



# Amazon Redshift Data Sharing

Build scalable multi-tenant architectures

**Milind Oke**

Analytics Specialist Solutions Architect  
Amazon Web Services

# Introduction to Amazon Redshift



# Why Amazon Redshift for your data needs?



FAST, EASY, AND SECURE ANALYTICS AT SCALE

Easy analytics  
for everyone



Focus on getting from data to insights in seconds without worrying about infrastructure

Analyze all  
your data



Get insights running real-time and predictive analytics on complex, scaled data across your operational databases, data lake, and data warehouse

Best price  
performance  
at any scale



Gain up to 3x better price performance than other cloud data warehouses, and dynamically scale to improve query speed for complex and critical workloads

# Amazon Redshift innovates to meet your needs



## Easy analytics for everyone

NEW!



Serverless

NEW!



Query Editor v2

Updated!



Automated DW Management

NEW!



Automatic Materialized Views



Data API

Updated!



Redshift Advisor

Updated!



Cloud Formation templates

NEW!



Grafana Plugin



## Analyze all your data

Updated!



Data sharing

NEW!



ADX Integration

NEW!



Streaming Ingestion

Updated!



Amazon Redshift ML

Updated!



Federated Query

Updated!



Geospatial enhancements



SUPER data type with JSON



## Performance at any scale

Updated!



RA3 nodes & managed storage



AQUA

Updated!



Workload Management enhancements

NEW!



Concurrency scaling for Writes

Updated!



SQL Enhancements & migration support



Security, Governance, & Compliance



# 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, ...

# Redshift cluster architecture



## Leader node

- SQL endpoint
- Stores metadata
- Coordinates parallel SQL processing & ML optimizations
- Leader node is no-charge for clusters with 2+ nodes

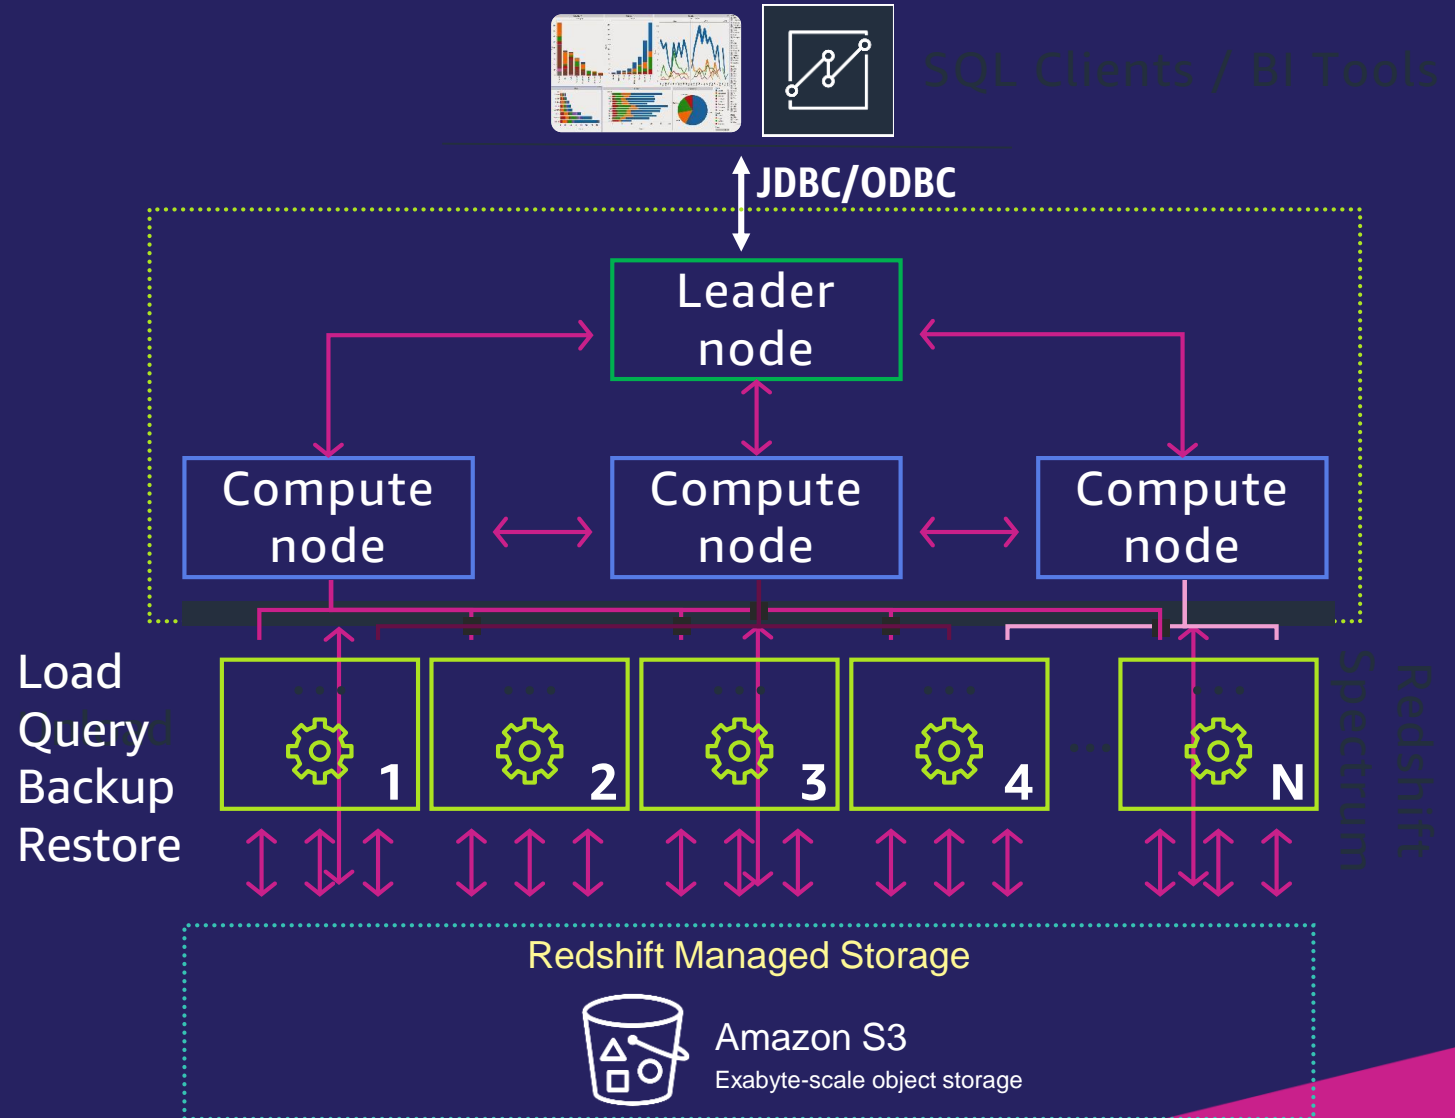
## Compute nodes

- Local, columnar storage
- Executes queries in parallel
- Load, unload, backup, restore from S3

## Amazon Redshift Spectrum nodes

- Execute queries directly against data lake

## Massively parallel, shared nothing architecture



# Amazon Redshift Data Sharing



# Data Sharing



## A SECURE AND EASY WAY TO SHARE DATA ACROSS YOUR AMAZON REDSHIFT CLUSTERS

Instant, granular, high-performance data access without data copies/movement

Live and consistently updating views of data across all consumers

Secure and governed collaboration within and across organizations

Workloads accessing shared data are isolated from each other

**NEW** Cross-account\* data sharing

**NEW PREVIEW** Cross-Region\* data sharing

**NEW** Performance Enhancements: Result Caching & Concurrency scaling for consumer clusters





# Data Sharing – Steps for **producer** cluster



Amazon Redshift > Databases > ds\_finance

## ds\_finance [Info](#)

**Action required**  
To authorize or remove authorization for this datashare, choose one or more data consumers from the **Data consumers** table.

### General information

Datashare name ds_finance	Publicly accessible Enabled	Datashare type Datashare
Datashare status <b>Action required</b>	Producer namespace 293ef5dd-377e-4...	Amazon Resource Name (ARN) arn:aws:redshift:us-east-2:497316421912:datashare:293ef5dd-377e-4...

### Data consumers (1/1) [Info](#)

The following data consumers can access this datashare. To add or remove data consumers, edit the datashare.

<input checked="" type="checkbox"/>	Data consumer ID	Consumer type	Consumer added date	Status change date	Consumer status
<input checked="" type="checkbox"/>	497316421912	Account	January 14, 2022, 16:37 (UTC-05:00)	January 14, 2022, 16:37 (UTC-05:00)	Pending authorization

**Remove authorization** **Authorize**

### Datashare objects [Info](#)

Objects in this datashare shown below are read-only.

**Schemas** | Tables and views | User-defined functions

#### Schemas (0)

1 schema, 0 tables and views, 0 user-defined functions selected

# Data Sharing – Steps for **consumer** cluster



The screenshot displays the Amazon Redshift console interface. On the left, the navigation pane shows 'Amazon Redshift provisioned clusters' and 'Query editor'. The main area is divided into three sections:

- Database Selection:** 'Select database' is set to 'db\_ds\_finance' (highlighted with a red circle). 'Select schema' is set to 'pivot\_schema'. Below, a table list shows 'pivot\_source' expanded with columns: 'primary\_key', 'pivot\_key', and 'pivot\_value'.
- Query Editor:** The SQL query is: 

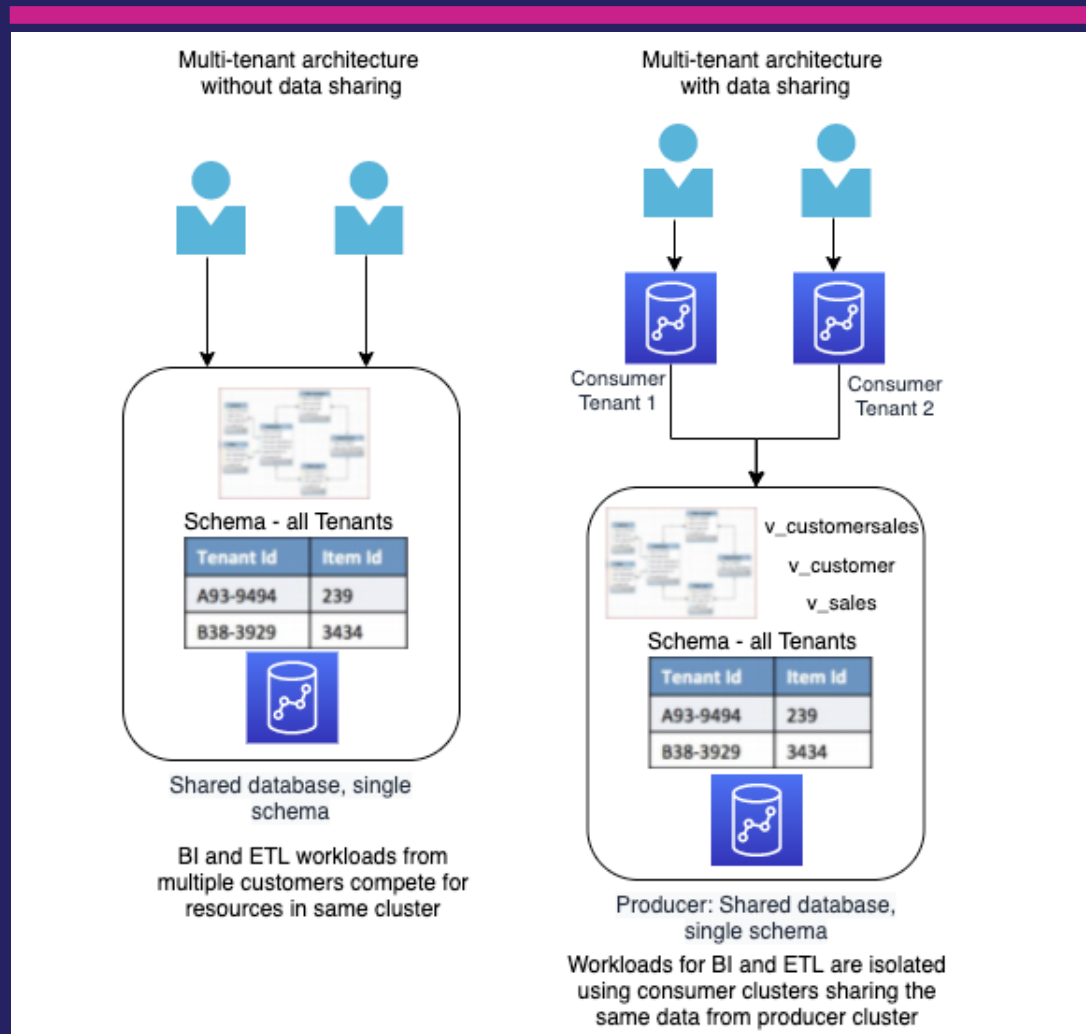
```
1 SELECT * FROM "db_ds_finance"."pivot_schema"."pivot_source";
2
3 SELECT *
4 FROM "db_ds_finance"."pivot_schema"."pivot_source"
5 PIVOT
6 ( MAX(pivot_value)
7   FOR pivot_key IN ('model100_binary','model100_rank')
8 )
9 ;
```

 The first line is highlighted with a red circle. Below the editor are buttons for 'Run', 'Save', 'Schedule', and 'Clear'.
- Query Results:** Shows 'Query 512' completed on January 14, 2022 at 18:12:46. It displays 6 rows returned. The results table is as follows:

primary_key	pivot_key	pivot_value
1007	model100_binary	200
1007	model100_rank	100

# Multi-tenant patterns in Amazon Redshift using Data Sharing

# Amazon Redshift Data Sharing via Pool Model



## Pool Model

Data is stored in a single database schema for all tenants, and a new column is used to scope and control access to individual tenant data. Access to the multi-tenant data is controlled using views.

1. Create new table TENANT and capture the consumer clusters namespace
2. Add TENANT ID to all tables in DWH
3. Create views for each table with logic to filter rows by each tenant's namespace
4. Create data share and add views to shares
5. Consumers will see row level filtered results

### Amazon Redshift provisioned clusters

Redshift serverless **New**

**Provisioned clusters dashboard**

Clusters

- Reserved nodes
- Snapshots

Query editor

- Query editor v2
- Queries and loads

Datashares

Configurations

- Advisor
- AWS Marketplace
- Alarms
- Events
- What's new

Get consistent, fast query performance for thousands of concurrent queries and users. [View tutorial](#)

Amazon Redshift > Provisioned clusters dashboard

## Provisioned clusters dashboard [Info](#)

[Try Amazon Redshift Serverless \(Preview\)](#) [Purchase reserved nodes](#) [Create cluster](#)

### Resources overview

Resource data for US East (N. Virginia) Region.

Total nodes	On-demand nodes	Reserved nodes	Reserved nodes available	Automated snapshots	Manual snapshots
4	4	0	0 (0 of 0 used)	2	0

### Cluster overview (2) Available

Cluster	Status
redshift-dmw-producer	Available

[View all clusters](#)

### Cluster metrics

[Any clusters](#) [Last hour](#) [View in CloudWatch](#)

- Number of queries** | Database connections | Disk space used | CPU utilization

10

### Datashares

Authorize other AWS accounts to access datashares created in this AWS account. Associate or decline datashares from other AWS accounts.

<a href="#">Require authorization</a>	<a href="#">Require association</a>
0	0

### Alarms (0)

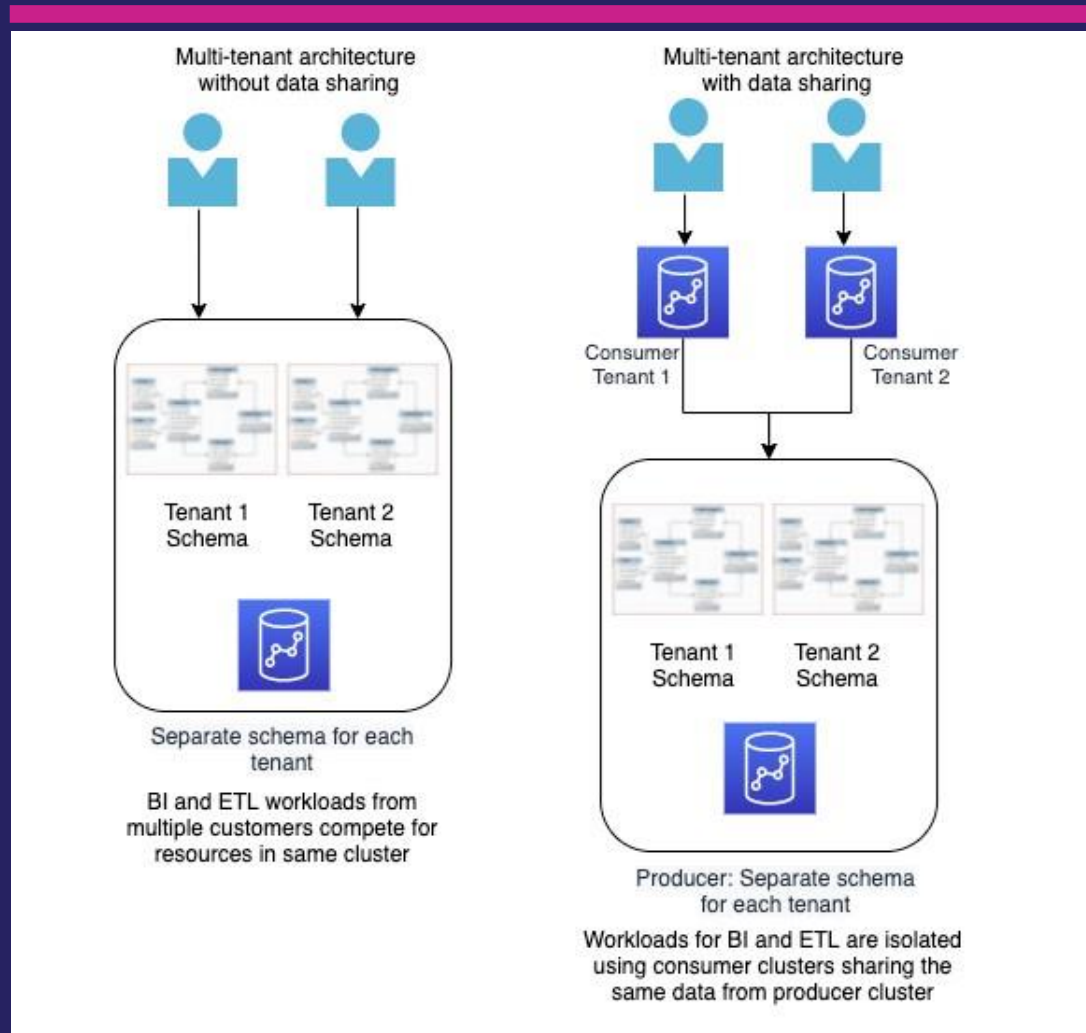
[View in CloudWatch](#)

Alarm name

No ongoing alarms

### Events (37)

# Amazon Redshift Data Sharing via Bridge Model

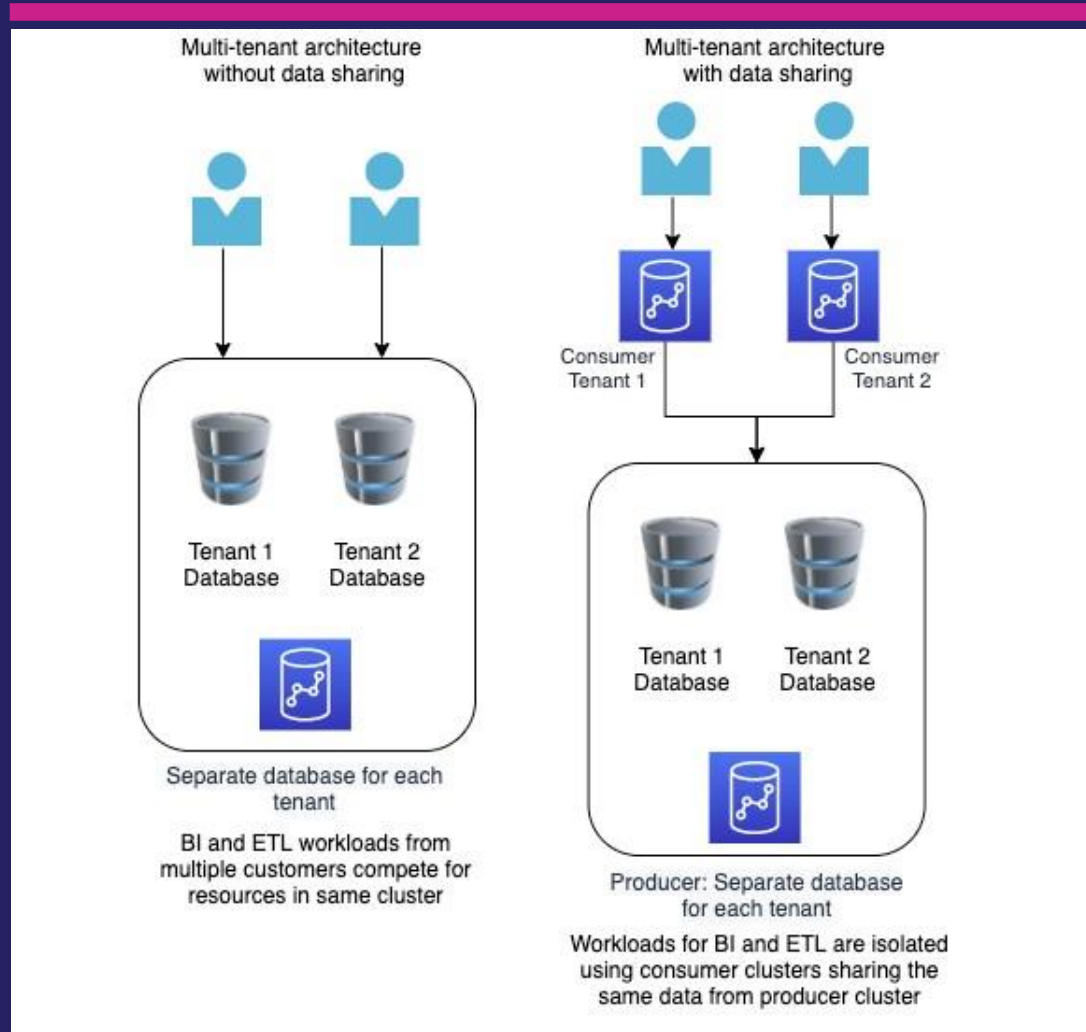


## Bridge Model

Storage and access to data for each tenant is controlled at individual schema level in the same database.

1. Create a schema for first tenant
2. Create separate schema for second tenant
3. Create all the tables in each schema individually
4. Create a datashare and add first schema
5. Create separate datashare and add the second schema

# Amazon Redshift Data Sharing via Silo Model



## Silo Model

Storage and access control to data for each tenant is maintained in separate databases.

1. Create a database for first tenant
2. Create separate database for second tenant
3. Create the schema with its tables in first database
4. Create separate schema with its tables in second database
5. Create a datashare and add first database
6. Create separate datashare and add the second database

# Considerations for choosing between Pool, Bridge, Silo

Things to consider	Pool	Bridge	Silo
Level of separation for tenant's data	Views	Schema	Database
Complexity of ETL pipelines	Low	Low	Medium
Quotas and Limits	100,000 tables* per cluster	9,900 schemas per database	60 databases per cluster
Chargeback to consumer accounts	Yes	Yes	Yes
Scalability	High	High	High

\* for clusters with RA3.4XL, and RA3.16XL node types

Refer blogs for details

<https://aws.amazon.com/blogs/big-data/implementing-multi-tenant-patterns-in-amazon-redshift-using-data-sharing/>

<https://aws.amazon.com/blogs/aws/cross-account-data-sharing-for-amazon-redshift/>

<https://aws.amazon.com/blogs/big-data/sharing-amazon-redshift-data-securely-across-amazon-redshift-clusters-for-workload-isolation/>





# AWS Data Exchange for Amazon Redshift



# AWS Data Exchange for Amazon Redshift



FIND, SUBSCRIBE TO, AND QUERY THIRD-PARTY DATA IN AN AMAZON REDSHIFT WAREHOUSE IN MINUTES



## AWS Data Exchange

### Data Provider

Automatic entitlements, billing, and payment management

Multiple data delivery methods on a single product



## Amazon Redshift

### Data Consumer

Query data in minutes, with no ETL required

Use the latest and greatest data as soon as its updated

## FACTSET

"We're really excited about the integration of Amazon Redshift and AWS Data Exchange. This will further remove friction associated with data onboarding. It's going to make things like billing more transparent and easier to manage for our clients. And it's going to really improve the user experience for those interested in browsing content and finding the data that they need, and sourcing it from trusted partners like FactSet."

**Peter Dorsey,**  
VP of Cloud Strategy,  
FactSet

# Console Home Info

Actions ▾

### Recently visited Info

- AWS Data Exchange
- Amazon Redshift
- IAM
- Key Management Service
- Systems Manager
- Trusted Advisor
- AWS Well-Architected Tool
- VPC
- Amazon SageMaker
- IAM
- Personal Health Dashboard
- CloudFormation
- AWS SQL Workbench
- Billing

### Welcome to AWS

- Getting started with AWS**  
Learn the fundamentals and find valuable information to get the most out of AWS.
- Training and certification**  
Learn from AWS experts and advance your skills and knowledge.
- What's new with AWS?**  
Discover new AWS services, features, and Regions.

### AWS Health Info

Open issues: **0** Past 7 days

Scheduled changes: **0** Upcoming and past 7 days

Other notifications: **0** Past 7 days

### Cost and usage Info

Current month costs: **\$28.38**

Forecasted month end costs: **\$51.06** Down 99% from last month

Last month costs: **\$4,683.62**

### Top costs for current month

Amazon QuickSight	\$15.55
Amazon SageMaker	\$12.54
Amazon Simple Storage Service	\$0.23
Amazon GuardDuty	\$0.04
Amazon Redshift	\$0.03



# Thank you!