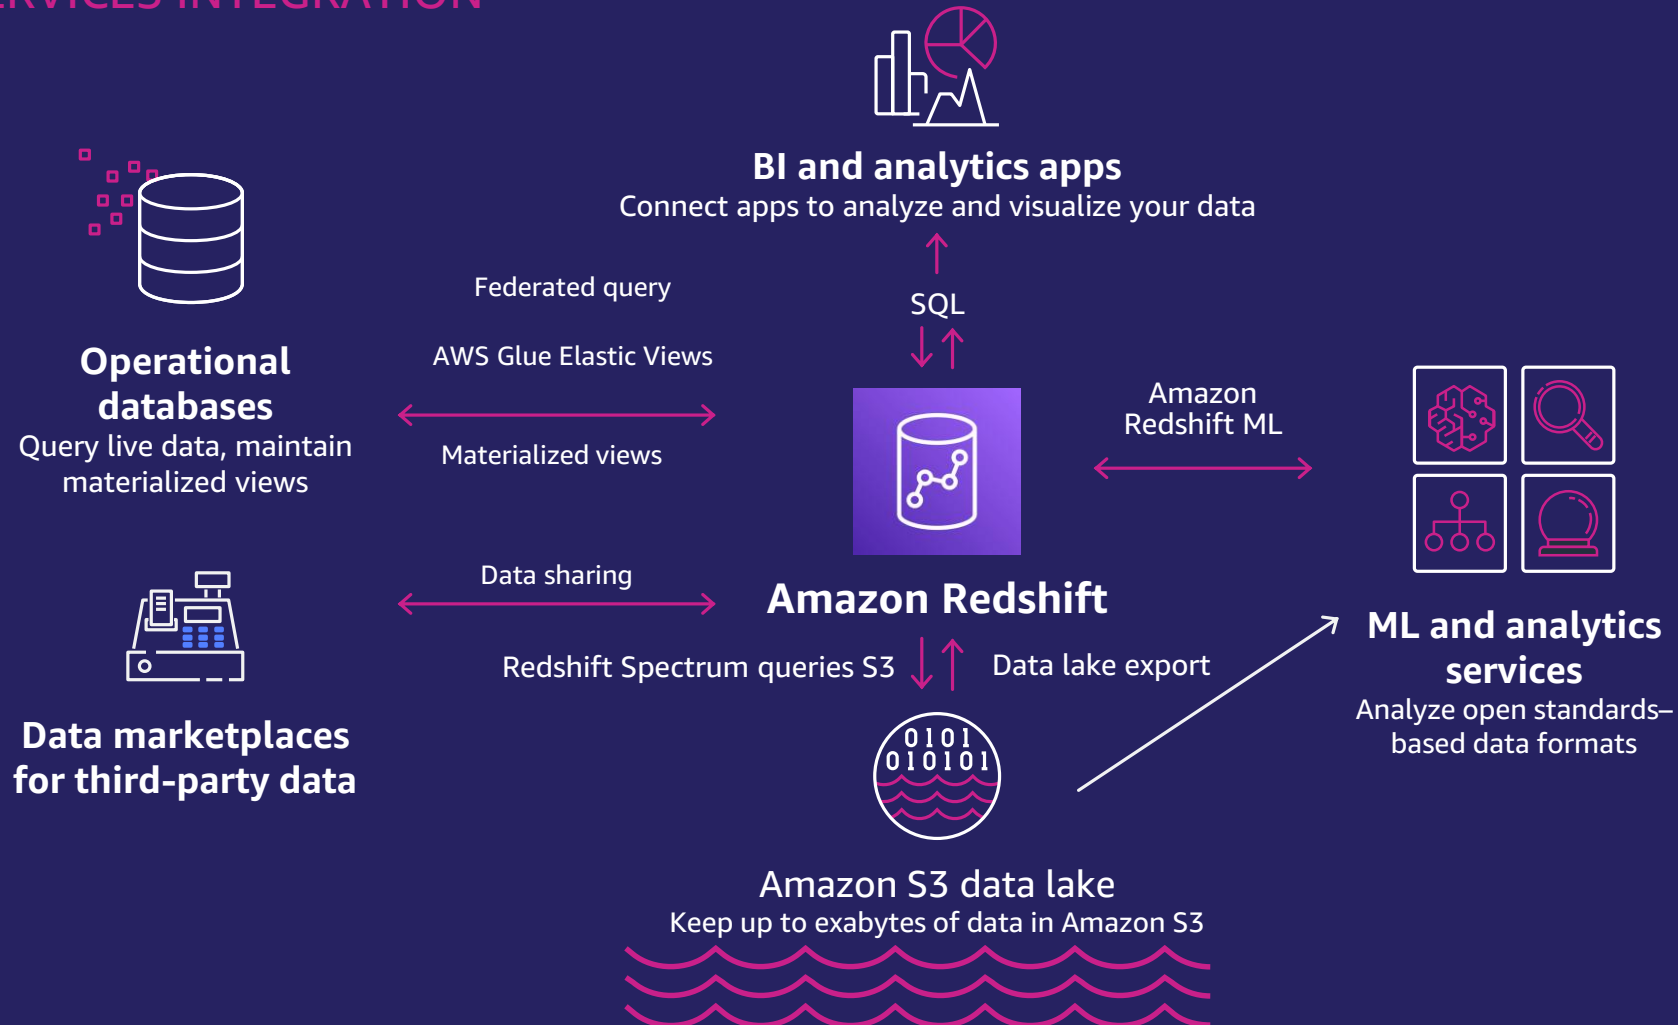
Scale compute and storage with a few clicks, plus minimal downtime for your application

# Redshift Overview



# Analyze all your data

WITH AWS SERVICES INTEGRATION



# What is the Amazon Redshift service?



*Not a standalone data warehouse, but a data warehouse that breaks down the silos to keep data “free”*

Fully Managed	Automated maintenance & workload management; cost-effective cloud data warehouse	Scale	Query GBs to EBs; auto scaling; independent compute/storage scaling
Superior Speed	Extensive machine learning based optimizations and features	Highly-rated & Most Popular	Tens of thousands of deployments; highly-rated by agencies such as Gartner
Highly-resilient	Service SLA: 99.9%	RDS, ML, and Data Lake Integration	Query data in-place in your data lake and RDS; ACID and ANSI SQL
Secure	End-to-end encryption; SSO; compliance with SOC 1/2/3, HIPAA, FedRamp & more	Commitment to Open Formats	Query open formats in place Import/Export Parquet & CSV; Query ORC, Avro, JSON, ...

# Most comprehensive and open Analytics portfolio

## Data, visualization, engagement, & machine learning

NEW



Data Exchange



QuickSight



Pinpoint



SageMaker



Comprehend



Lex



Polly



Rekognition



Translate

+ many more

## Analytics



Redshift



EMR (Spark & Hadoop)



AWS Glue  
(Spark & Python)



Athena



Elasticsearch  
Service



Kinesis Data  
Analytics

## Data lake infrastructure & management



S3/Glacier



Lake  
Formation



AWS Glue

## Data movement

Database Migration Service | Snowball | Snowmobile | Kinesis Data Streams | Kinesis Data Firehose | Managed Streaming for Apache Kafka



# Redshift Seamless Integration



# Seamless Integration with AWS Services

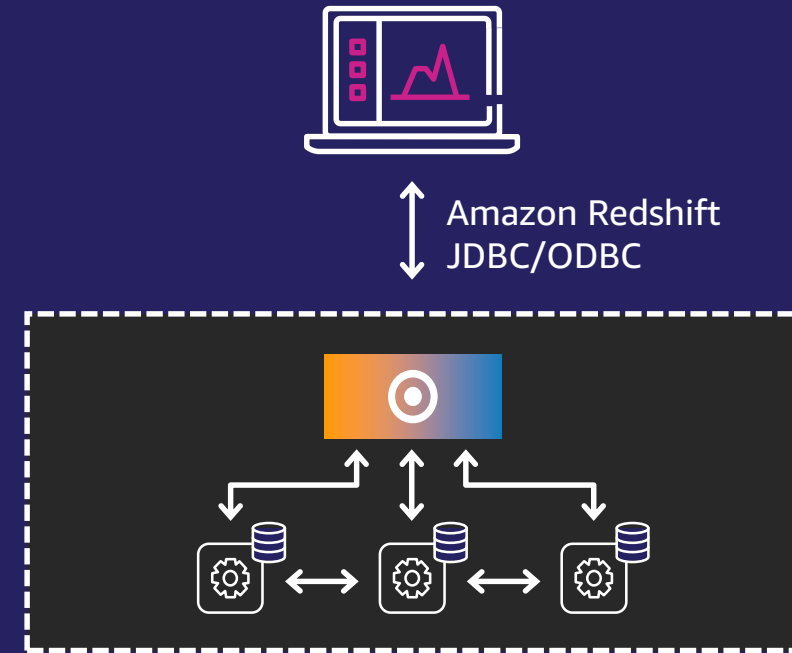
———— Migration | ETL | Machine Learning | Data Visualization ————

- AWS Database Migration Service (DMS)  
*Fully-managed migration service*
- AWS Glue  
*Fully managed extract, transform, and load (ETL) service*
- Amazon EMR  
*Cloud-native big data platform*
- Amazon Sagemaker  
*Fully managed service to build, train, and deploy machine learning (ML) models*
- Amazon Quicksight  
*Fast, cloud-powered business intelligence service*  
*and more...*

# Amazon Redshift Federated Query

## Unified analytics across databases, data warehouse, and data lake

- ✓ Use case: Integrate operational data with DW and data lake for real-time analytics
- ✓ Analytics on operational data without data movement and ETL delays
- ✓ Query and join data from one or more Amazon RDS and Aurora PostgreSQL databases
- ✓ Flexible and easy way to ingest data avoiding complex ETL pipelines
- ✓ Intelligent distribution of computation to remote sources to optimize performance
- ✓ Amazon RDS for MySQL and Amazon Aurora MySQL Compatible Edition support



Amazon Aurora  
MySQL-  
Compatible,  
PostgreSQL-  
Compatible



Amazon S3  
data lake



Amazon RDS  
MySQL-  
Compatible,  
PostgreSQL-  
Compatible



# Demo

- RDS for MySQL
- Amazon Redshift Federated Query
- Amazon Redshift
- AWS Glue
- Amazon QuickSight

# Q&A



# Resources

- [RDS Product Documentation](#)
- [Amazon Redshift Documentation](#)
- [Amazon Redshift Federated Query Blog](#)



**Thank you!**