AWS re:Invent

NOV. 29 – DEC. 3, 2021 | LAS VEGAS, NV

ANT216

Introducing Amazon Redshift Serverless

Yan Leshinsky
VP Amazon Redshift
AWS

Neeraja Rentachintala
Principal Product Manager
AWS

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda



Challenges and business opportunities with data



Amazon Redshift evolution



Amazon Redshift Serverless : introduction, key capabilities



Amazon Redshift Serverless: use cases, and benefits



Demo



Summary

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

“In 2020, 64.2ZB of data was created or replicated. . . . The amount of digital data created over the next five years will be greater than twice the amount of data created since the advent of digital storage.”

Source: IDC, Worldwide Global DataSphere Forecast, 2021–2025, Doc # US46410421, March 2021



“Analytics requirements are coming in from all corners; we don’t have enough specialized skills”



“How do I derive value from growing, diverse data living in data silos?”

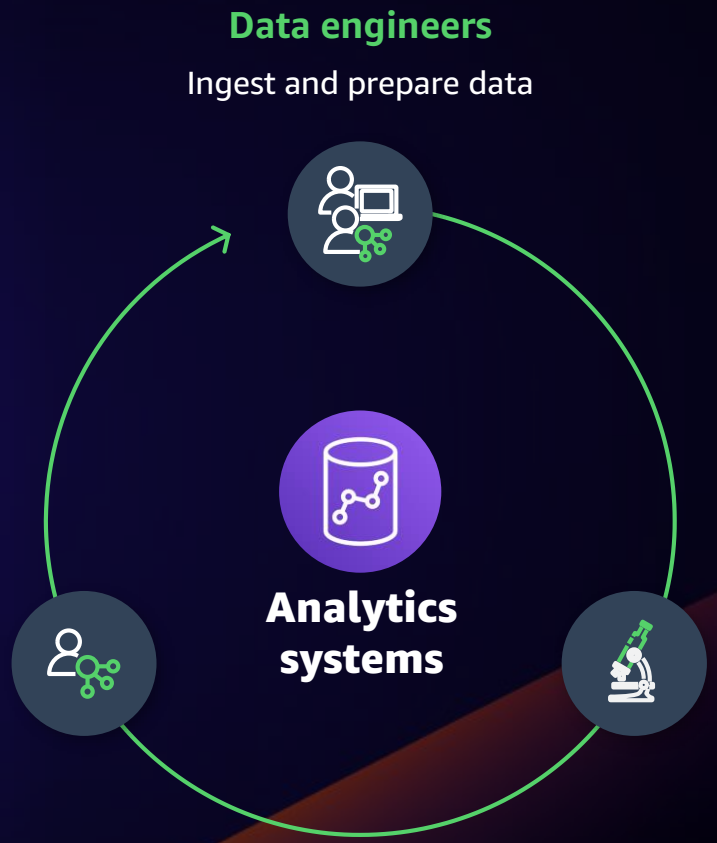


“Performance has to be consistently fast; costs are hitting the roof”

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data can work more effectively for diverse users through easy-to-deploy, self-serve, and auto-scaling analytics systems



Improve operational efficiency; make more informed decisions; accelerate innovation

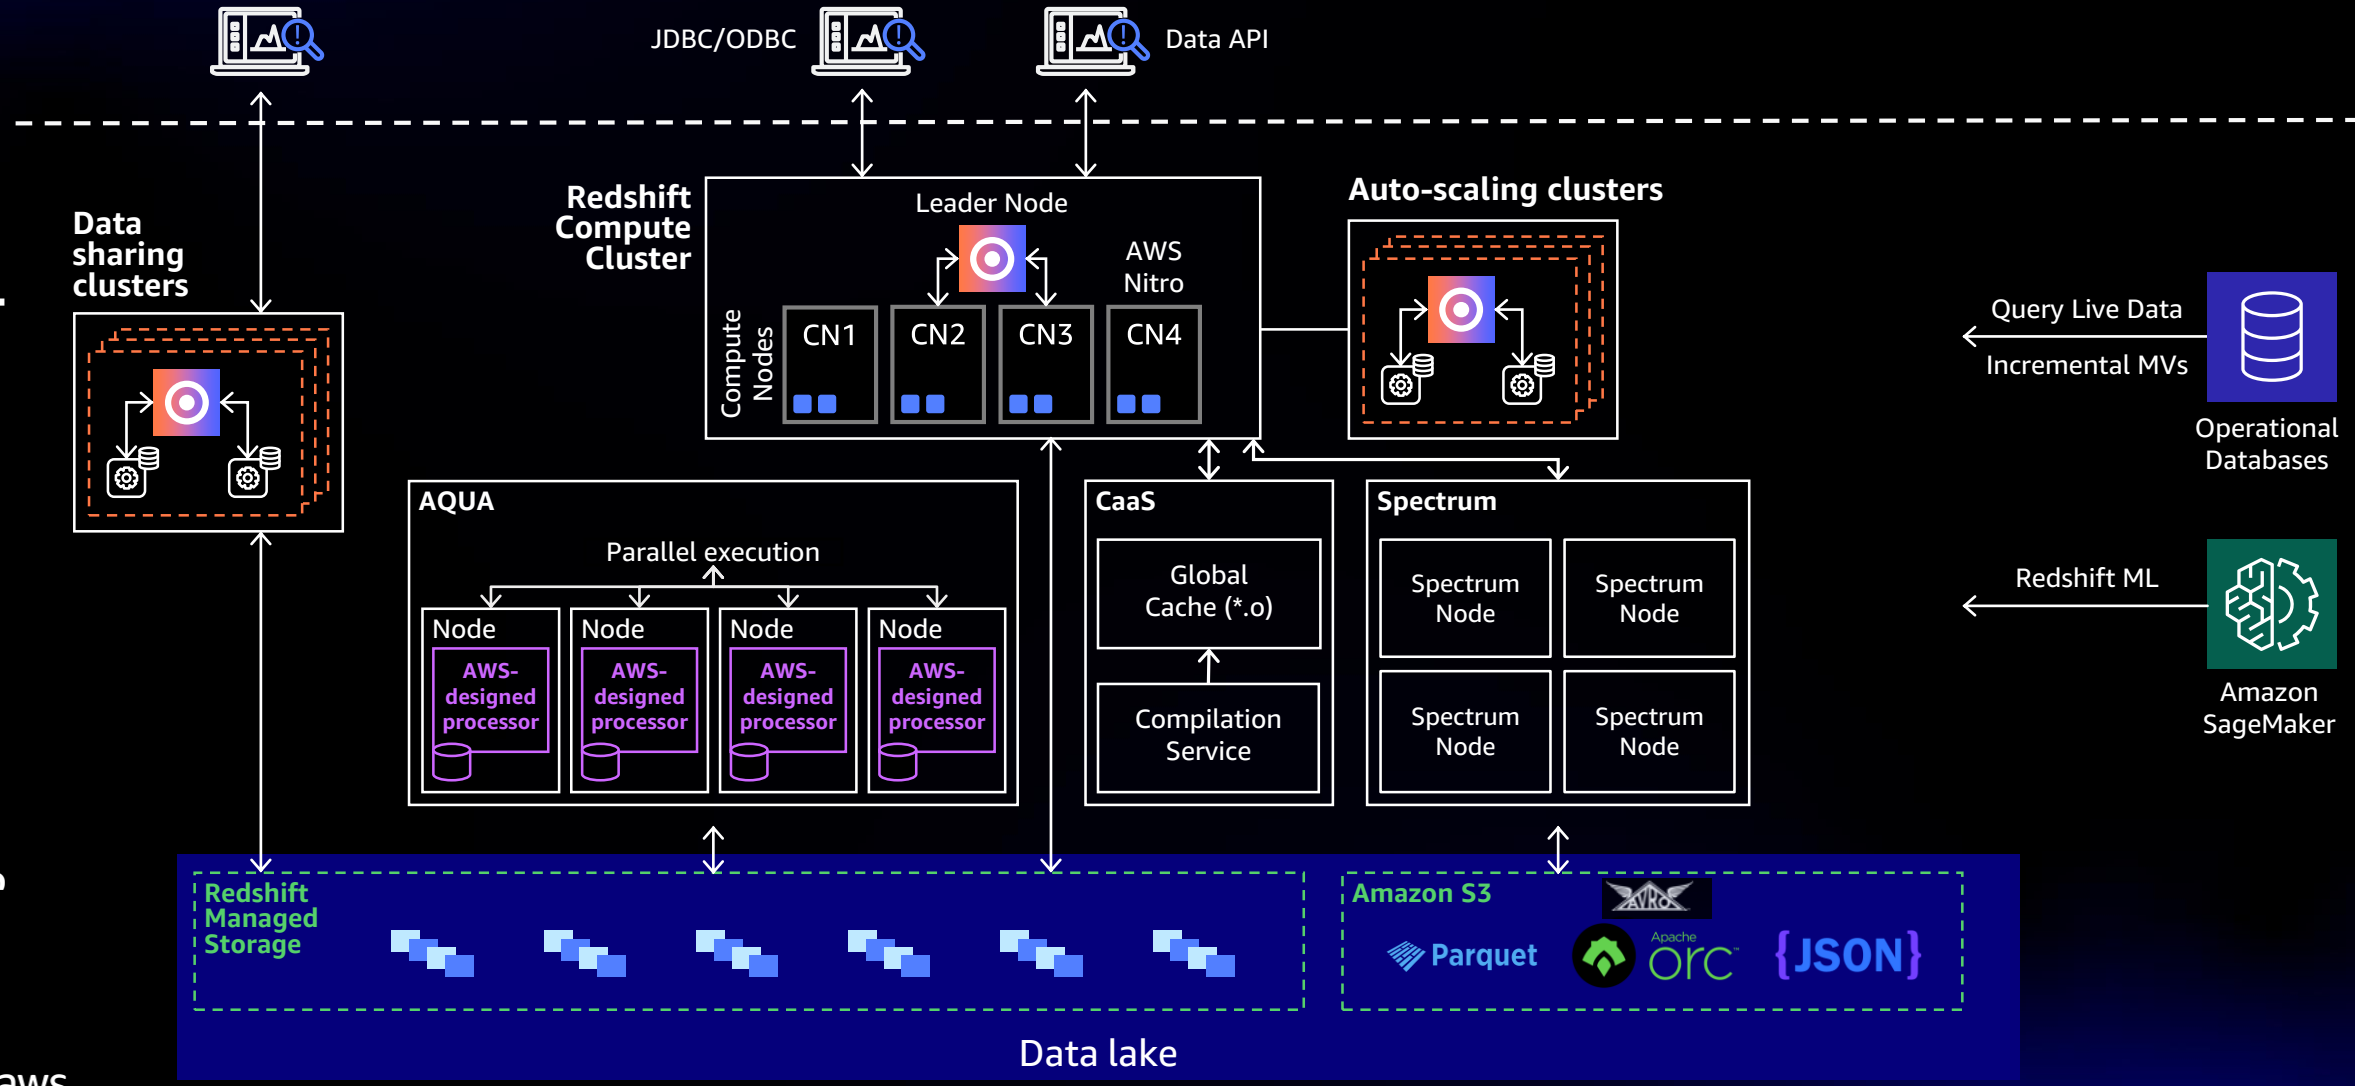
By making 10% more data accessible, a typical Fortune 1000 company will see a ***$65 million increase in net income***[1]

[1] Dykes, "The Four Key Pillars to Fostering a Data-Driven Culture"

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift architecture



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Tens of thousands of customers process exabytes of data with Amazon Redshift daily



Zynga

Doubled ETL performance and scaled to process >5.3TB of game data daily



Schneider Electric

Supports tens of thousands of users through Redshift concurrency scaling and RA3 nodes to support a green future



Magellan Rx

Redshift ML for predictive analytics to forecast future drug costs and identify trends



SiriusXM

Moved ~80TB of reporting and ETL workloads to Amazon Redshift as part of its AWS migration, saving $1M annually



NASDAQ

Easily support the jump from 30 billion records to 70 billion records a day because of the flexibility and scalability of Amazon S3 and Amazon Redshift



aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why Amazon Redshift for your data needs

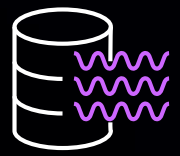
FAST, EASY, AND SECURE ANALYTICS AT SCALE

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

Analyze all your data



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

Easy analytics for everyone



Focus on getting from data to insights in seconds and delivering on your business outcomes, without worrying about infrastructure

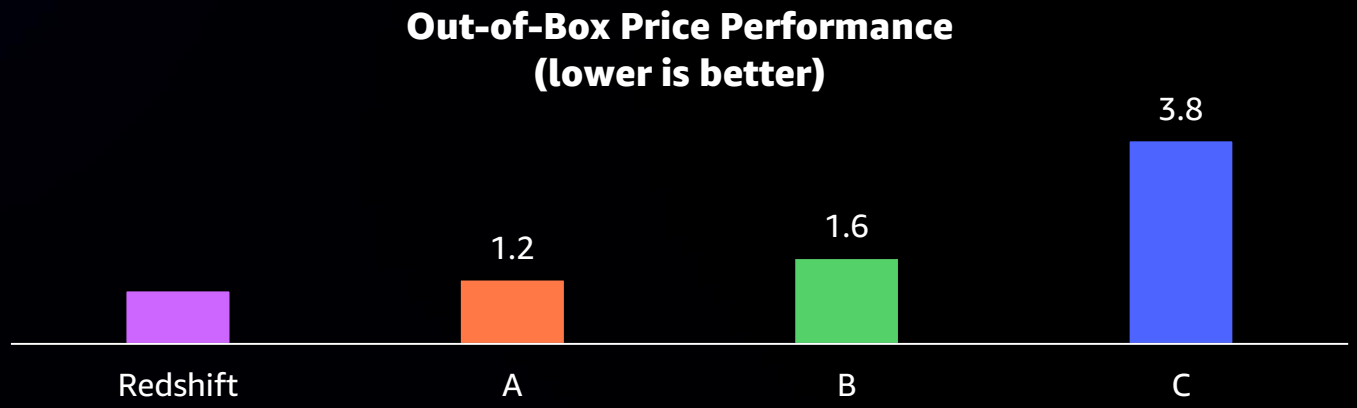
aws

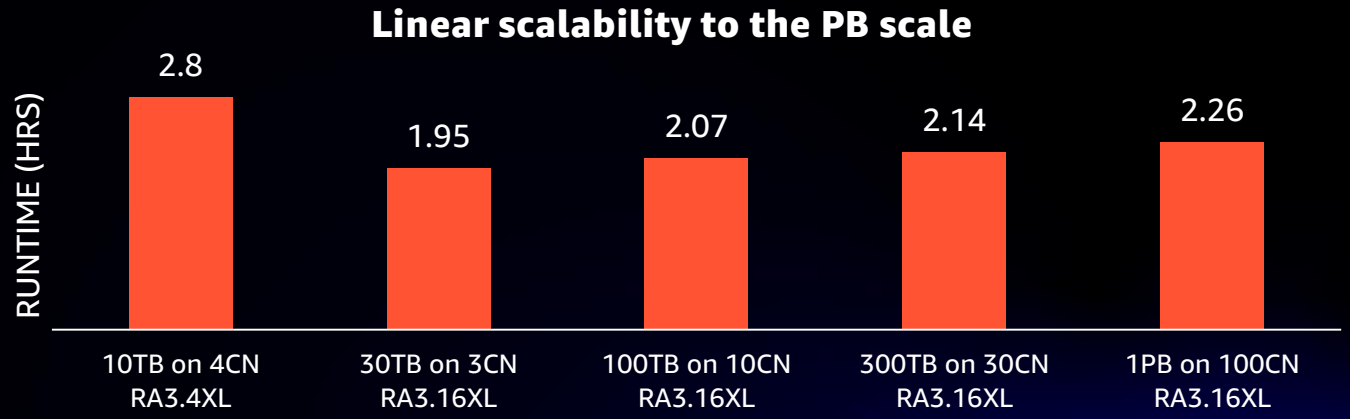
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Performance at any scale

WHILE KEEPING YOUR COSTS PREDICTABLE

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



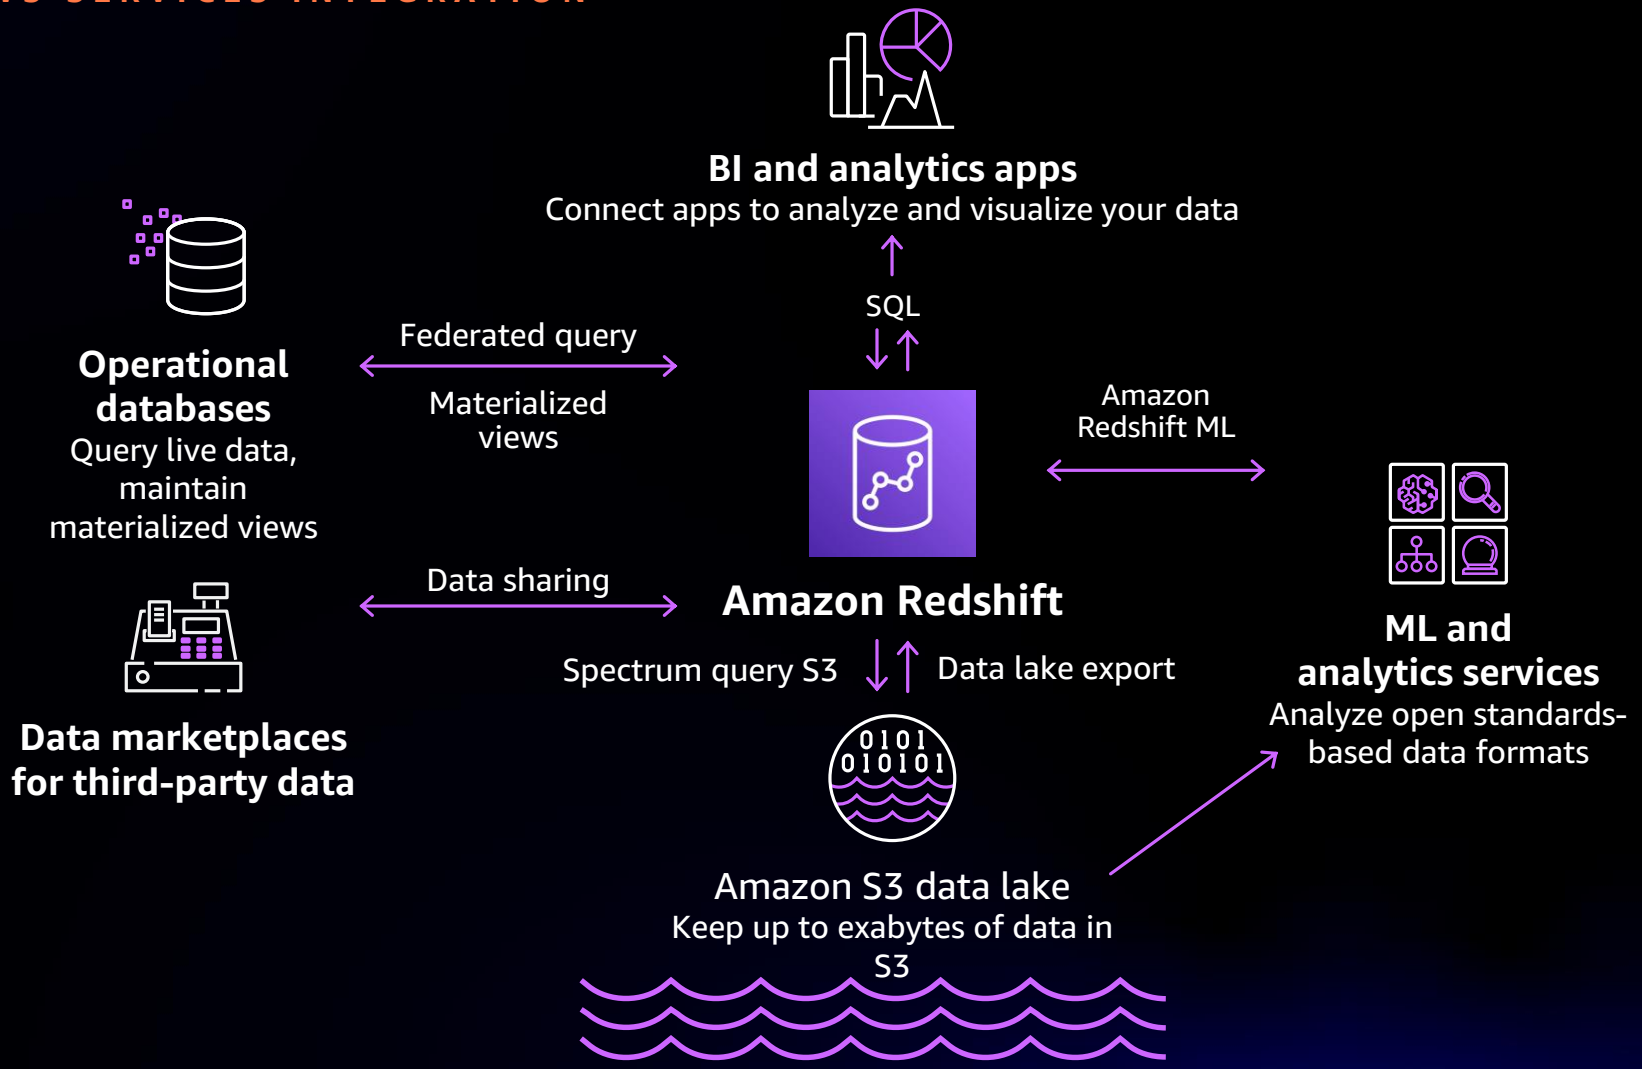


© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Analyze all your data

CUT THROUGH DATA SILOS WITH AWS SERVICES INTEGRATION

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



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Today, Amazon Redshift allows you to customize capacity according to your needs

1. **Flexible instance types for granular control price performance**
2. **Manual and automatic workload management with query priorities**
3. **Vertical and horizontal scaling for performance and concurrency**
4. **Pause/resume clusters to save on costs during idle periods**

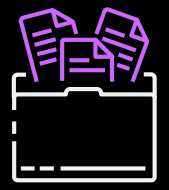
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Customers want to empower more users with Amazon Redshift



Users want to focus on insights, not managing the data warehouse

Access data without thinking about selecting instances, sizing, scaling, and tuning



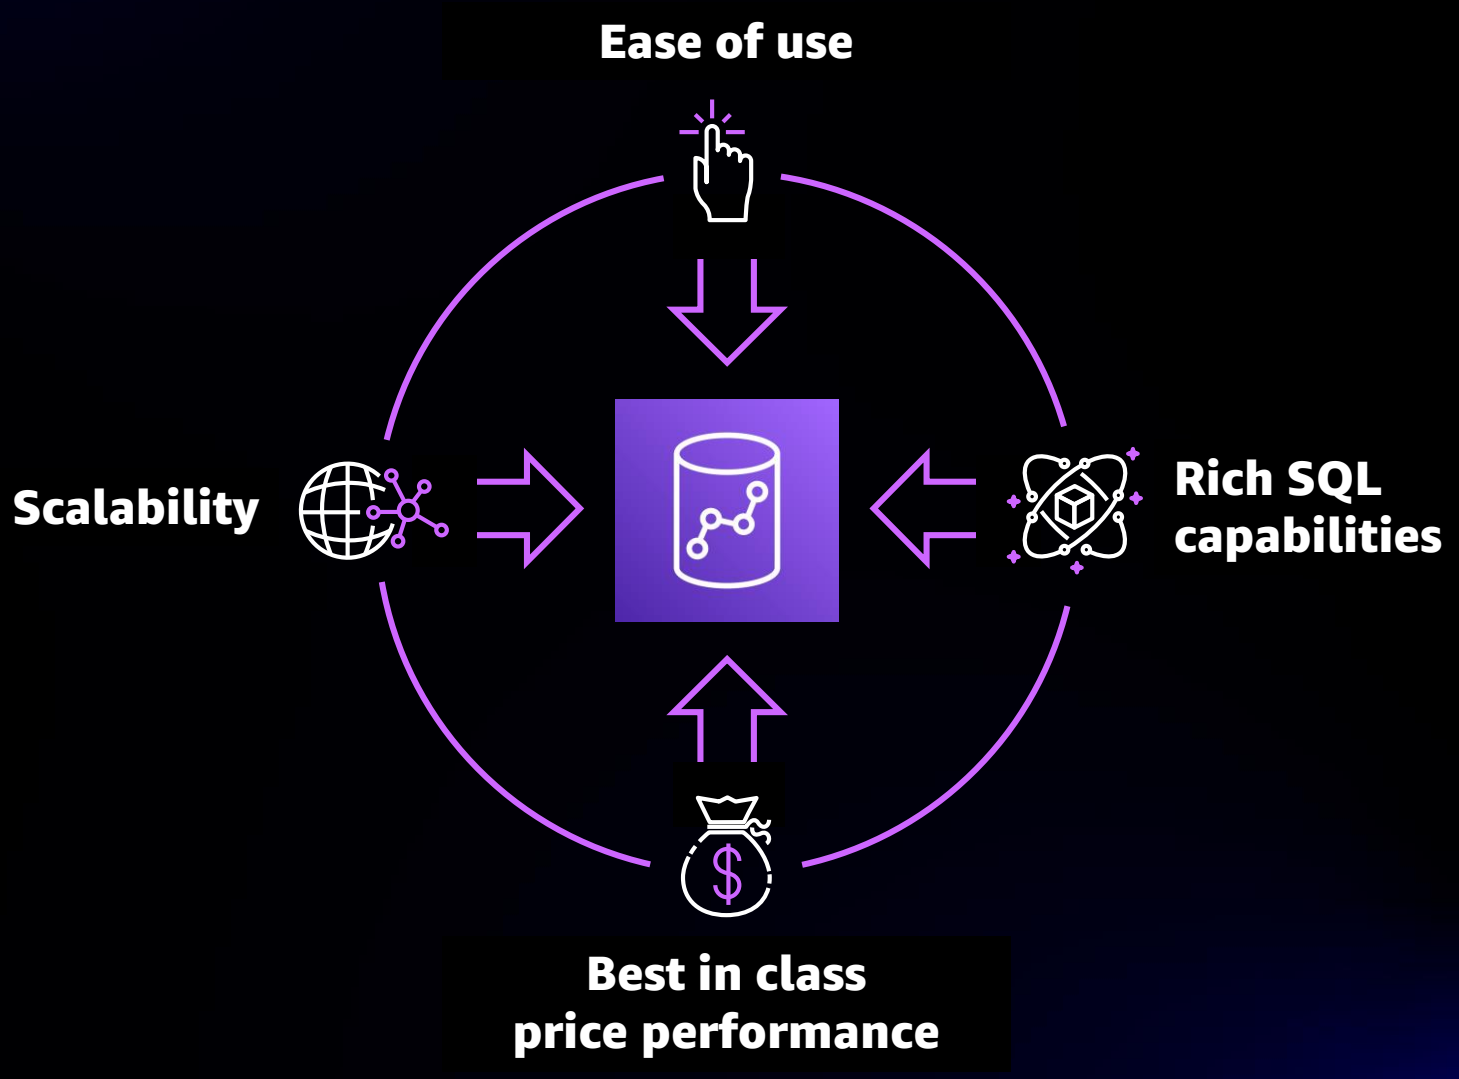
Unpredictable and varying workloads

Diverse use cases with heterogeneous compute requirements

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Need for analytics without compromises



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Easy analytics for everyone

Introducing Amazon Redshift Serverless

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift Serverless

Simplified user experience



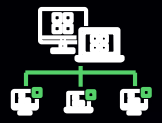
Run and scale analytics without having to manage data warehouse clusters

Intelligent and dynamic compute



Automatically provisions and scales data warehouse capacity to deliver consistently fast performance

All Redshift functionality and performance



Leverage Amazon Redshift's rich SQL capabilities, seamless data lake integration, as well as industry-leading price performance at scale

Pay for use



Pay for the compute capacity only for the workload duration on a per-second basis. No charges for idleness

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift Serverless key capabilities

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

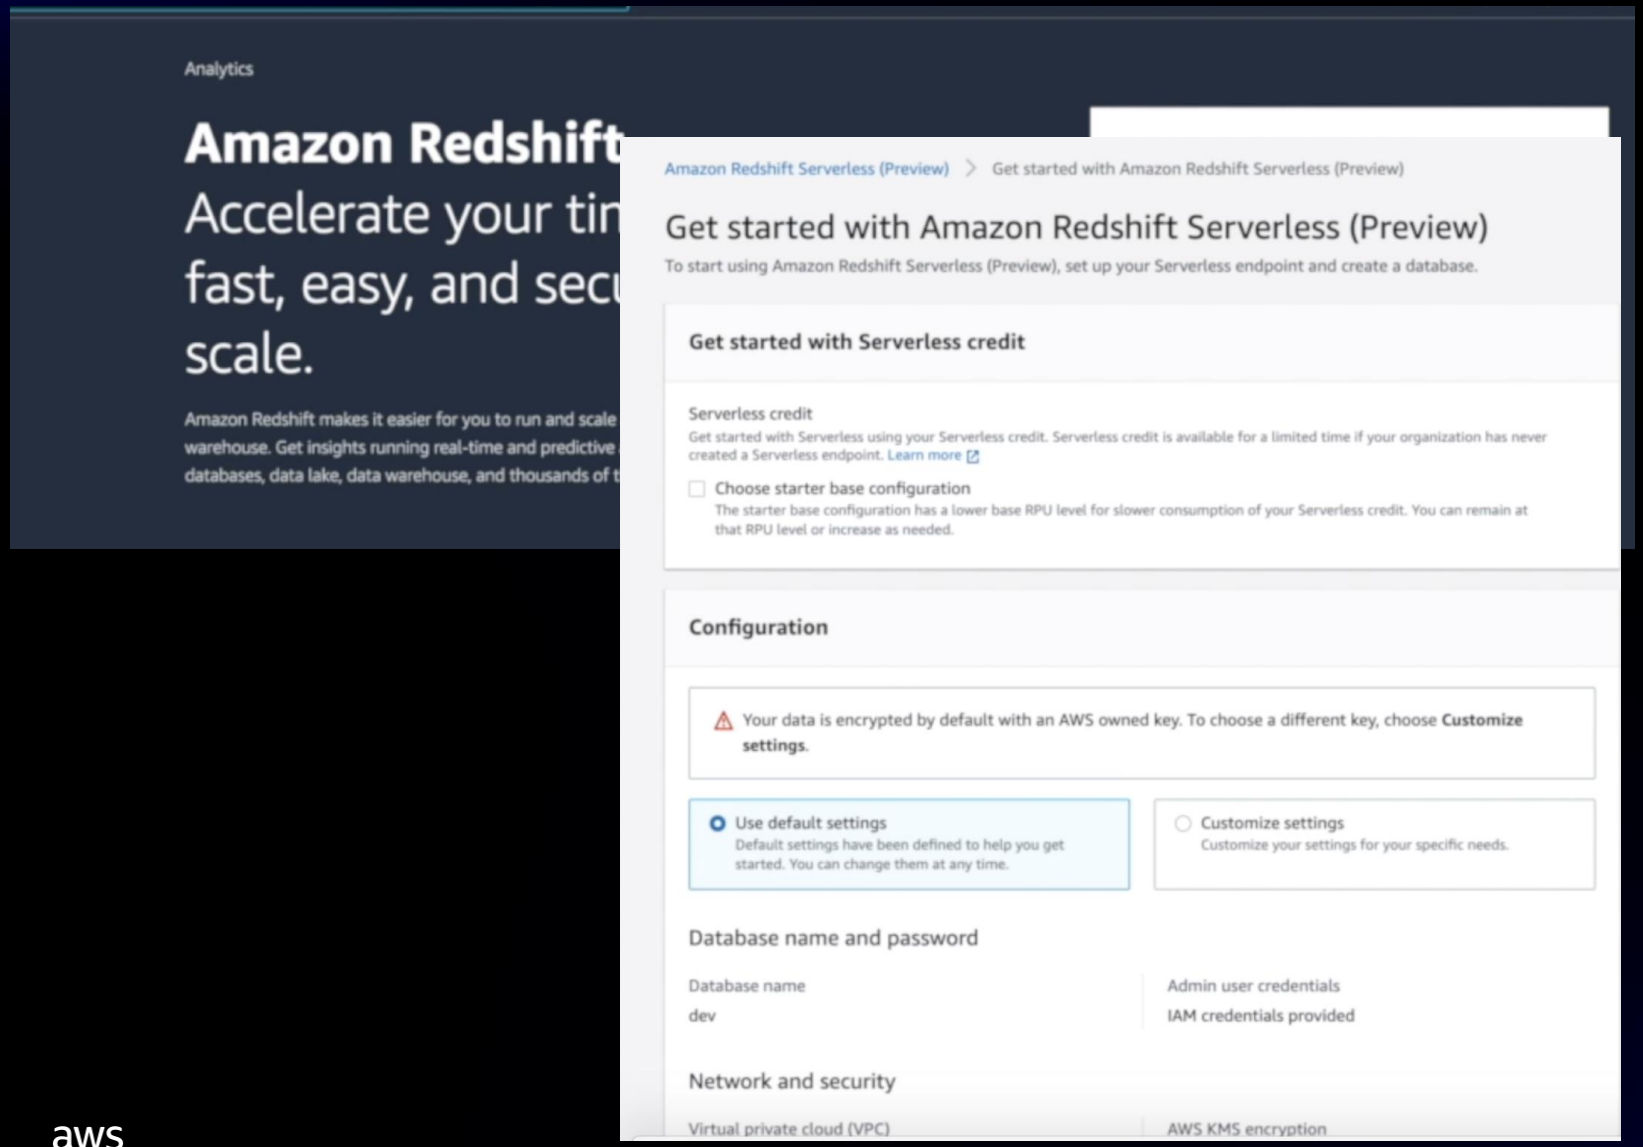
Getting started with serverless is easy



1. Enable Amazon Redshift serverless for your AWS account

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Getting started with serverless is easy

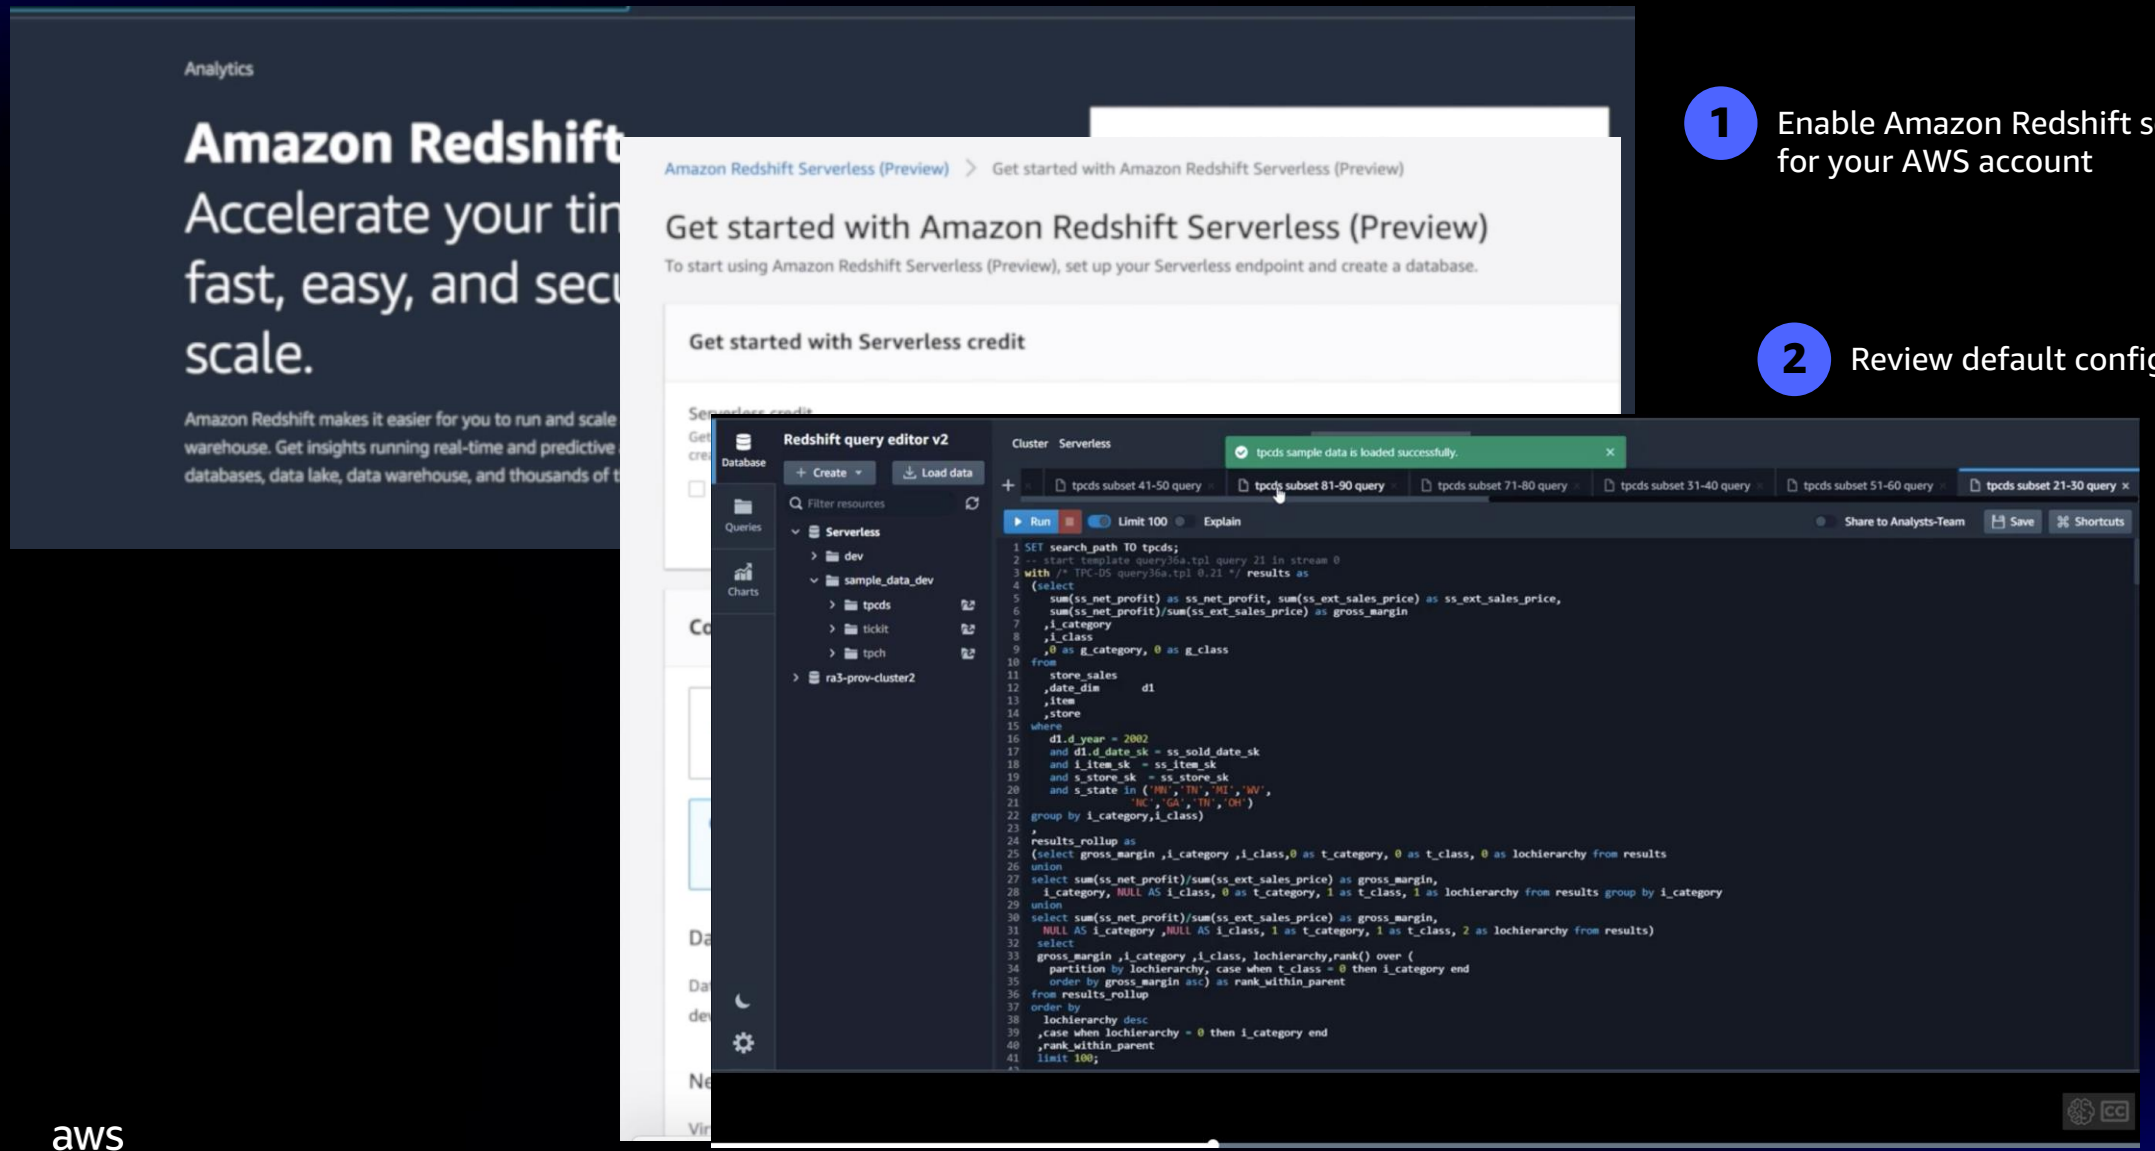


1. Enable Amazon Redshift serverless for your AWS account

2. Review default configuration

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Getting started with serverless is easy

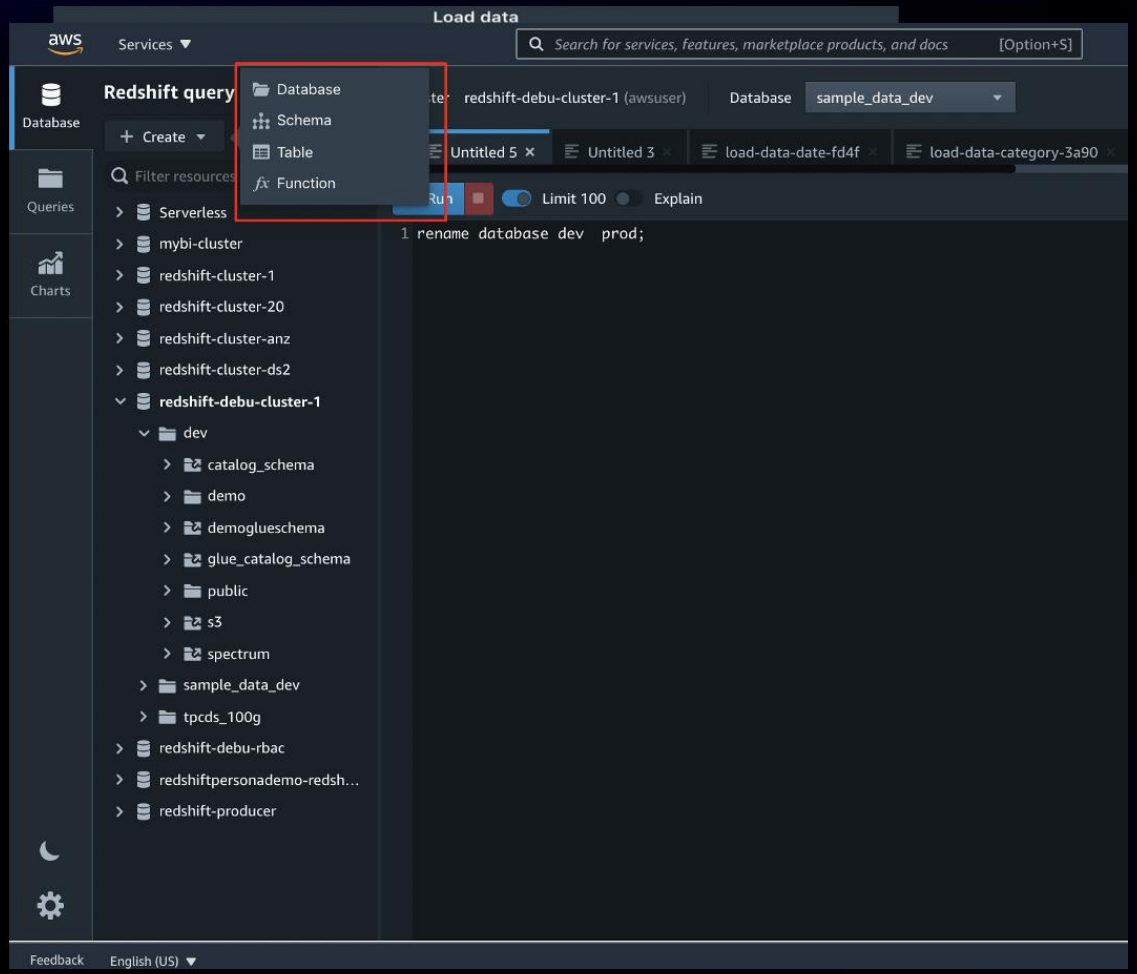


1. Enable Amazon Redshift serverless for your AWS account
2. Review default configuration
3. Connect from your favorite tools or Amazon Redshift Query Editor

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Leverage all the data for analysis

EASILY LOAD DATA INTO THE DATA WAREHOUSE

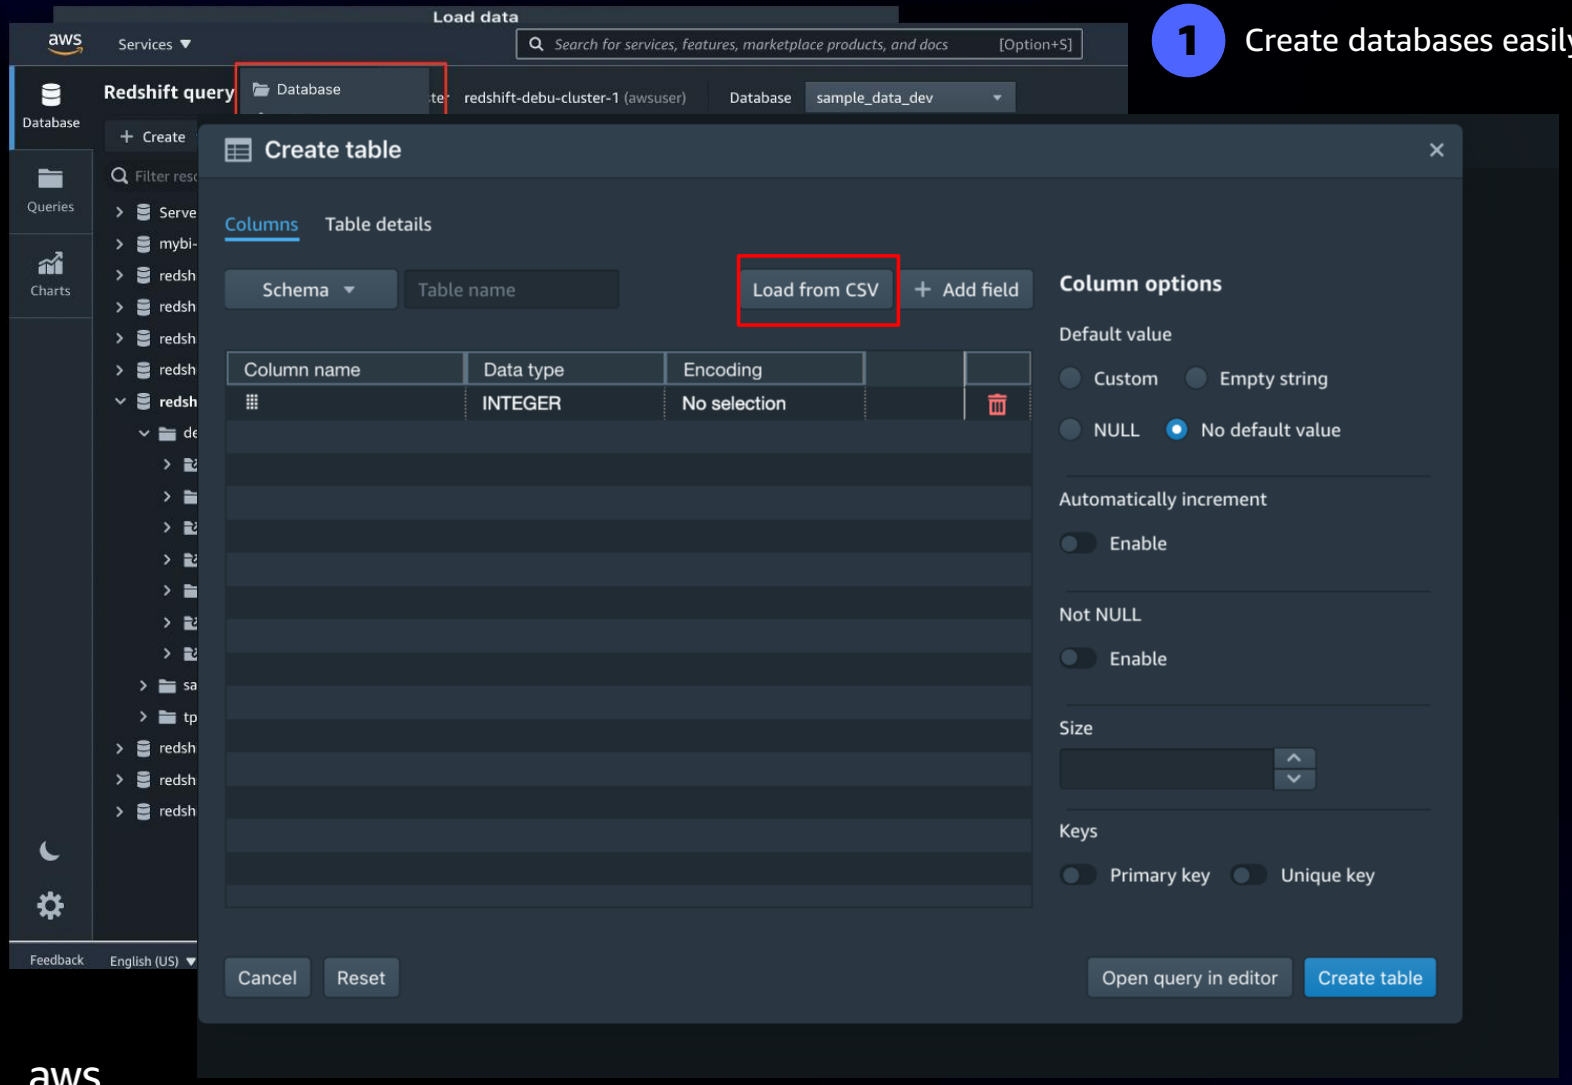


1. Create databases easily by loading data from existing sources

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Leverage all the data for analysis

EASILY LOAD DATA INTO THE DATA WAREHOUSE



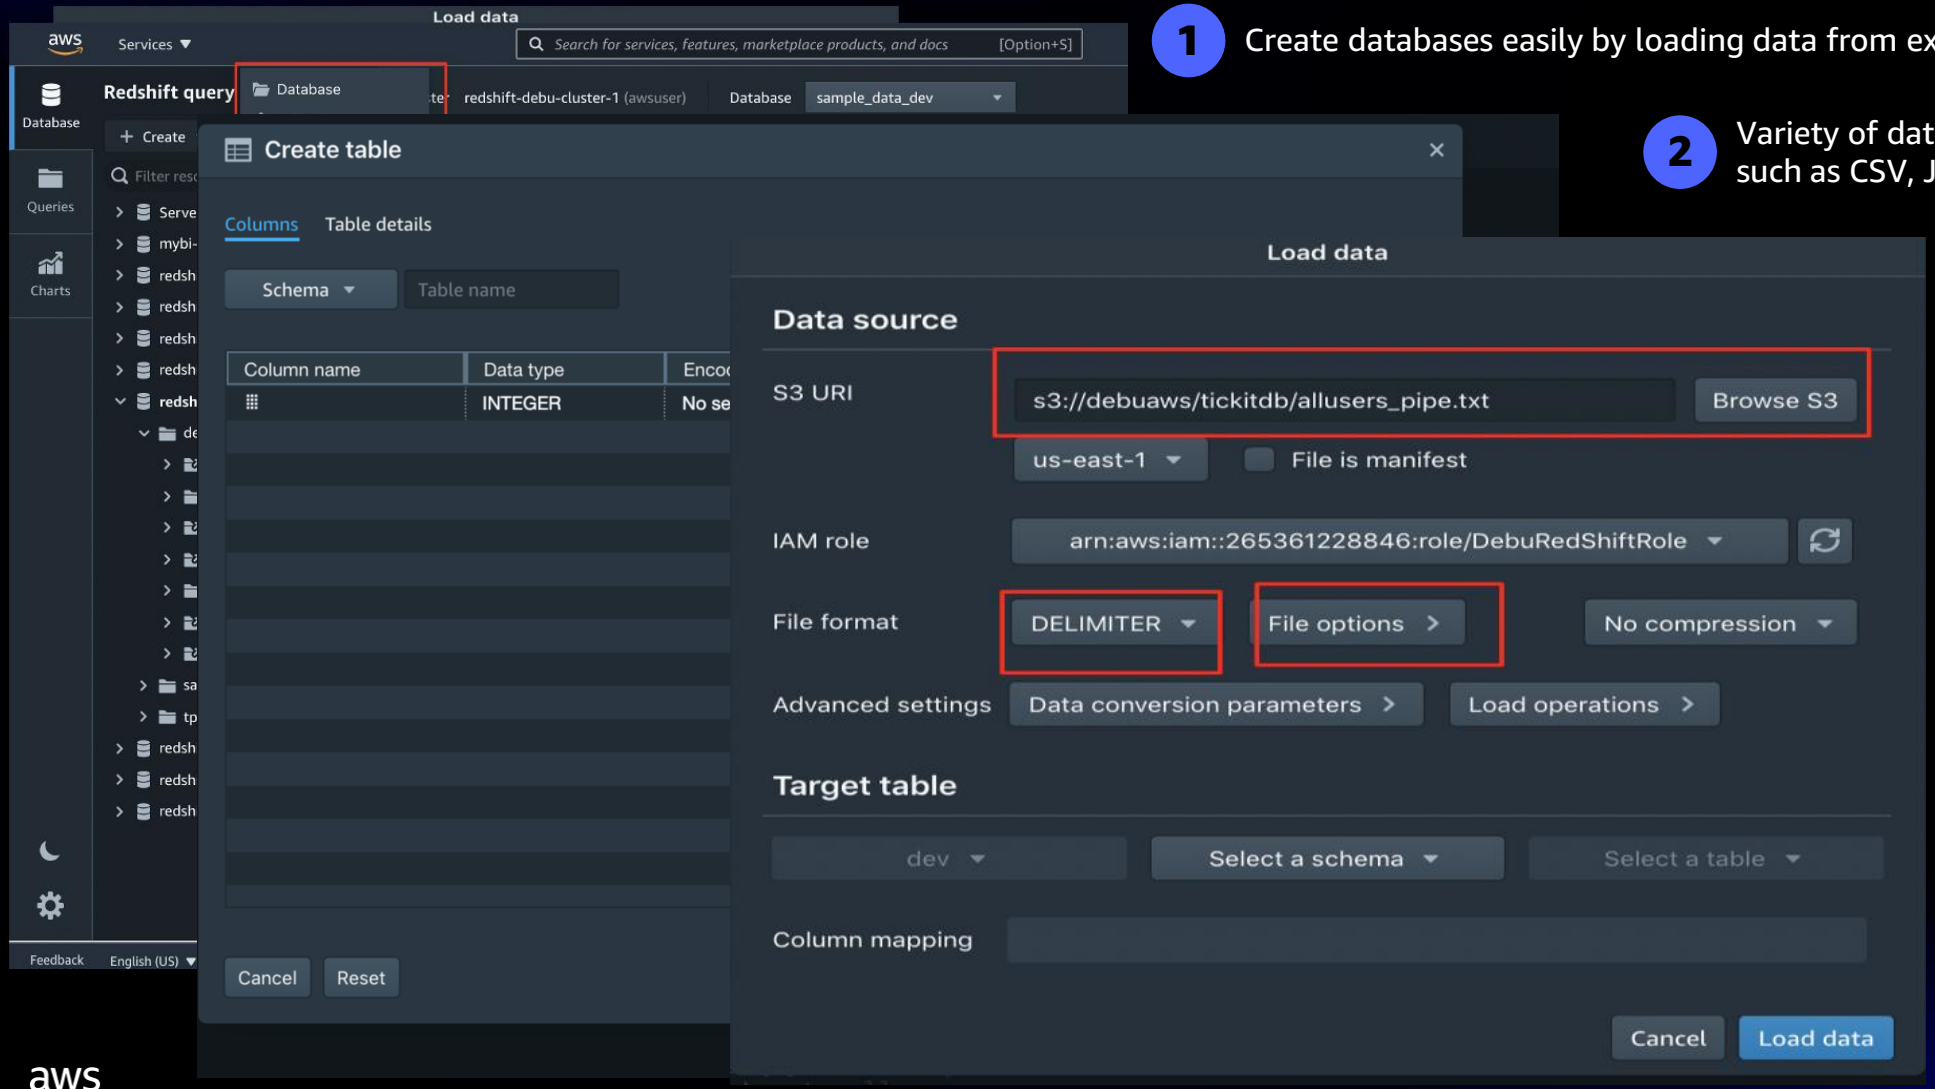
1. Create databases easily by loading data from existing sources

2. Variety of data formats support such as CSV, JSON

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Leverage all the data for analysis

EASILY LOAD DATA INTO THE DATA WAREHOUSE

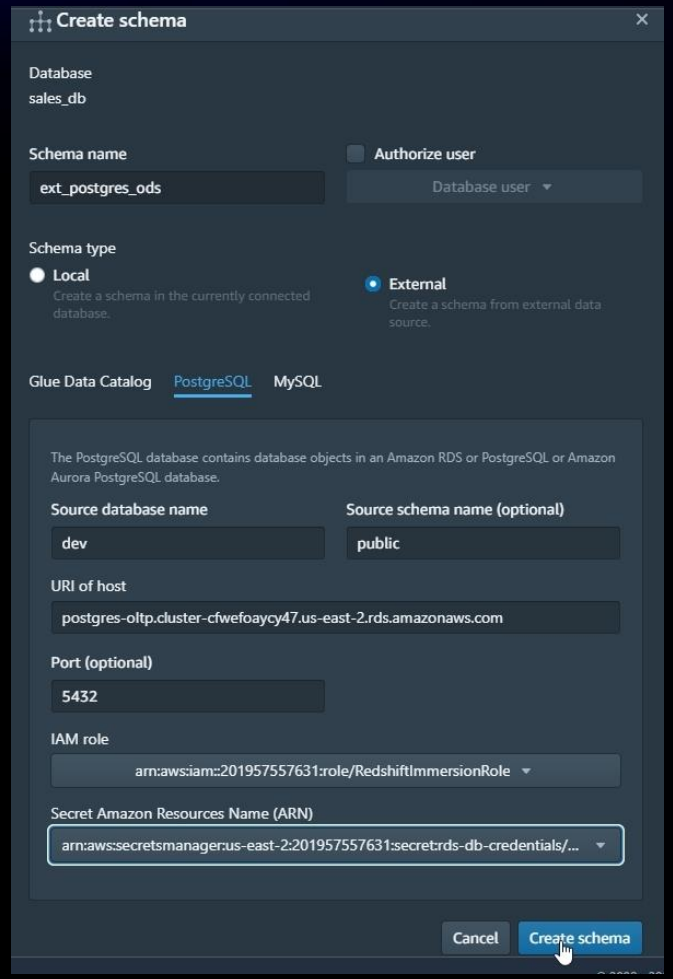


1. Create databases easily by loading data from existing sources
2. Variety of data formats support such as CSV, JSON
3. Browse and load data from S3 directly

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Leverage all the data for analysis

BROWSE AND QUERY DATA IN AMAZON S3 DATA LAKE AND OPERATIONAL DATABASES



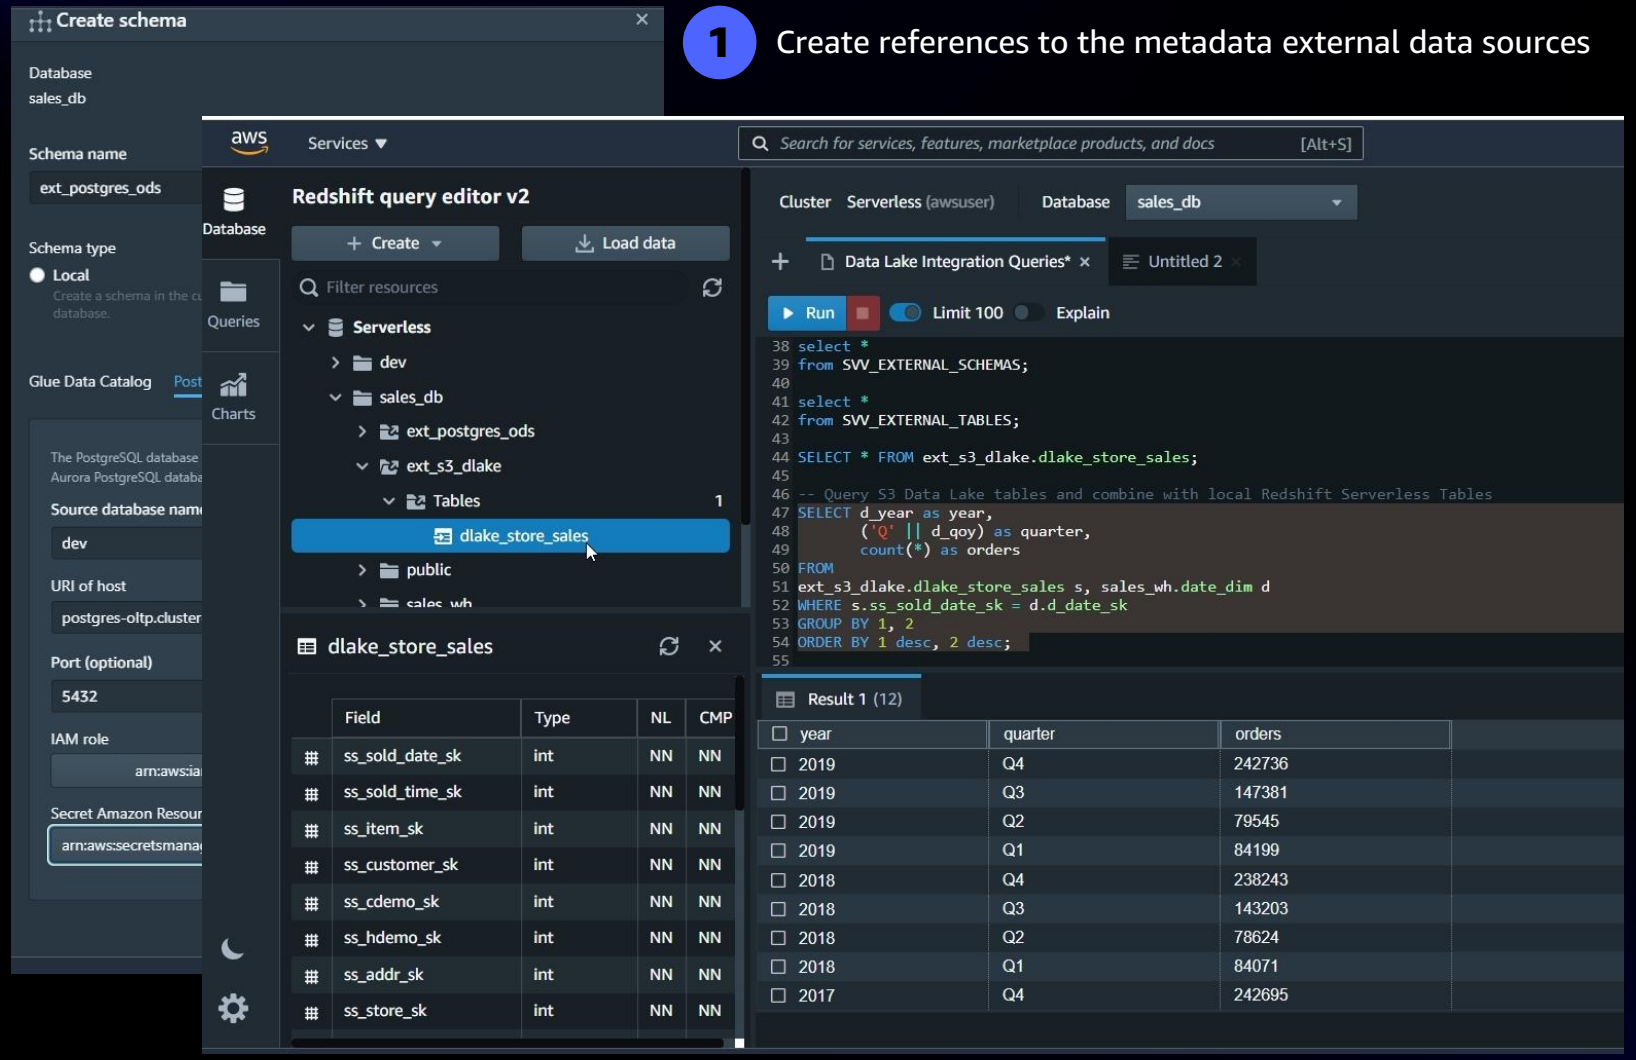
1. Create references to the metadata external data sources

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Leverage all the data for analysis

BROWSE AND QUERY DATA IN AMAZON S3 DATA LAKE AND OPERATIONAL DATABASES

1 Create references to the metadata external data sources



2 Browse and query data in Amazon S3 data lake via AWS Glue catalog

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

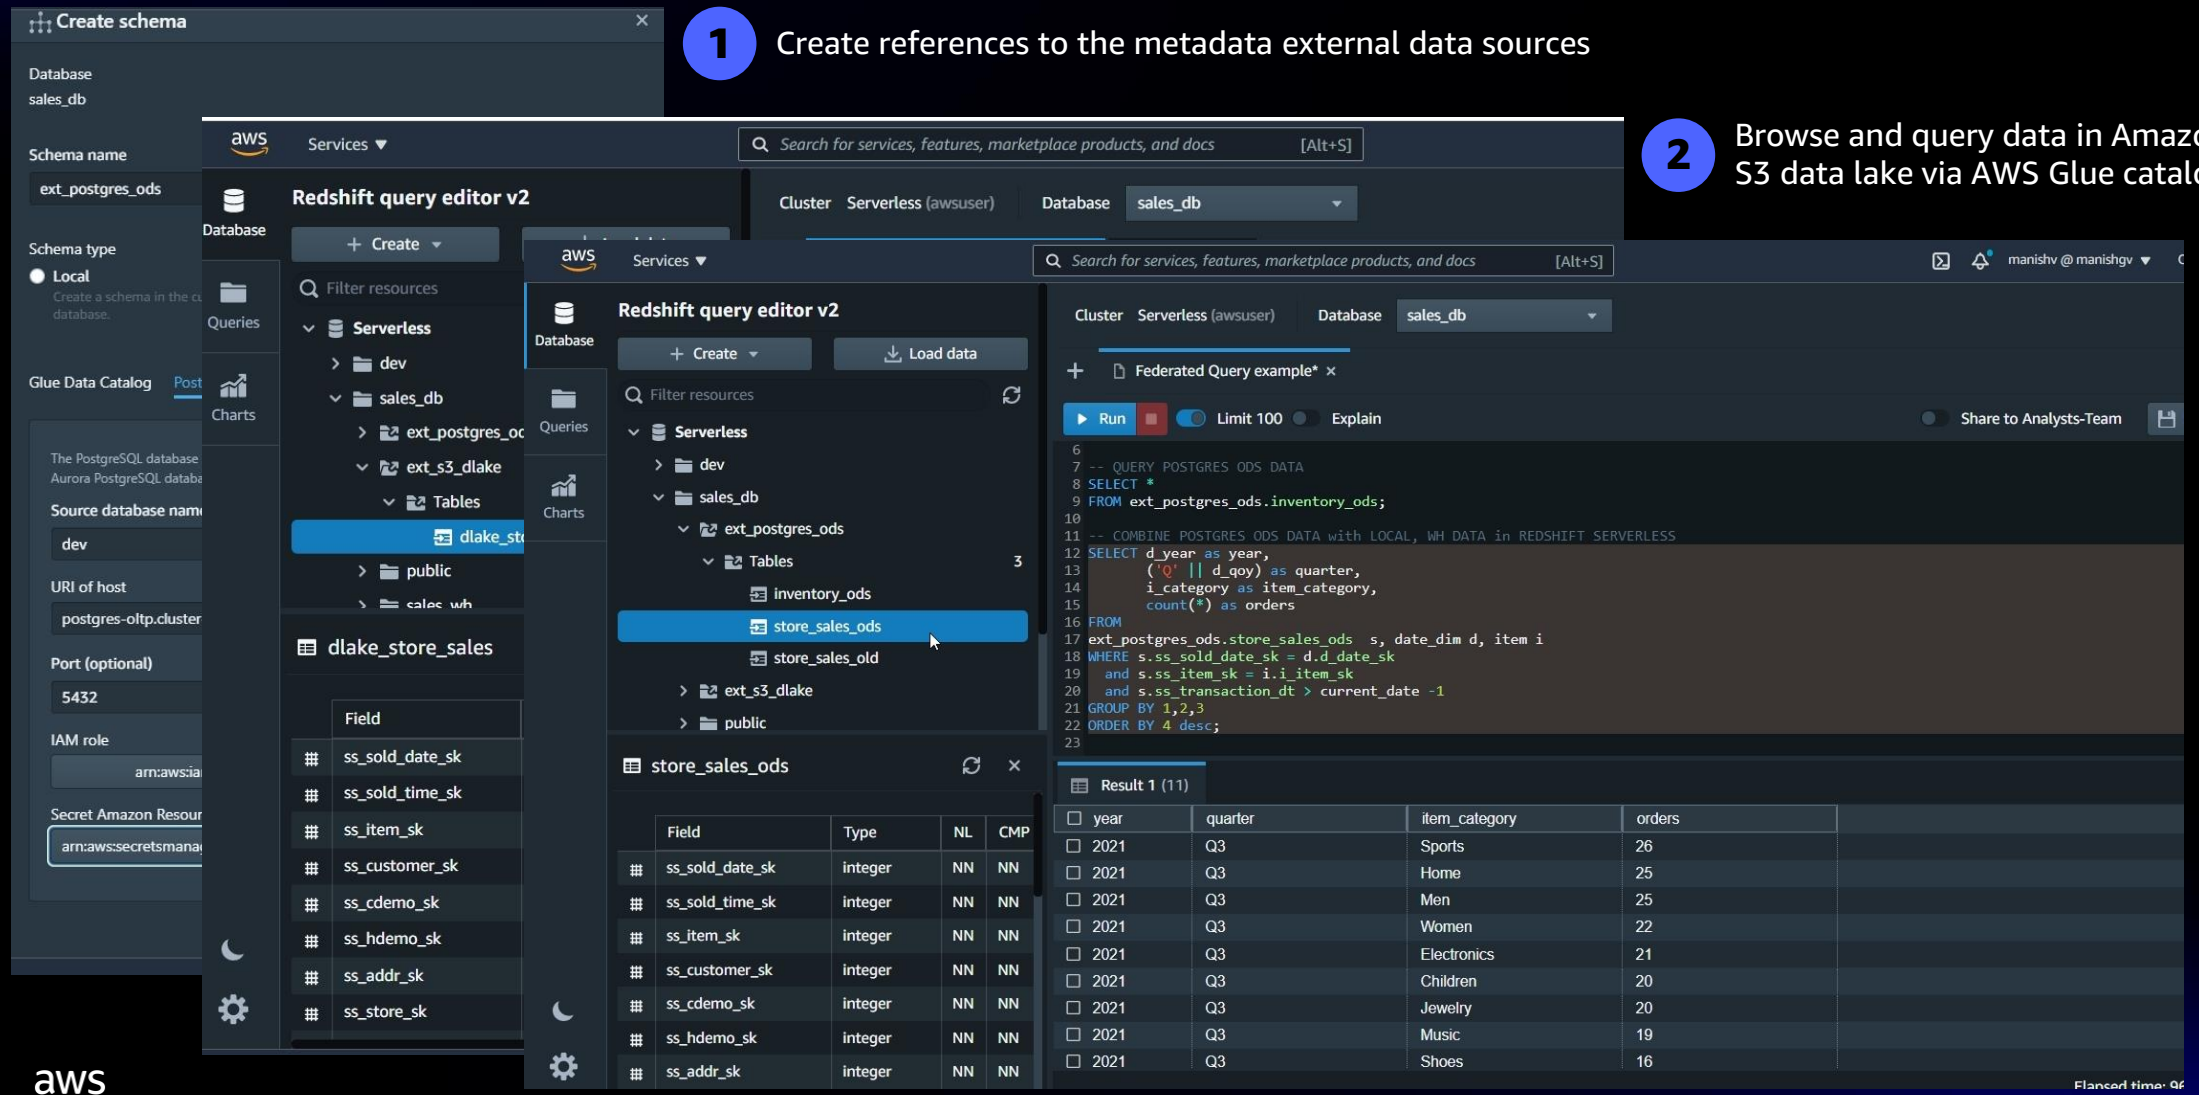
Leverage all the data for analysis

BROWSE AND QUERY DATA IN AMAZON S3 DATA LAKE AND OPERATIONAL DATABASES

1 Create references to the metadata external data sources

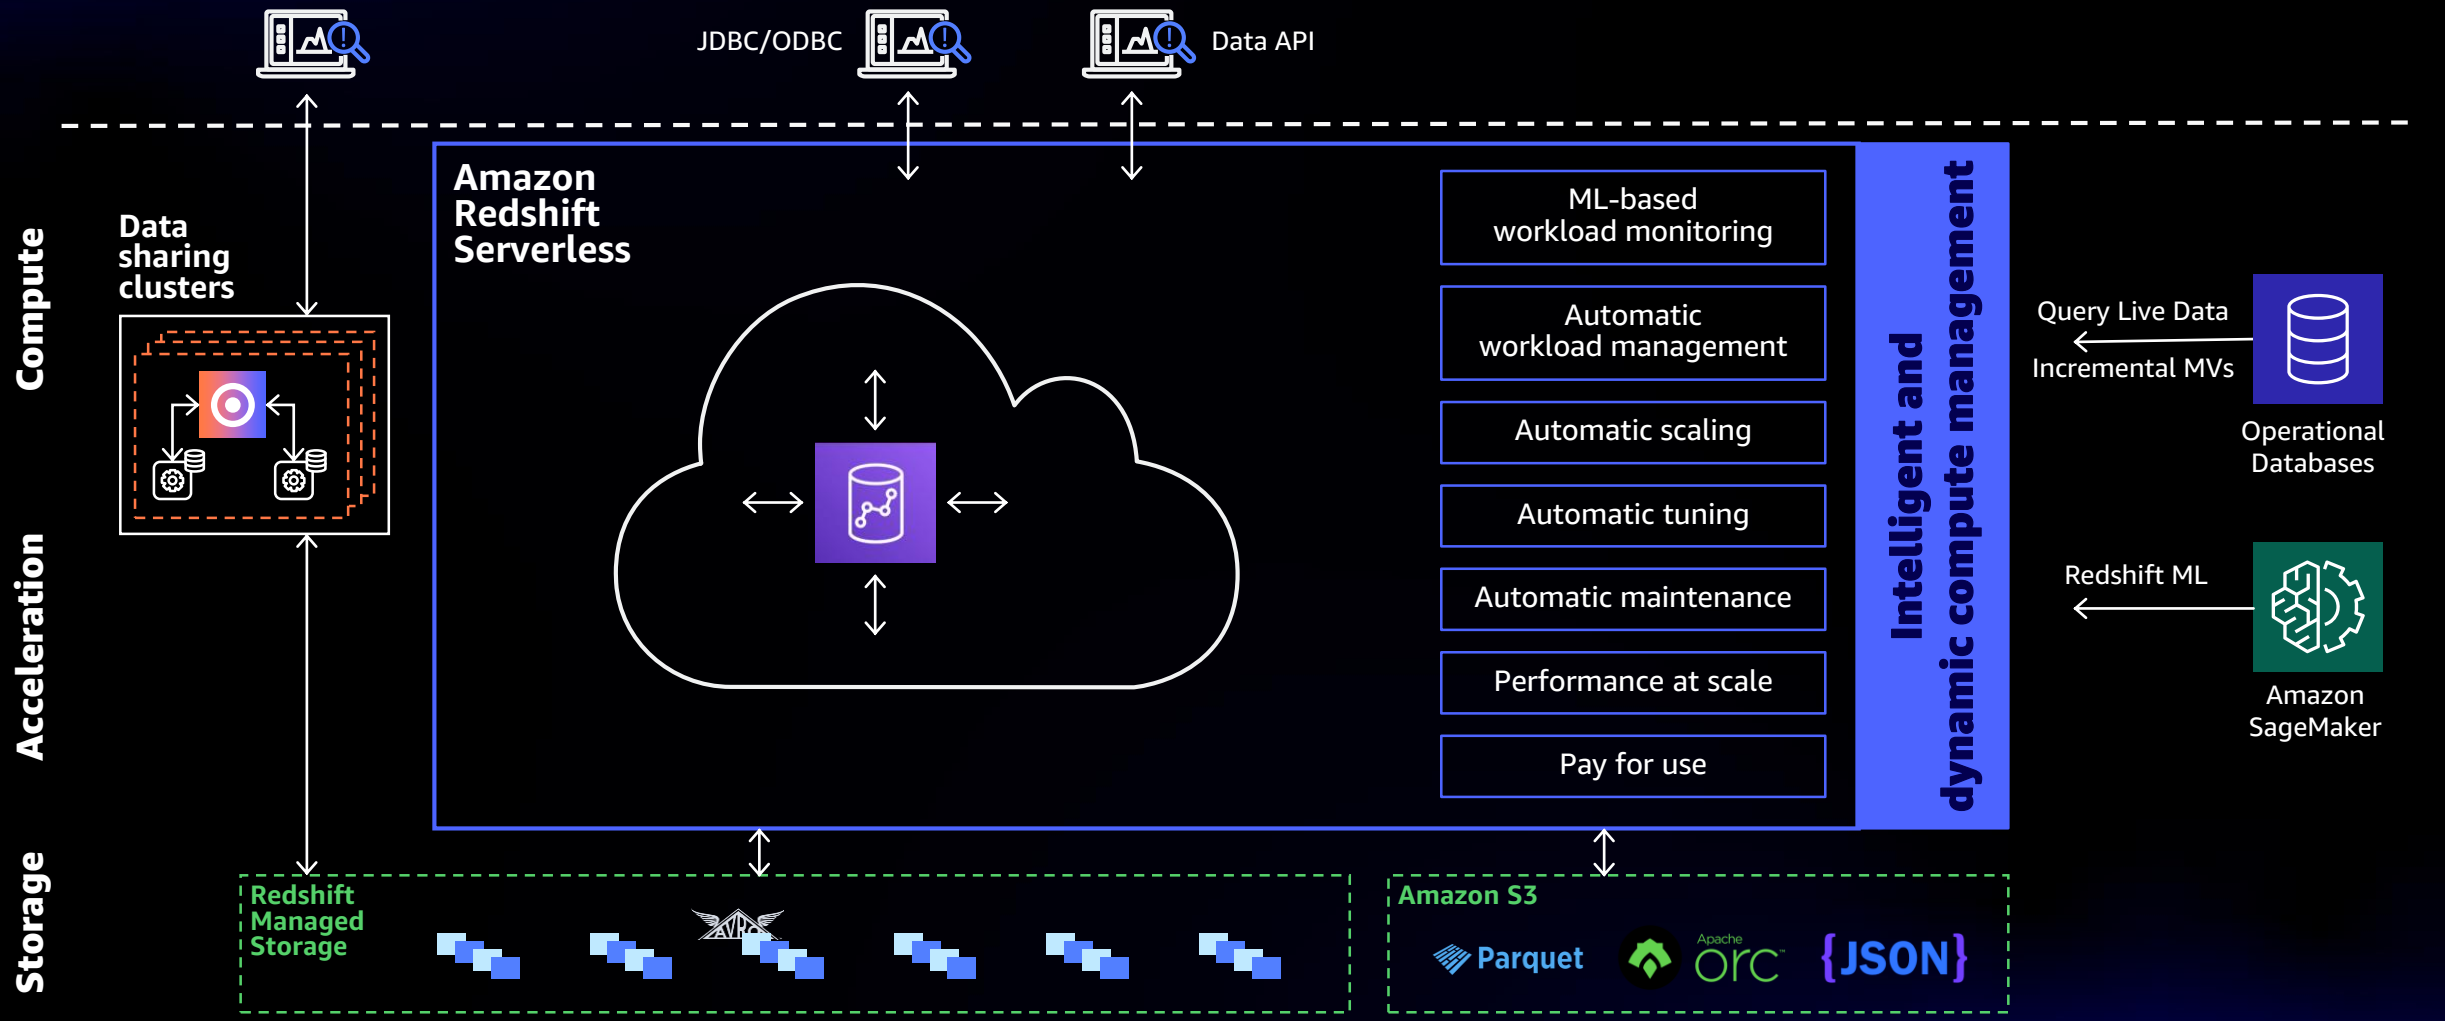
2 Browse and query data in Amazon S3 data lake via AWS Glue catalog

3 Browse and query data in Amazon Aurora/ Amazon RDS databases



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Intelligent and dynamic compute management



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Unified billing

Compute



- New normalized compute unit – Redshift Processing Unit (RPU)
- Usage is billed in RPU-hours, metered on per-second basis
- Base data warehouse, scaling capacity, data lake queries are part of same RPU-hours
- Concurrency scaling and spectrum charges do not apply

Storage



- Fixed GB-month rate pricing for the Redshift managed storage and user snapshots
- Restore their data warehouse to specific points in last 24 hours at a 30 minute granularity at free of charge

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Pay for use

Pay for the compute capacity **only** for the workload duration (metered on a per-second basis)

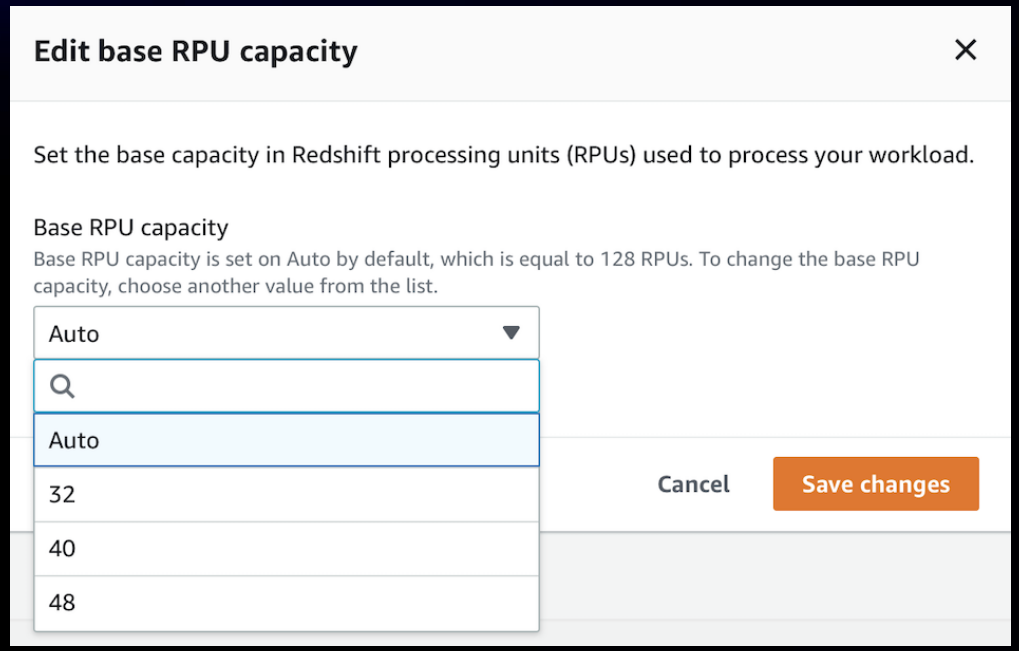


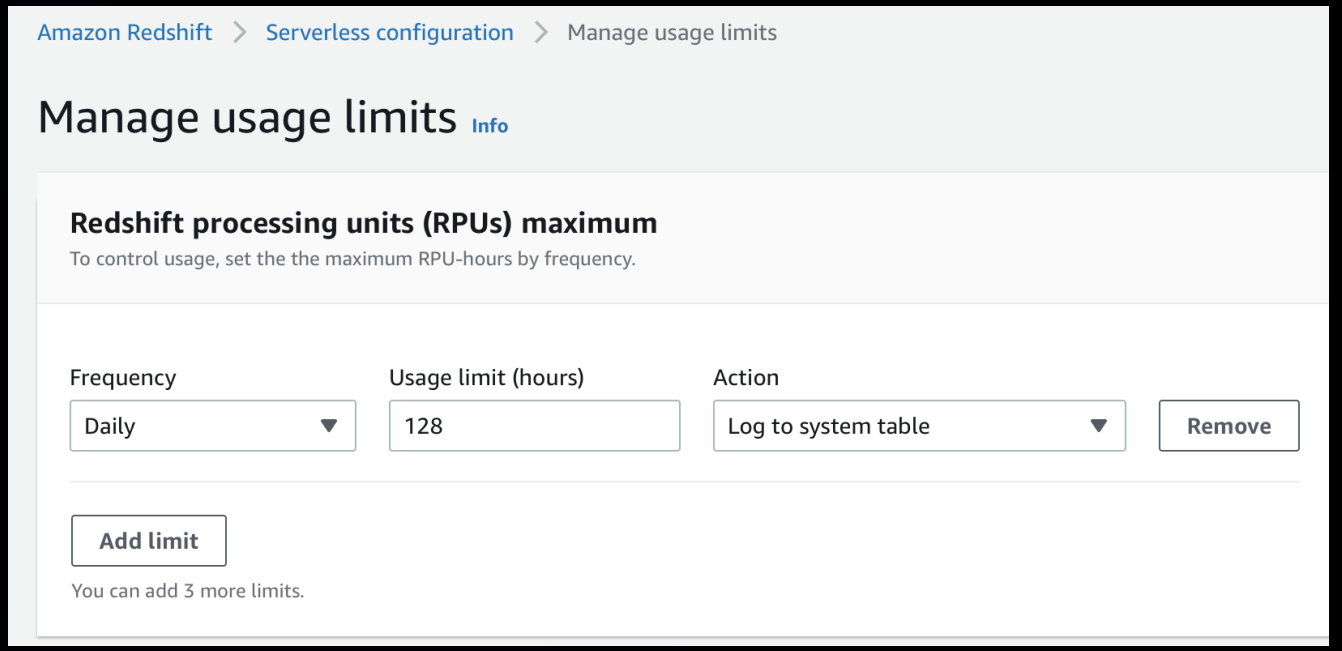
Billing duration	Query execution time
@2:03	3 minutes (for Q1, Q2, Q3)
@2:09	1 minute 10 seconds (for Q4)
@2:14	1 minute and 20 seconds (for Q5)
Total charges	5 minutes and 30 seconds

No charges for idleness

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Performance and cost controls



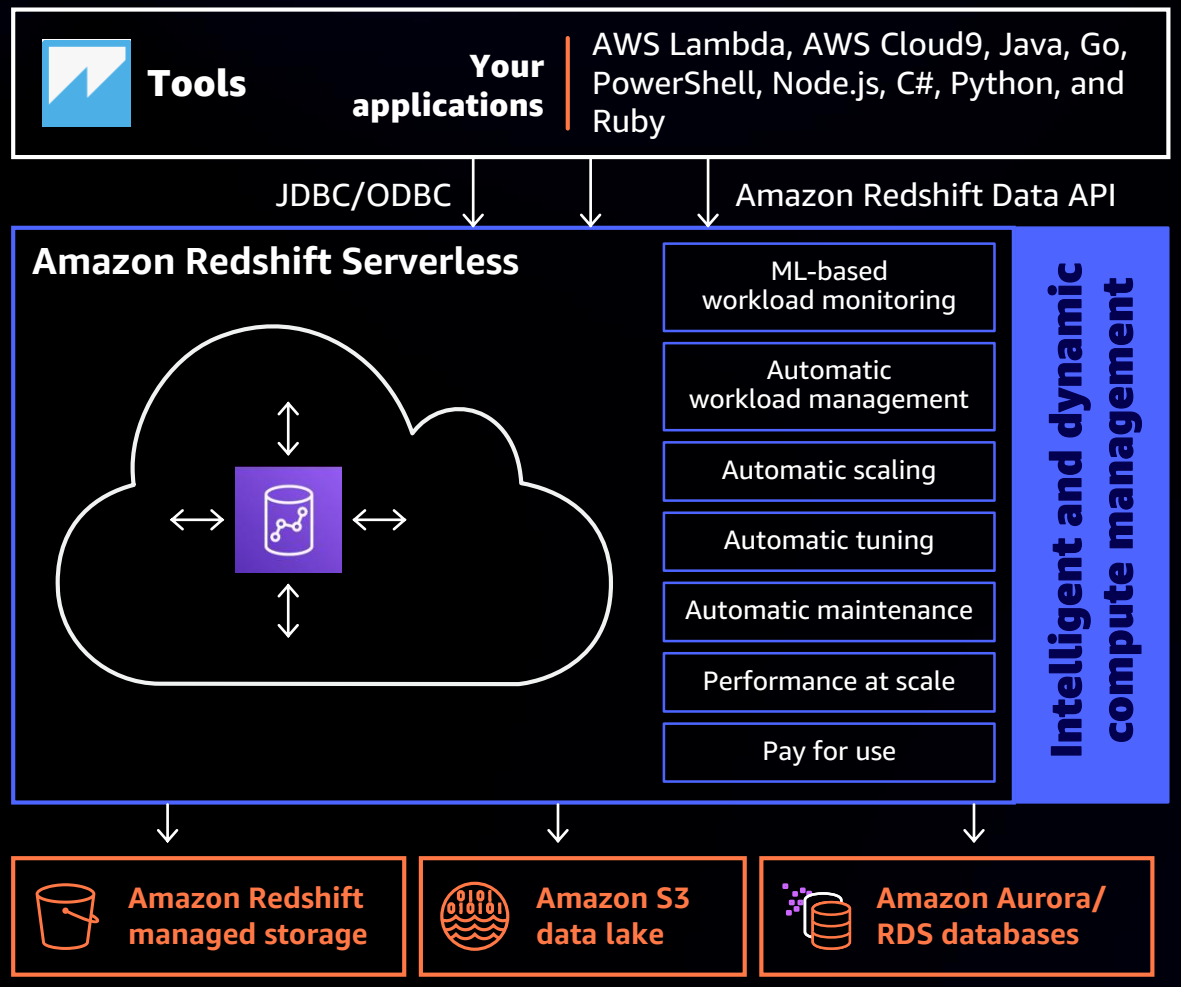


- Adjust base data warehouse compute based on workload performance requirements
- System applies workload management to ensure maximum throughput from the base data warehouse compute

- Set daily/weekly/monthly usage limit RPU-hours to put thresholds on costs
- System automatically scales within limits to handle additional users and query concurrency while providing consistent performance

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

No changes necessary for existing analytics and business intelligence applications



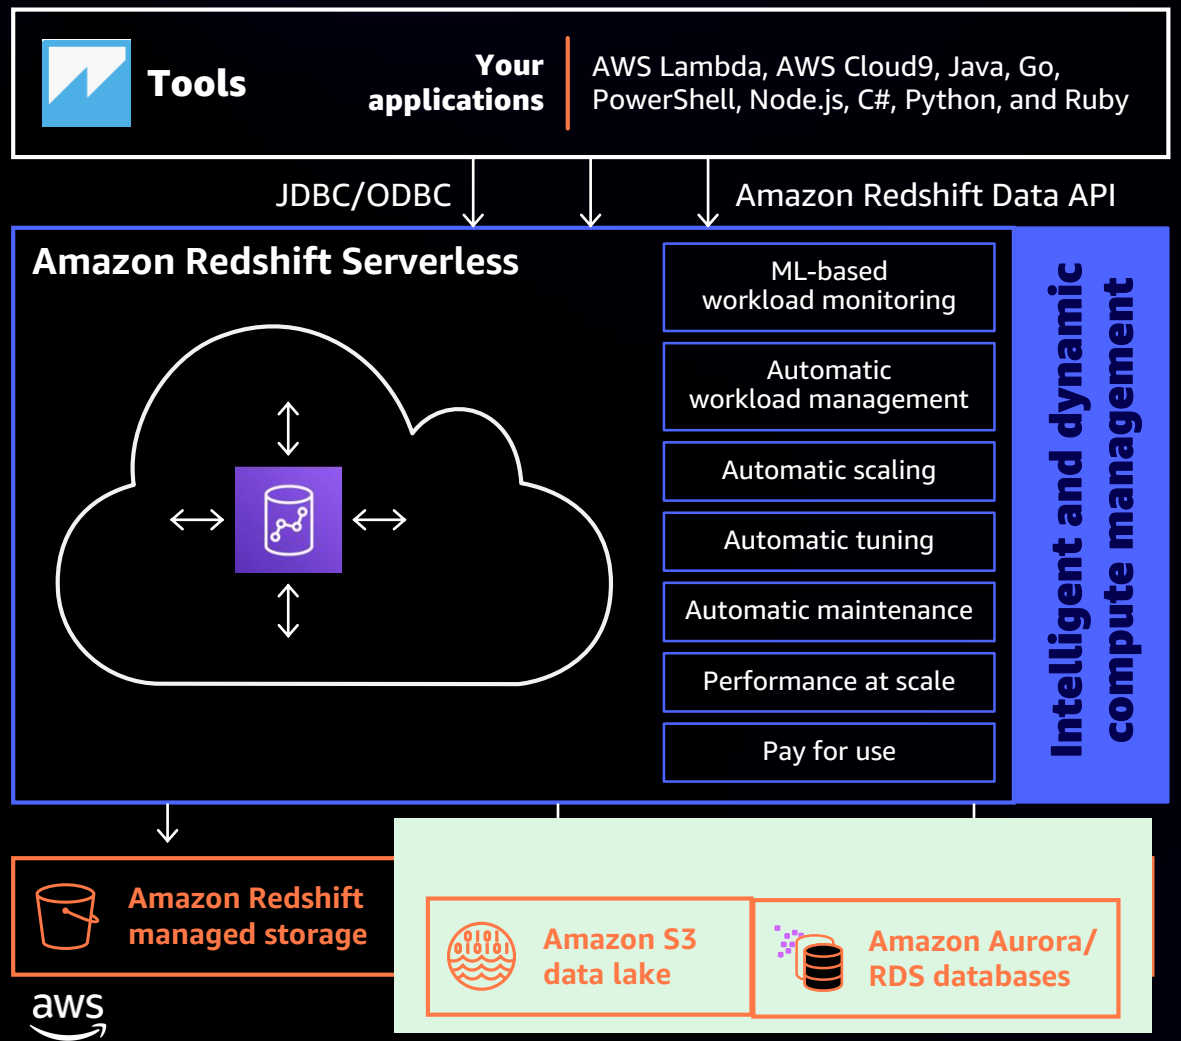
Simply point applications to the Amazon Redshift serverless end point and start running

All Redshift SQL functionality applies

- Security and user management
- Complex joins
- Semi-structured data
- Data sharing
- Machine learning functions
- Data lake queries
- Federated query
- Durability and transactional guarantees
- JDBC/ODBC and Data API
- And more

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Redshift Serverless: Query your data lake



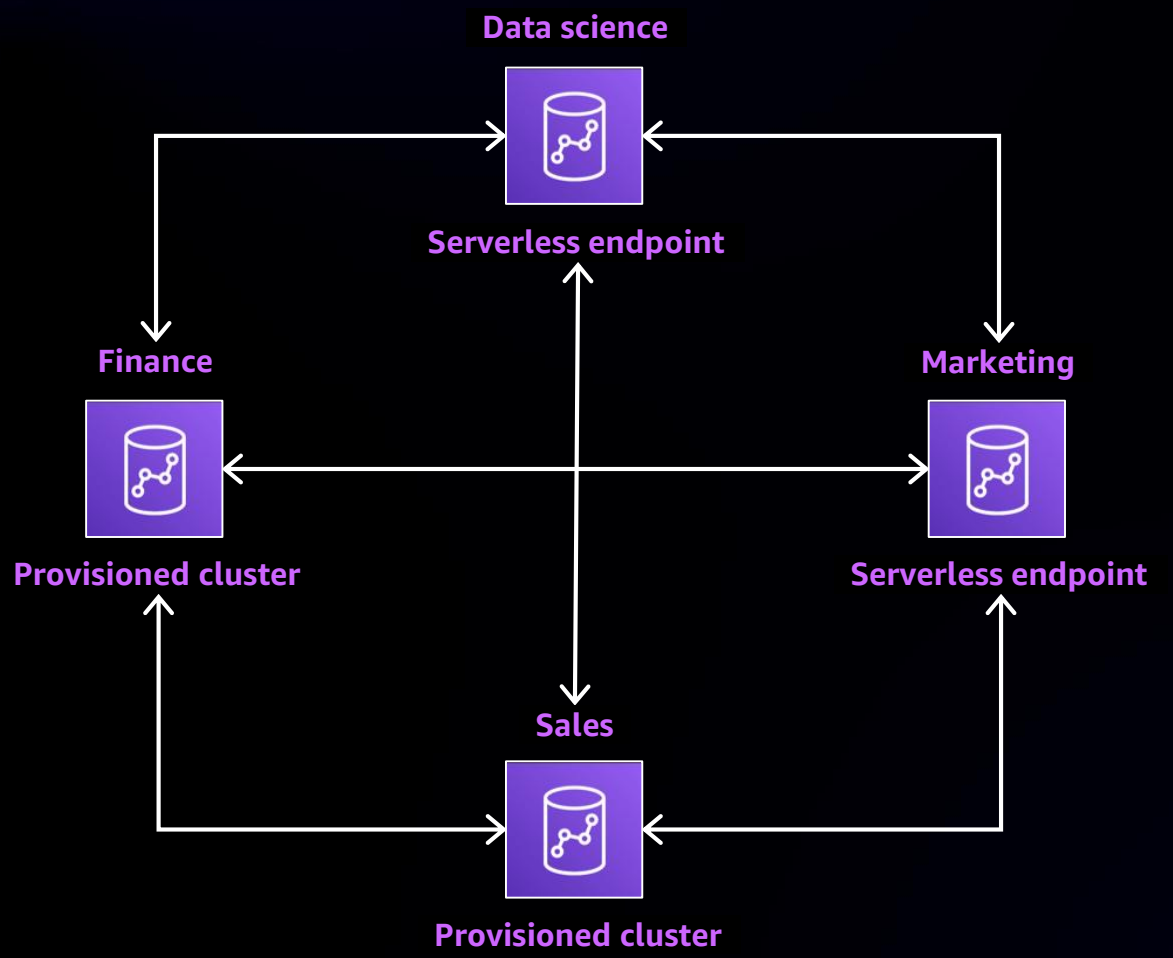
Query external data seamlessly from serverless end point

Billed unified into RPU-hours - No per TB price for scanning data lake data

Redshift Serverless routes data lake query through your VPC

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data sharing in Amazon Redshift Serverless



Share data seamlessly between provisioned cluster to serverless end point and vice versa

Share data across serverless end points across different AWS accounts

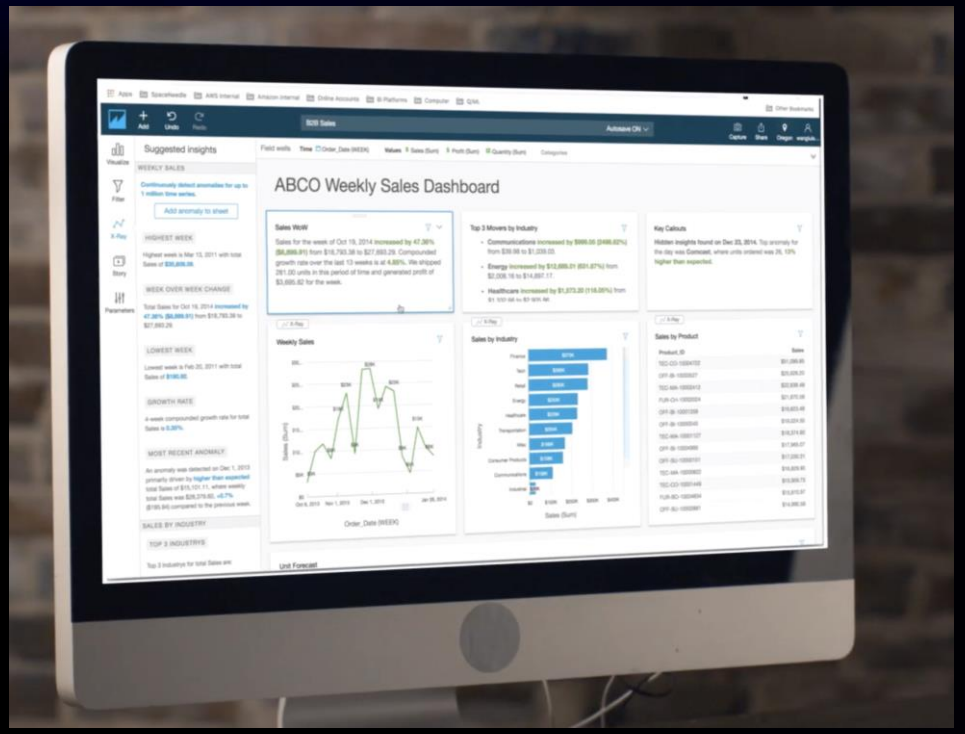
Instant, live data access within and across AWS accounts and across AWS regions

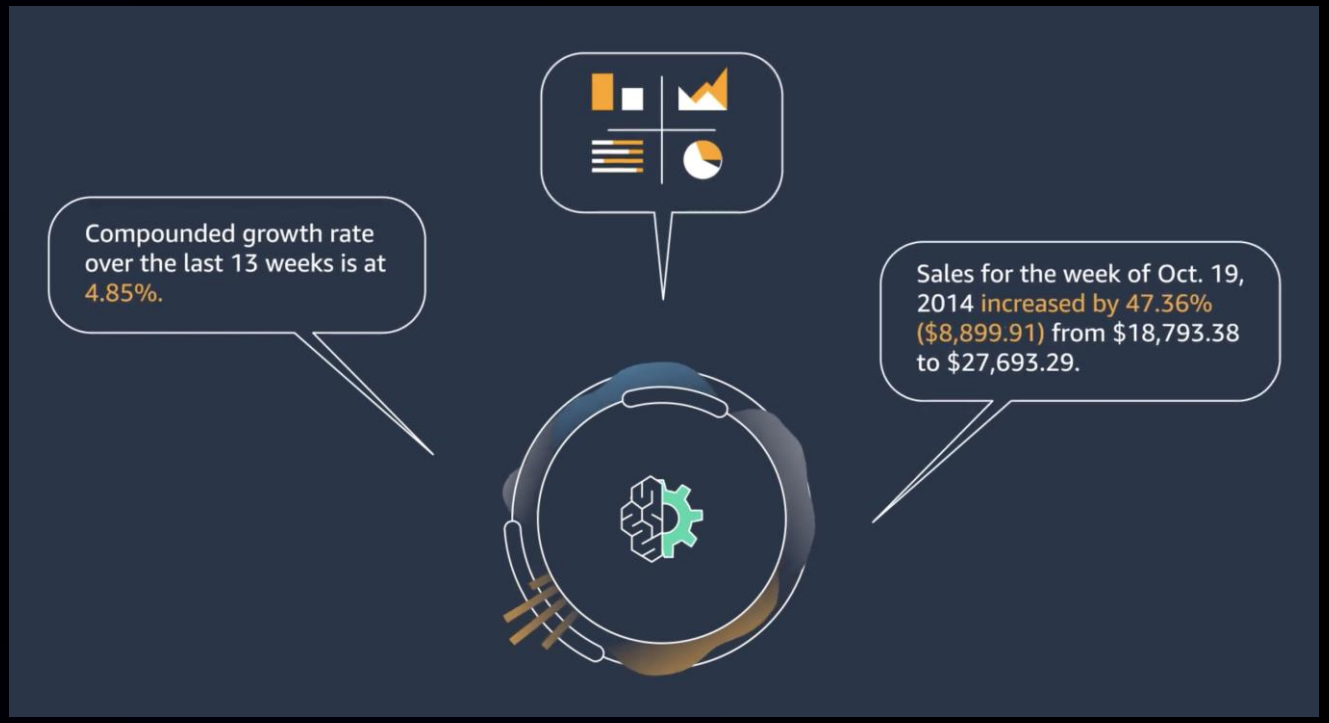
© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift Serverless use cases

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Serverless use cases: Easy analytics



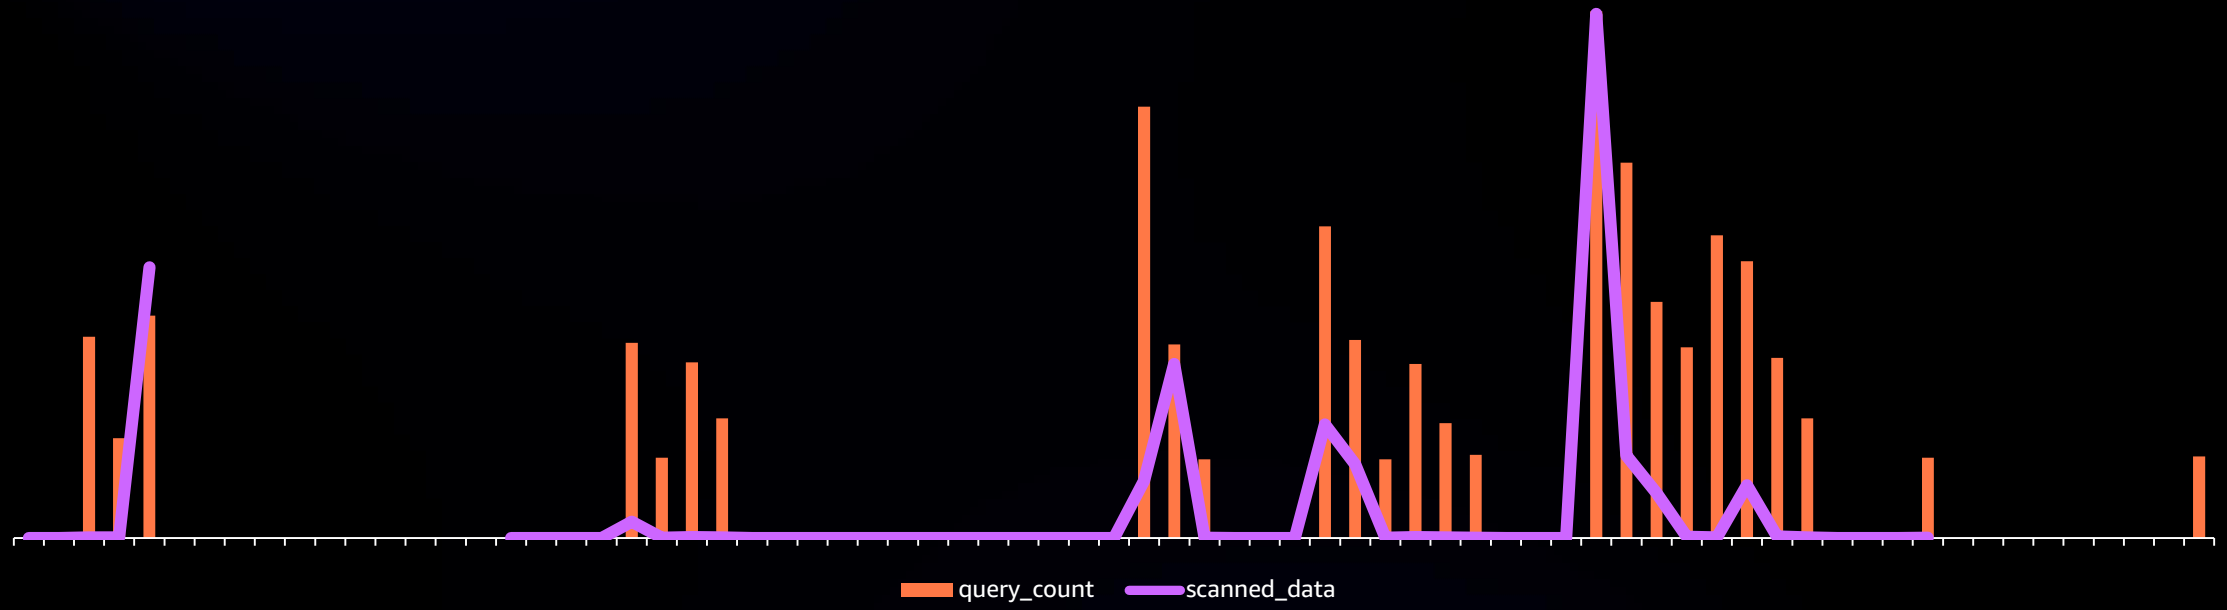


Why

- Get started easily without thinking about infrastructure
- No need to create and manage clusters by selecting node types and node count
- Example use cases—dev/test environments and ad-hoc business analytics
- Only pay for use

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Serverless use cases: Variable workloads

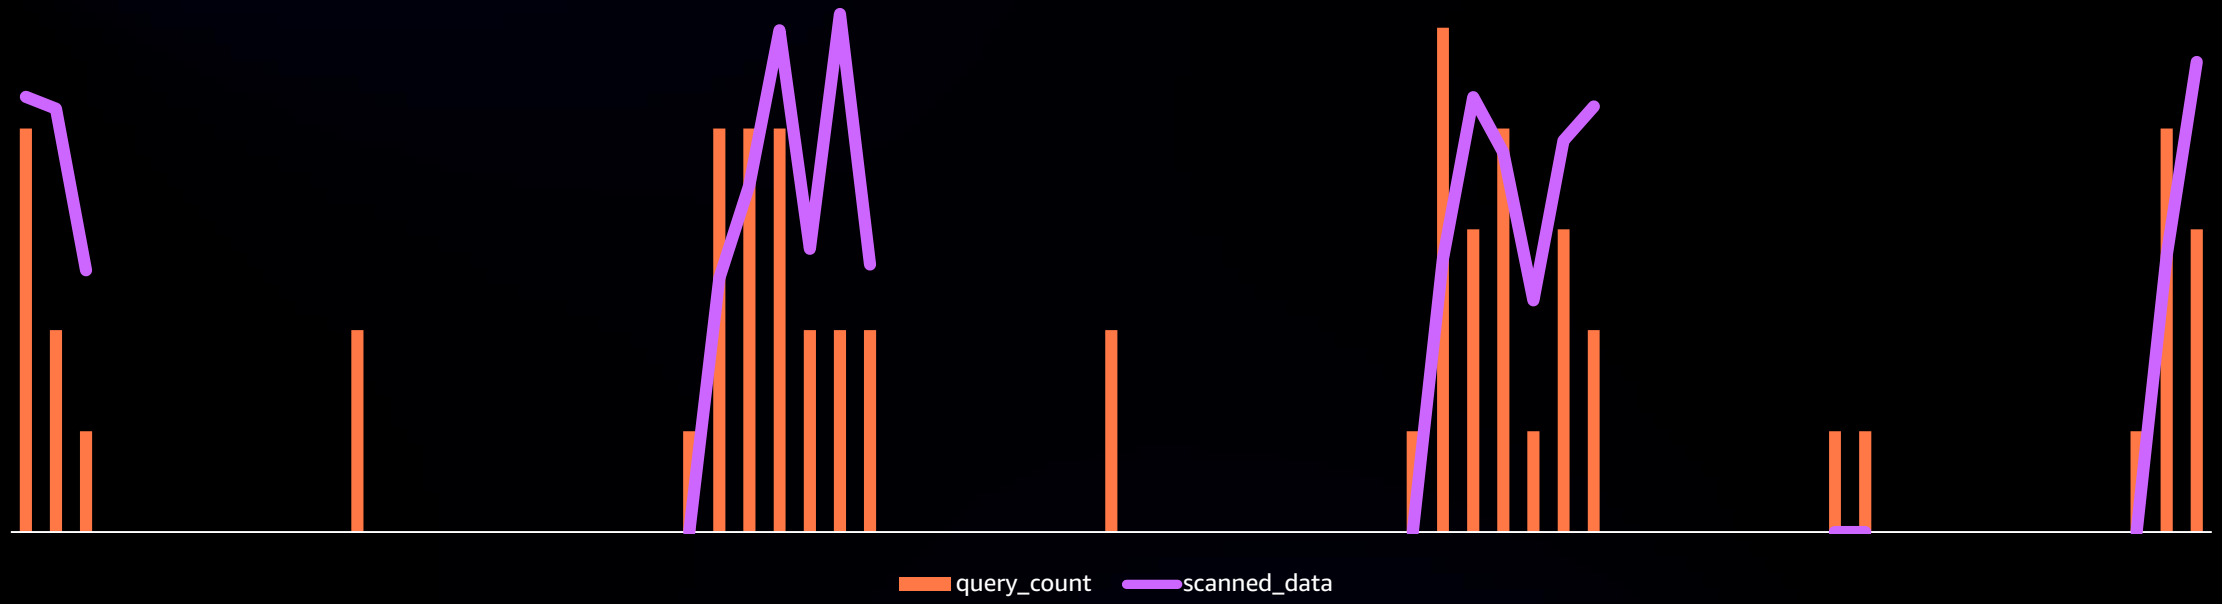


Why

- Scale up and down according to the workload variability to meet the unknown and unpredictable demands
- Only pay for use

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Serverless use cases: Periodic workloads

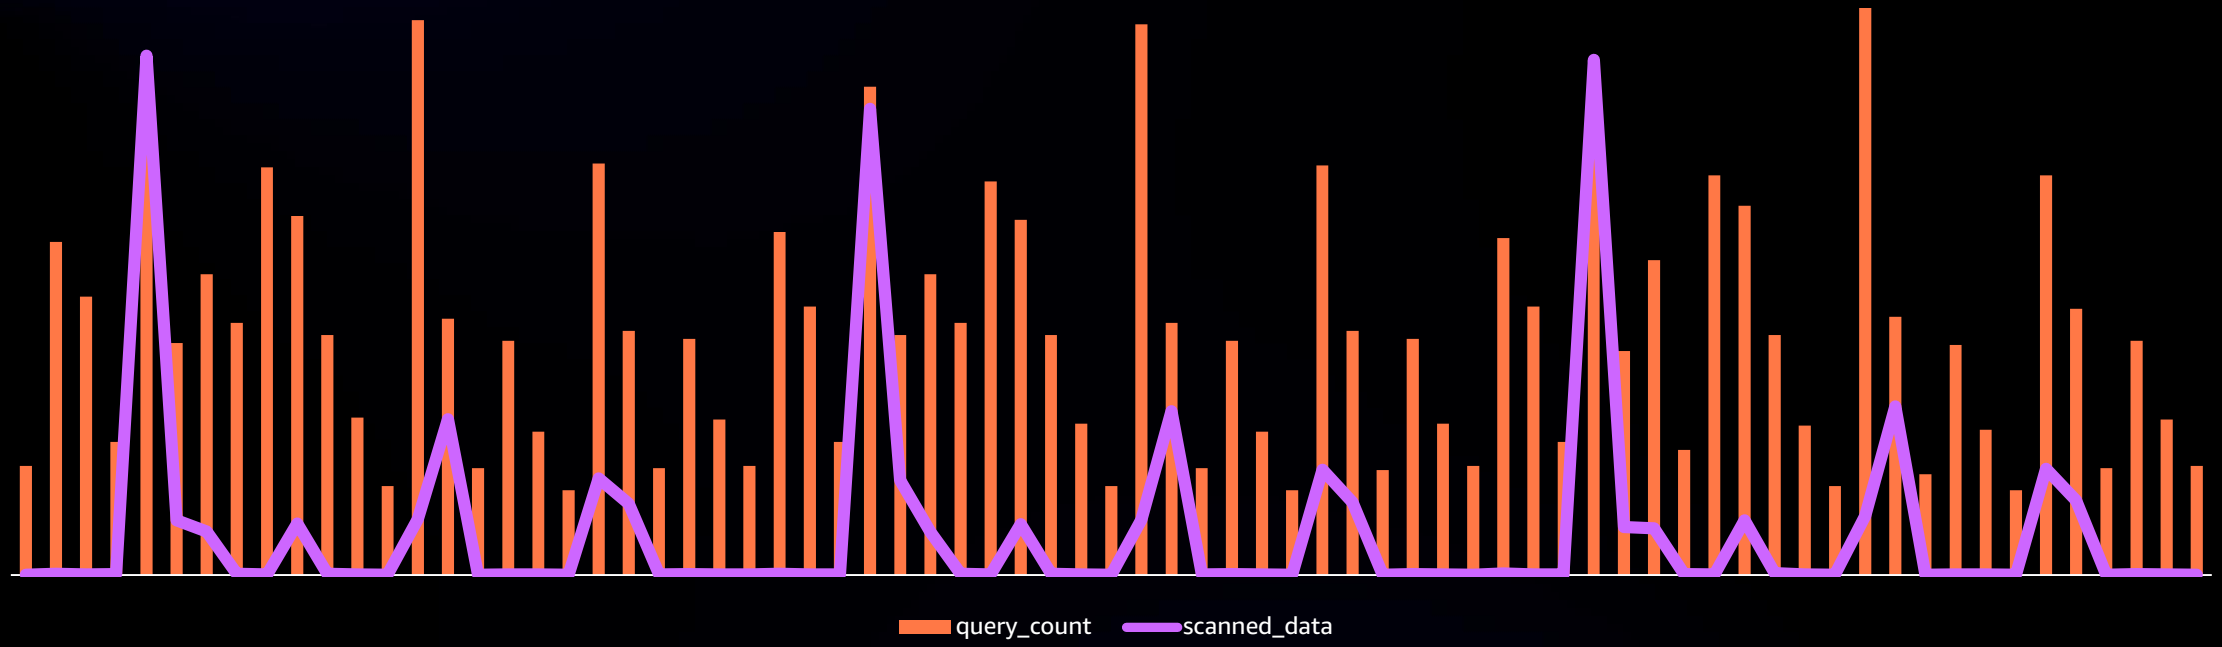


Why

- Automatic provisioning and scaling of resources so there is no manual intervention to create and manage compute
- No need to pay for idle durations

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Serverless use cases: Steady state with spikes



Why

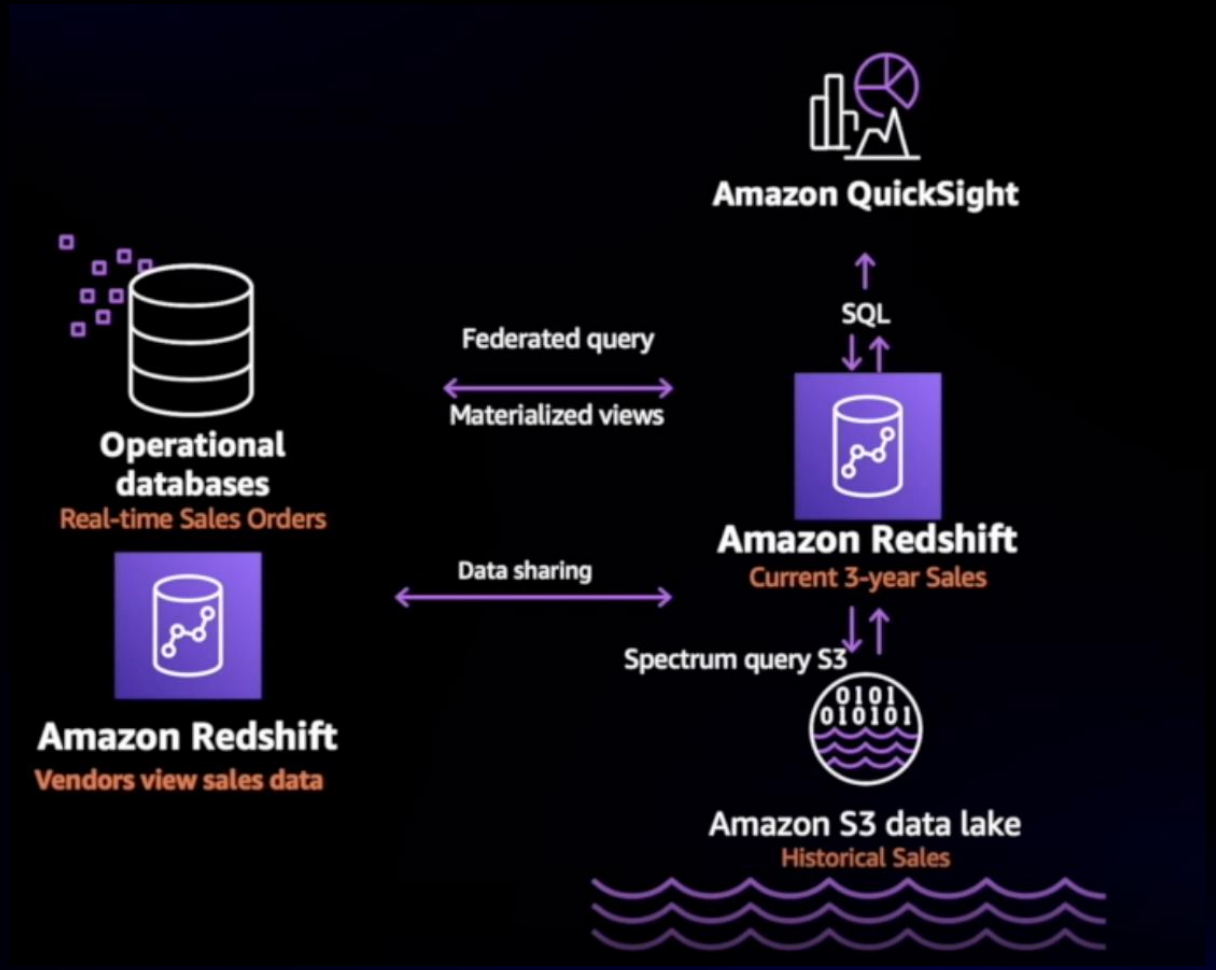
- Seamlessly and automatically scale up and down according to the workload spikes
- Consistent performance at scale

aws

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Demo Scenario

Analyze sales data and come up with a holiday promotion



© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Summary

- Gain insights from all your data, without having to manage data warehouse clusters—just pay for use
- Benefit from serverless simplicity without making any changes to existing analytics applications
- Continue to benefit from Amazon Redshift provisioned cluster for fine-grained customization and control over capacity to optimize cost and performance

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Next steps

Try Amazon Redshift Serverless preview today

- For questions – send email to redshift-preview-serverless@amazon.com

Learn more about Amazon Redshift Serverless at

- Feature page : http://aws.amazon.com/redshift/redshift-serverless
- Blog : https://aws.amazon.com/blogs/aws/introducing-amazon-redshift-serverless-run-analytics-at-any-scale-without-having-to-manage-data-warehouse-infrastructure

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Yan Leshinsky
yanlesh@amazon.com

Neeraja Rentachintala
neerajre@amazon.com

© 2021, Amazon Web Services, Inc. or its affiliates. All rights reserved.