
AWS
re:Invent
NOV. 28 – DEC. 2, 2022 | LAS VEGAS, NV

ANT320-R

Get better price performance in cloud data warehousing with Amazon Redshift

Ippokratis Pandis

VP/Distinguished Engineer
Amazon Web Services

Stefan Gromoll

Performance Engineering Manager
Amazon Web Services

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift: Best–in-class data warehouse

TENS OF THOUSANDS OF CUSTOMERS USE AMAZON REDSHIFT AND PROCESS EXABYTES OF DATA PER DAY

Easy, secure & reliable



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

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Why does price performance matter?



In the days of on-premises data warehousing . . .

Will it be fast enough for the next 1, 2, or 5 years?

What does it cost?

What do I do if I need better performance?
What if my data grows faster than I planned?
What if I have more users than I expected?
What if I have peaks of high usage that exceed my capacity?
What if my users are idle on weekends and on holidays?



Today in the cloud . . .

What is the **price performance**?

Does it **scale**?

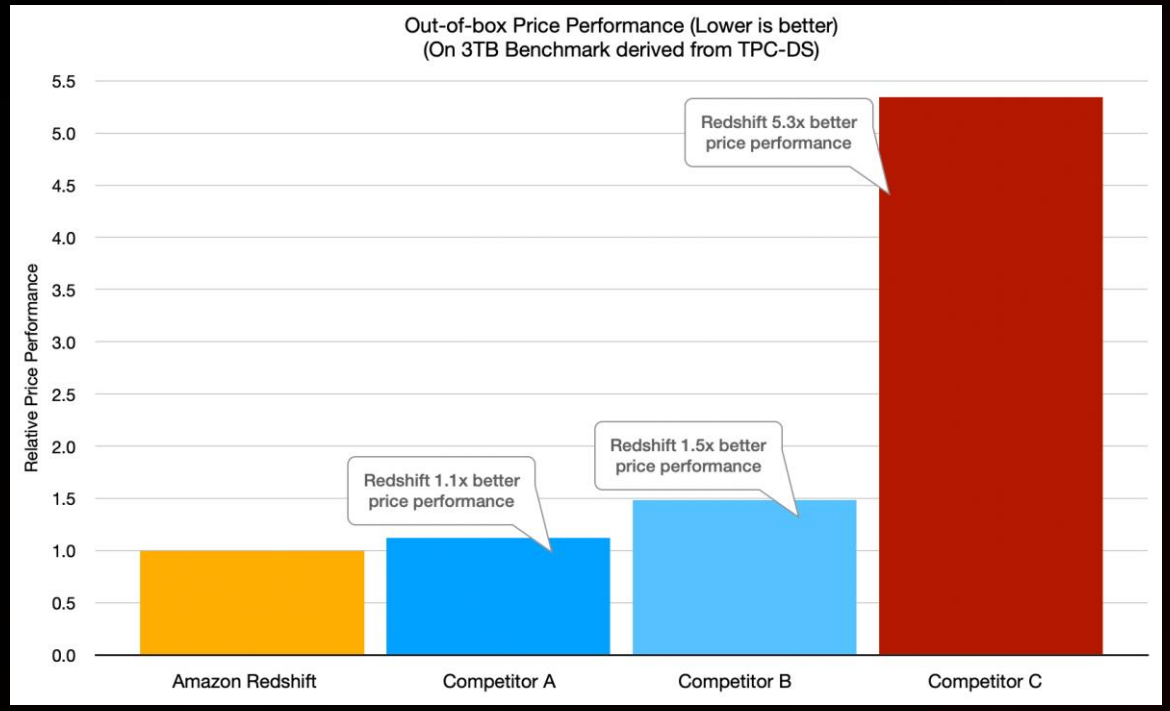
With a highly elastic warehouse like Amazon Redshift, price performance and scale tell you if a data warehouse is the most cost-effective

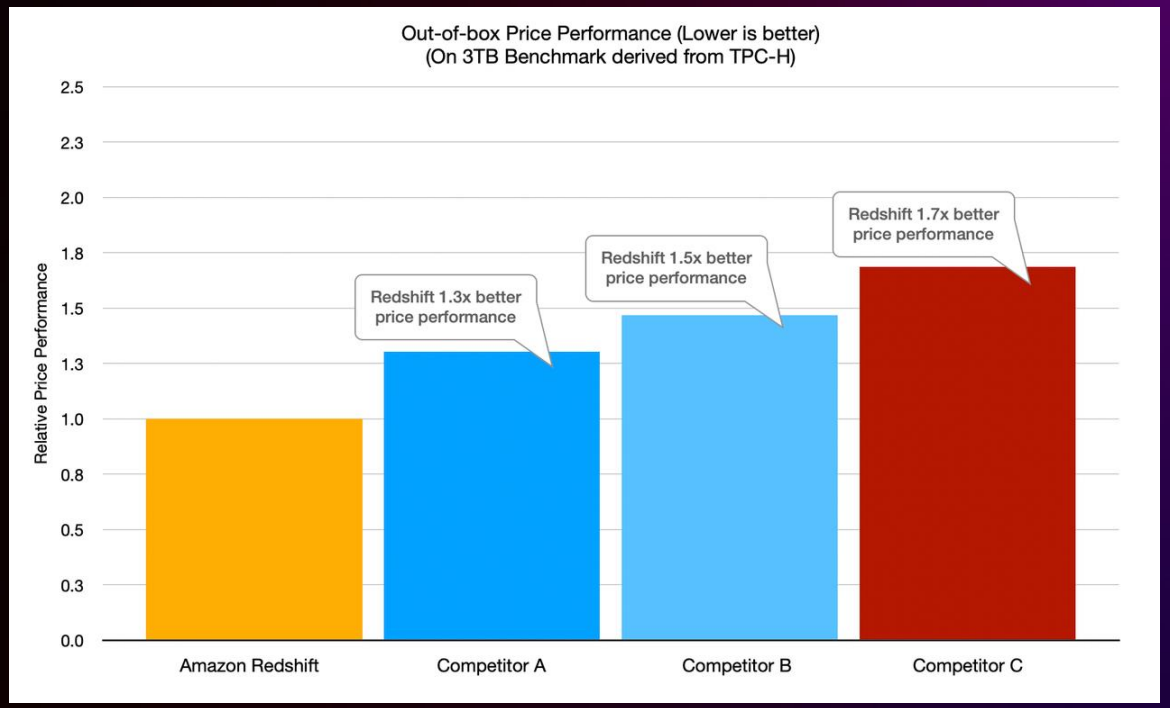
aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The best price performance

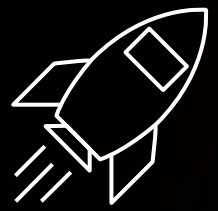
AMAZON REDSHIFT CONTINUES TO HAVE BETTER PRICE PERFORMANCE THAN OTHER CLOUD DATA WAREHOUSES





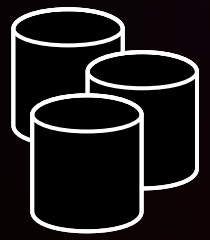
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

What is scale?



Scaling for workload performance

Scale your data warehouse so queries run faster



Scaling for growing data

Scale your data warehouse to get consistent performance as your data grows



Scaling for increased users

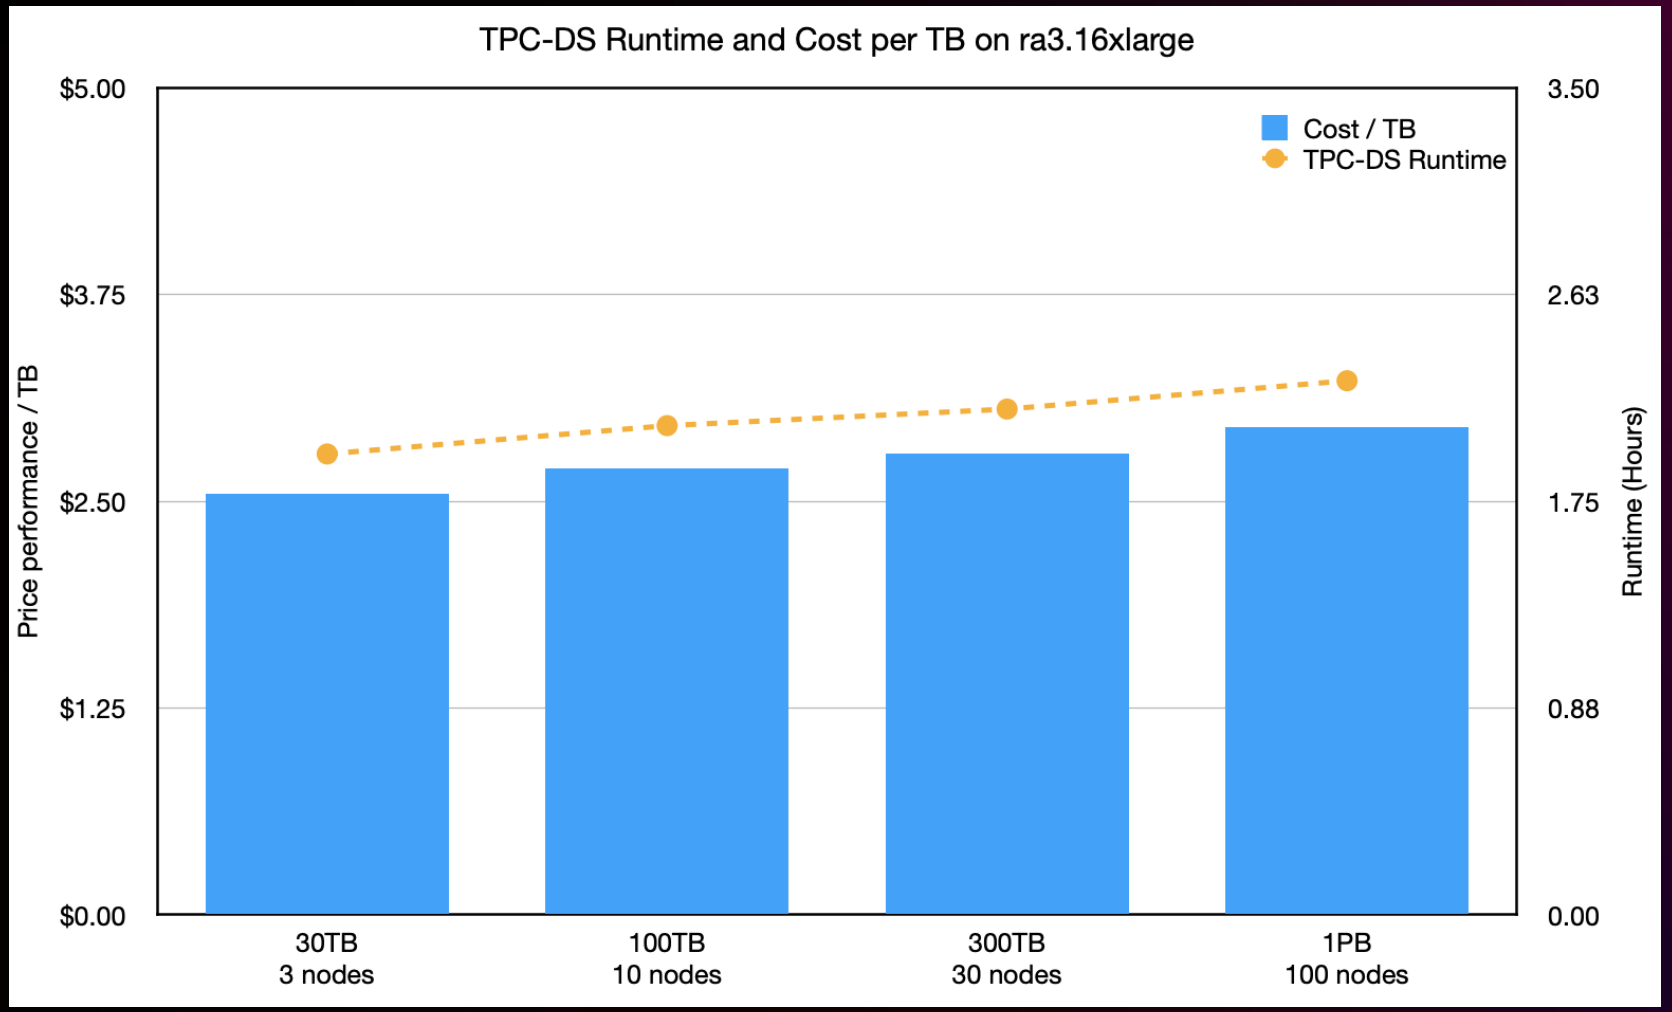
Scale your data warehouse to get increased query throughput as user base grows

A highly elastic system like Amazon Redshift can give near-linear price performance across wide ranges of performance, data, and query concurrencies

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Scale up as your data grows

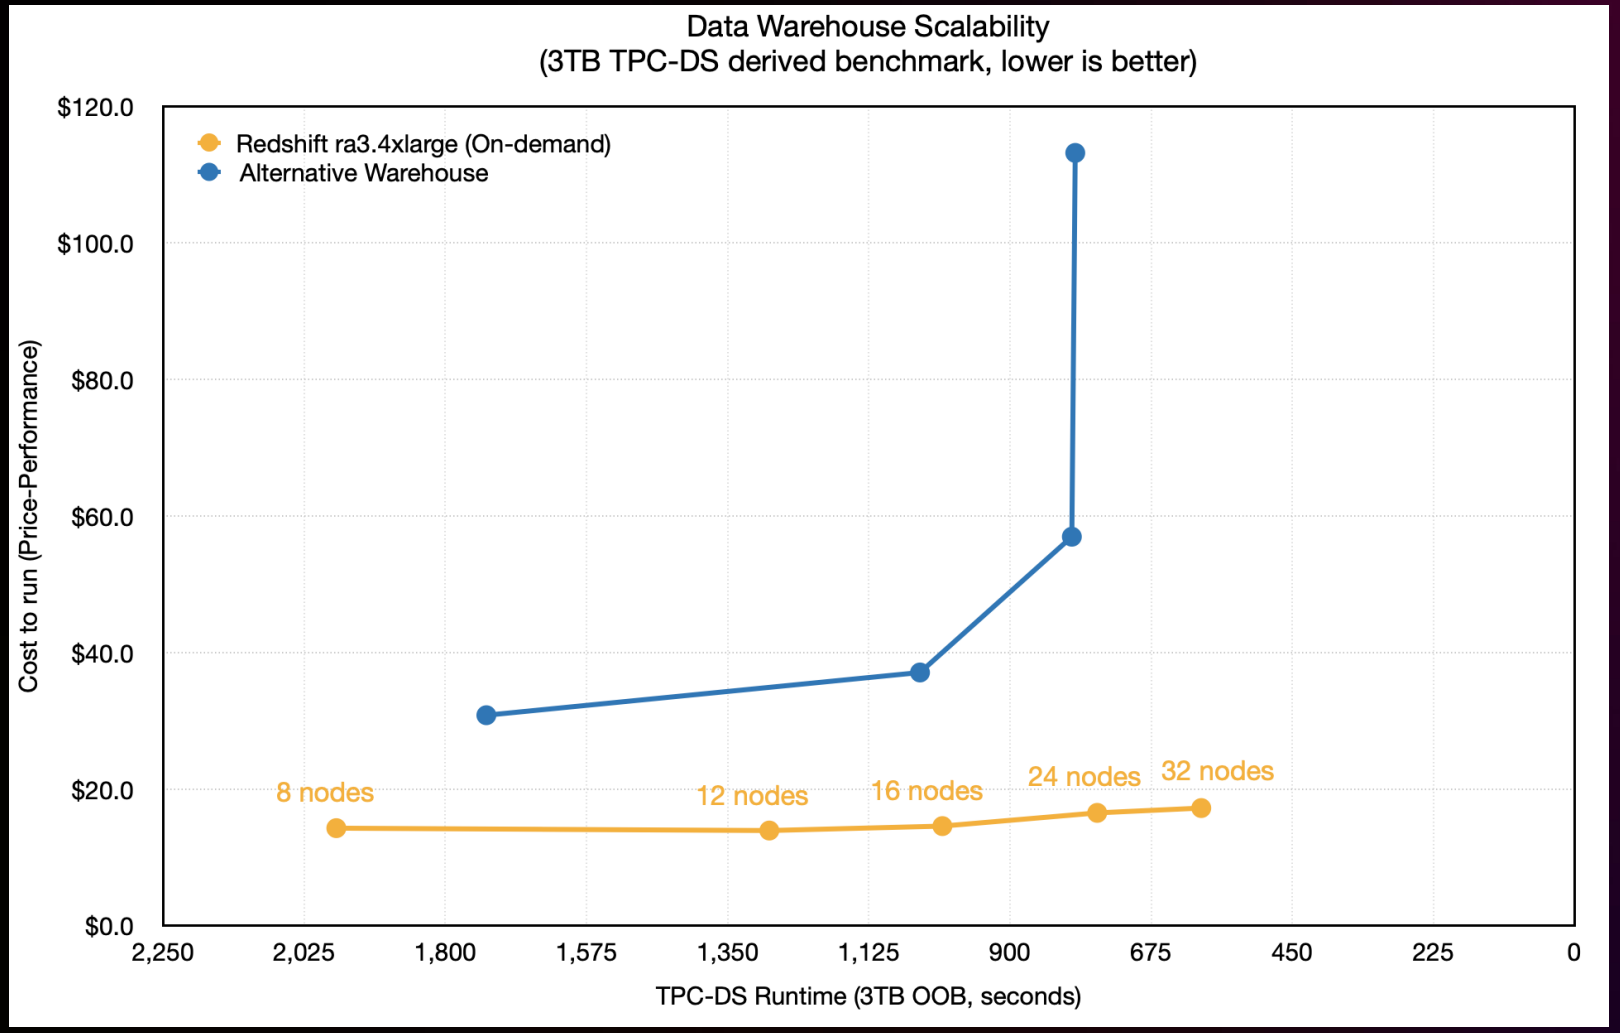
AMAZON REDSHIFT DELIVERS NEAR-LINEAR PRICE PERFORMANCE TO PETABYTES AND BEYOND



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Scale up to meet your performance SLA

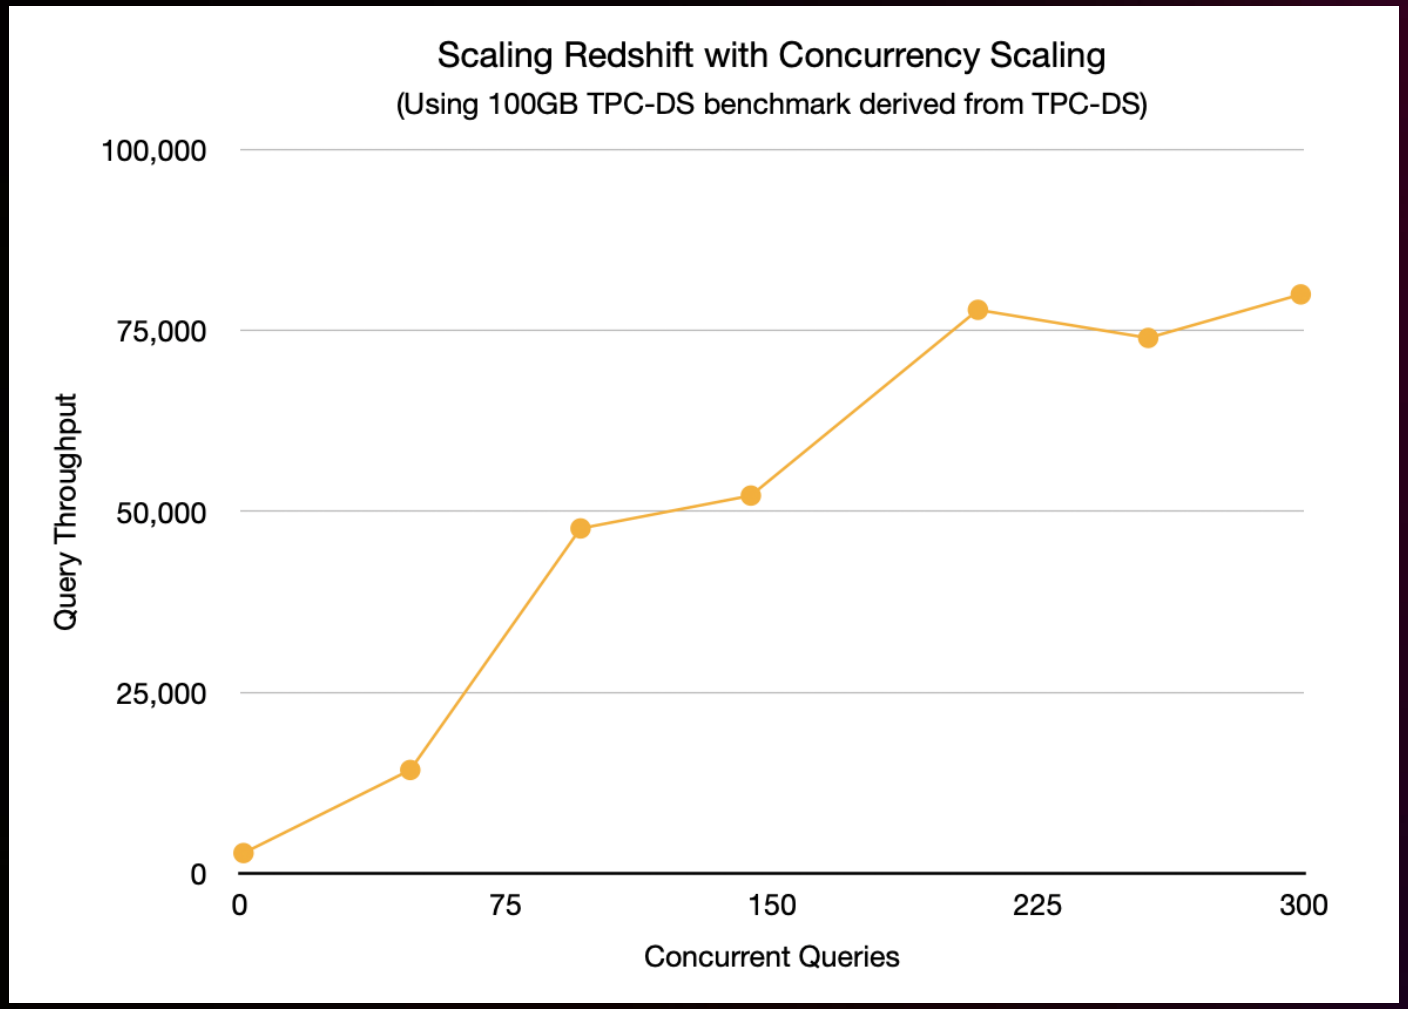
AMAZON REDSHIFT IS DESIGNED TO SCALE NEAR LINEARLY TO DELIVER PREDICTABLE PRICE PERFORMANCE OVER A WIDE RANGE



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Scale out to meet demand as user base grows

AMAZON REDSHIFT AUTOSCALES AS CONCURRENCY GROWS TO DELIVER CONSISTENT QUERY THROUGHPUT FOR SPIKY WORKLOADS



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Our approach to performance optimization

We have customer-obsessed focus on continuously improving real-world price-performance of our customers' workloads

- Performance optimization is never-ending
- How do we decide where to invest?



Customer feedback

New workloads and new applications built on Amazon Redshift



Fleet data

Benefit of large customer base is vast performance metadata



Anticipate future needs

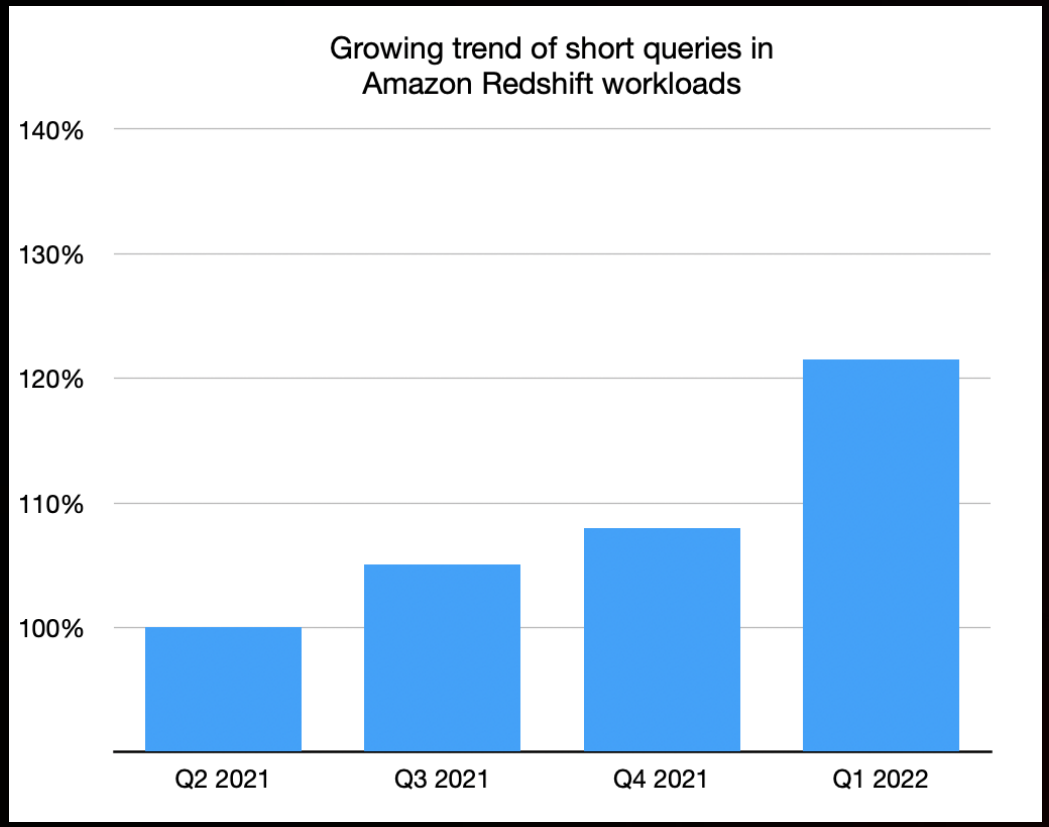
Invest in improvements that will improve price/perf for all customers (e.g. hardware acceleration)

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Example: Customer-obsessed performance

Request from customer: Can we bring new workload to Amazon Redshift with very high-throughput for very short <1 s queries?



Additional data from fleet confirms growing trend

Deep dive on entire short query lifetime:

1. Planning times small, but non-negligible for ~100ms query => reduce mutex contention in planner
2. Optimize query rewriter for short queries
3. Improve Concurrency Scaling algorithm to better burst short queries

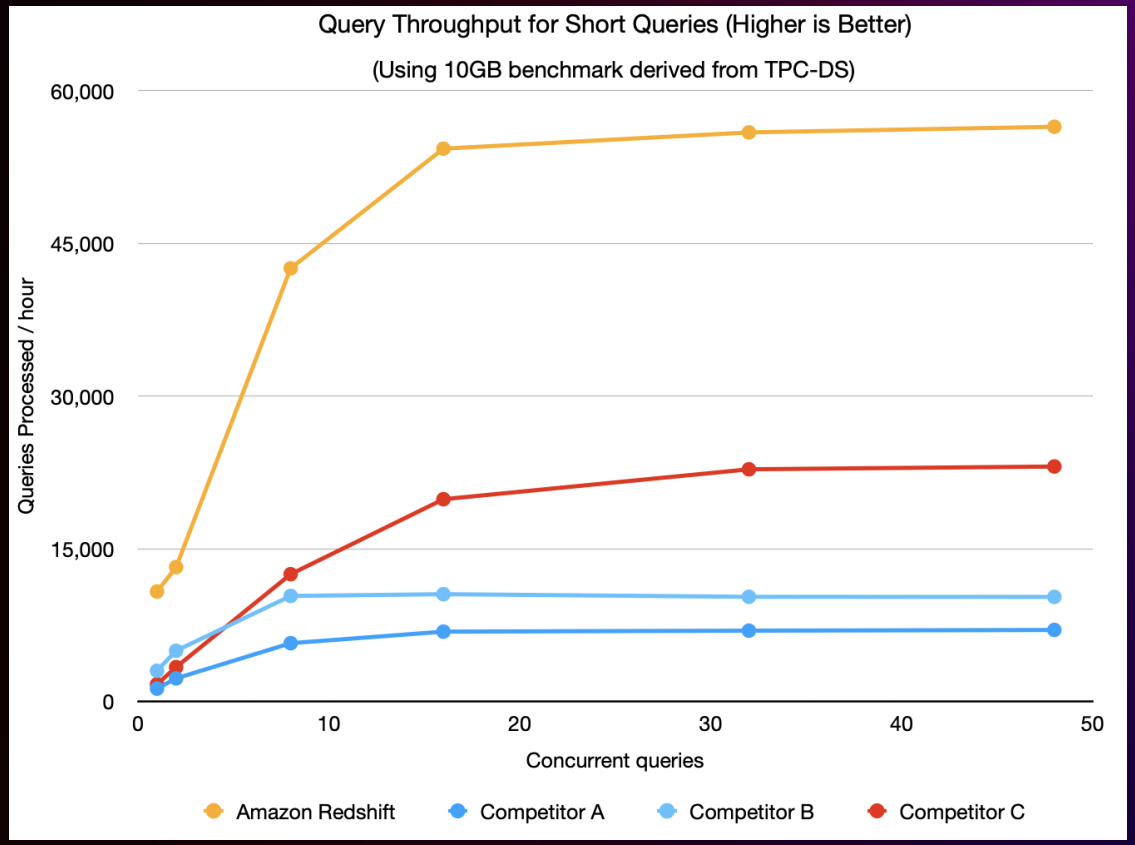
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Example: Customer-obsessed performance

Insight from our customers: Workloads with short queries (<1 sec)
at high concurrencies are becoming more common. Is there room for improvement?

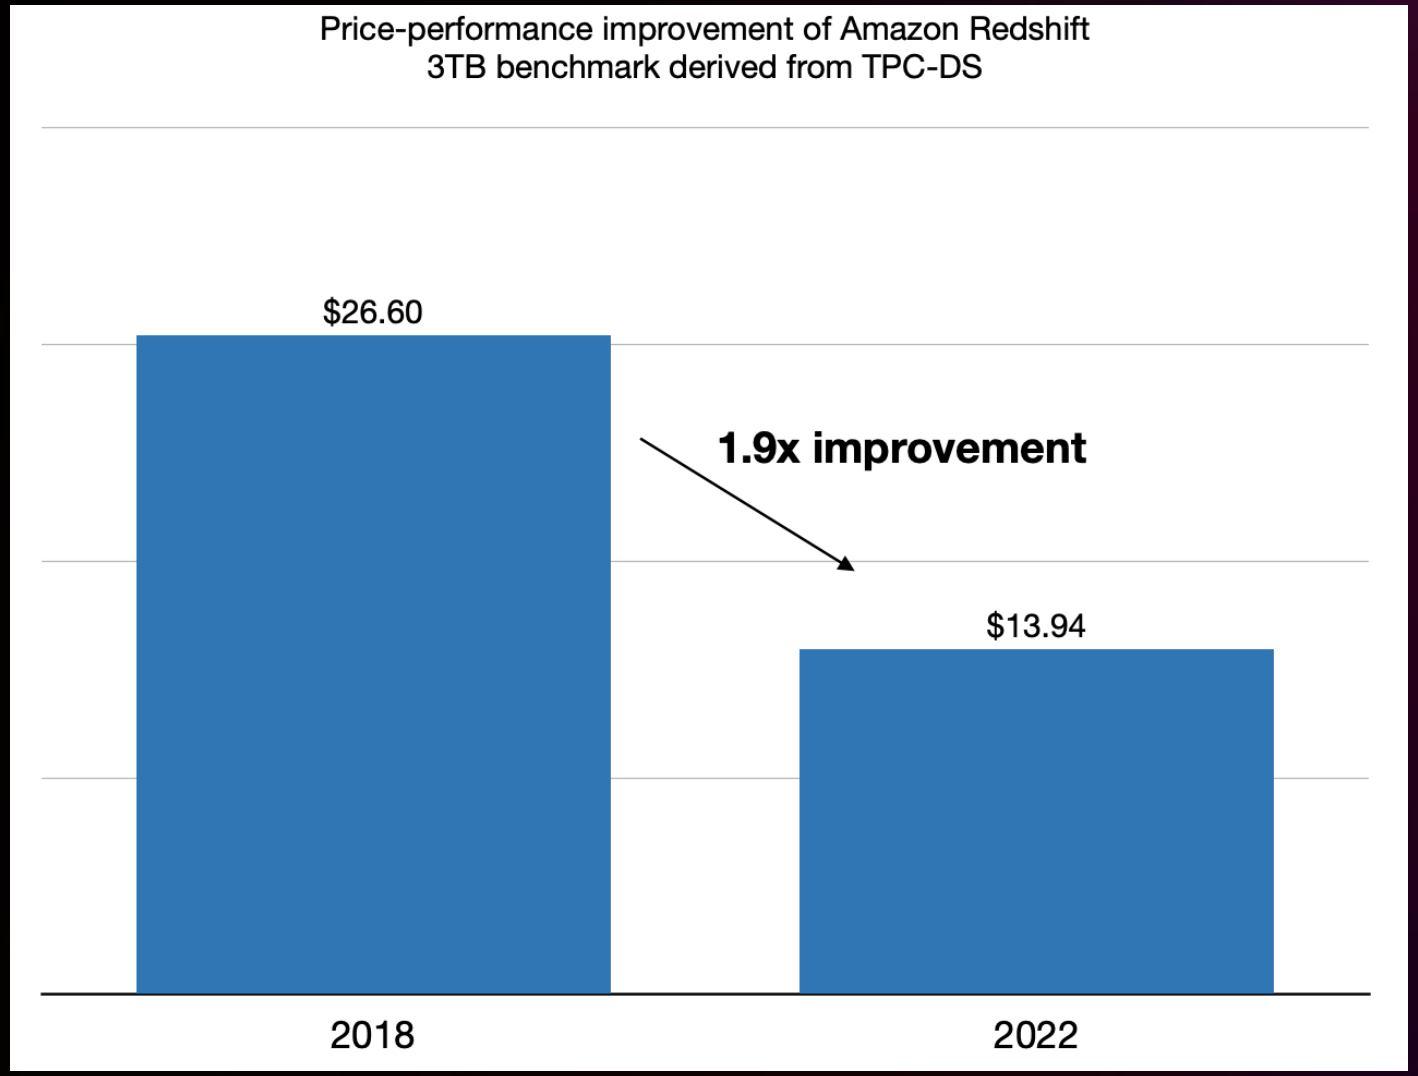
Deep dive on entire short query lifetime:

1. Planning times small, but non-negligible for ~100ms query => reduce mutex contention in planner
2. Optimize query rewriter for short queries
3. Improve Concurrency Scaling algorithm to better burst short queries



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift price performance over time



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The ideal cloud data warehouse . . .

Has performance that scales linearly over a wide range

- Adding 2x the resources gives predictable 2x performance

Has price that scales linearly

- Adding 2x the resources costs 2x as much

Ideally then price ÷ performance is consistent at any scale

- And, price ÷ performance is better than other options

The ideal cloud data warehouse **scales as needed**, delivers **consistent price-performance at any scale**, and **has the best price performance**

(Plus security, availability, features, pace of innovation . . .)

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

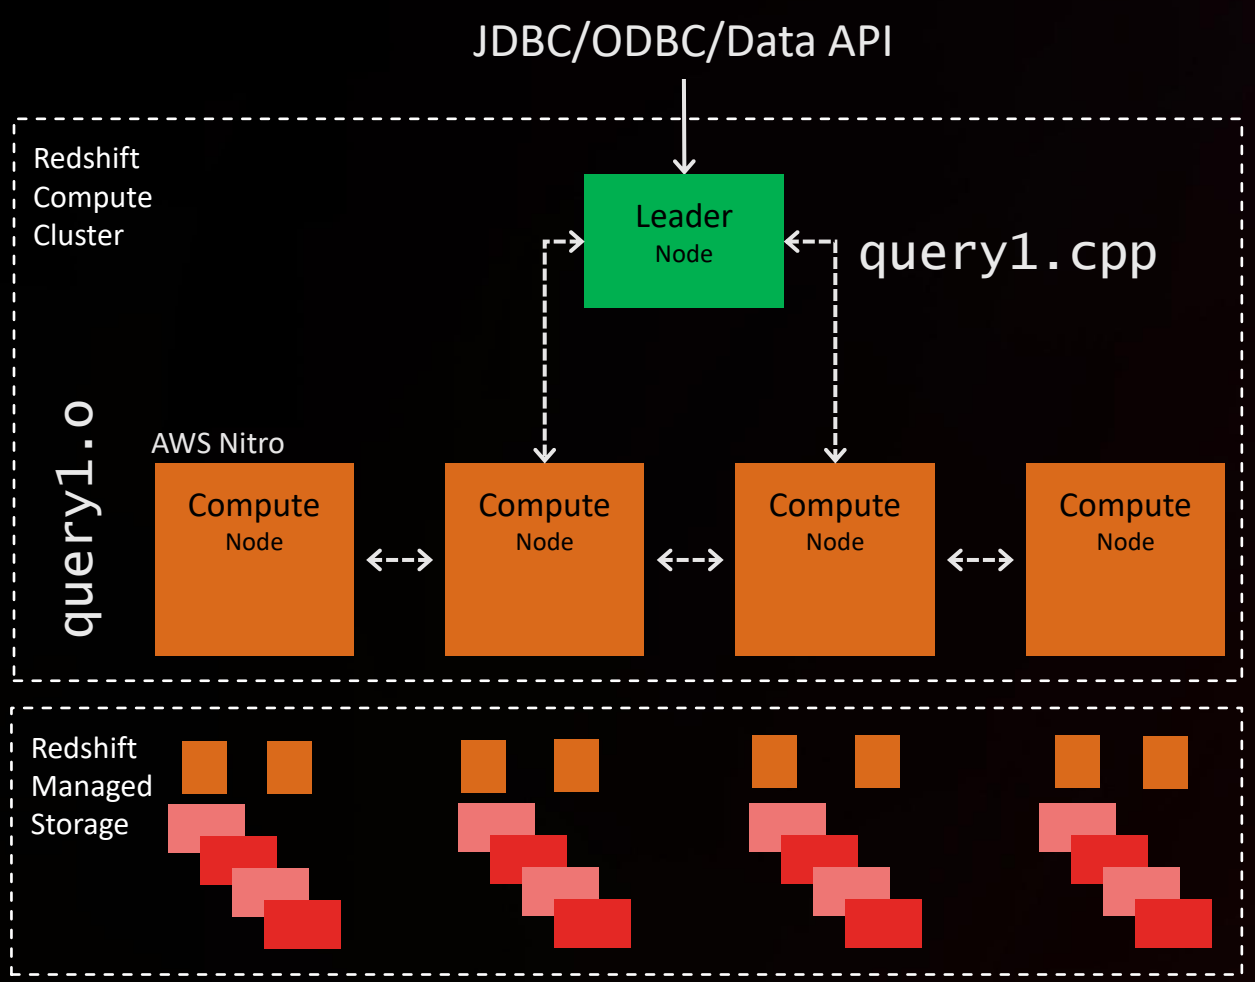
Agenda

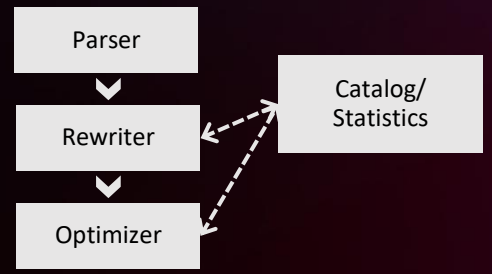
- Best-in-class price performance at any scale
- Looking Amazon Redshift under the hood
- Storage and compute elasticity
- Autonomics and Serverless
- Q&A



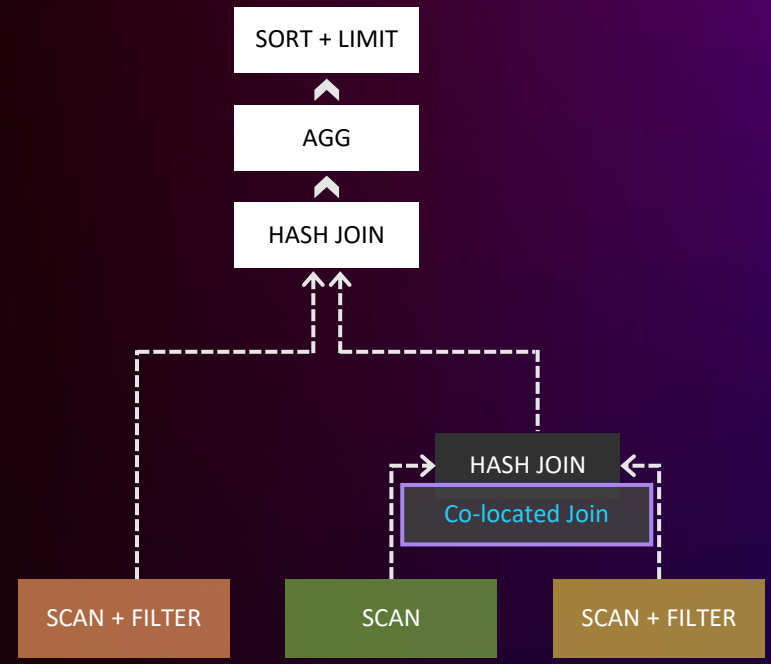
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Executing a query in Amazon Redshift





Min/max pruning
SIMD scans from
local-attached SSDs
AZ64 encoding



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Code generation, prefetching, vectorization…

```
SELECT sum(R.val) FROM R, S WHERE R.key = S.key AND R.val < 50
```

```
// Loop over the tuples of R.
while (scan_step->has_next()) {
  // Get next value for R.key.
  auto field1 = fetcher1.get_next();
  // Get next value for R.val.
  auto field2 = fetcher2.get_next();
  // Apply predicate R.val < 50.
  if (field2 < 50) {
    // Hash R.key and probe the hash table.
    size_t h1 = hash(field1) & (ht1_sz - 1);
    for (auto* p1 = ht1[h1];
         p1 != nullptr; p1 = p1->next) {
      // Evaluate the join condition R.key = S.key.
      if (field1 == p1->field1) sum1 += field2;
    }
  }
}
```

```
// Hash R.key and probe the hash table.
size_t h1 = hash(field1) & (ht1_sz -
1);
// Prefetch the current tuple.
prefetch(ht1[h1]);
// Push the current tuple in the
buffer.
prefetchbuffer1.push(field1, field2);
// Pop a prefetched earlier tuple.
h1, field1, field2 =
    prefetchbuffer1.pop();
// Probe the earlier tuple.
for (auto* p1 = ht1[h1];
     p1 != nullptr; p1 = p1->next) {
   ...
}
```

+ AVX2 Vectorization

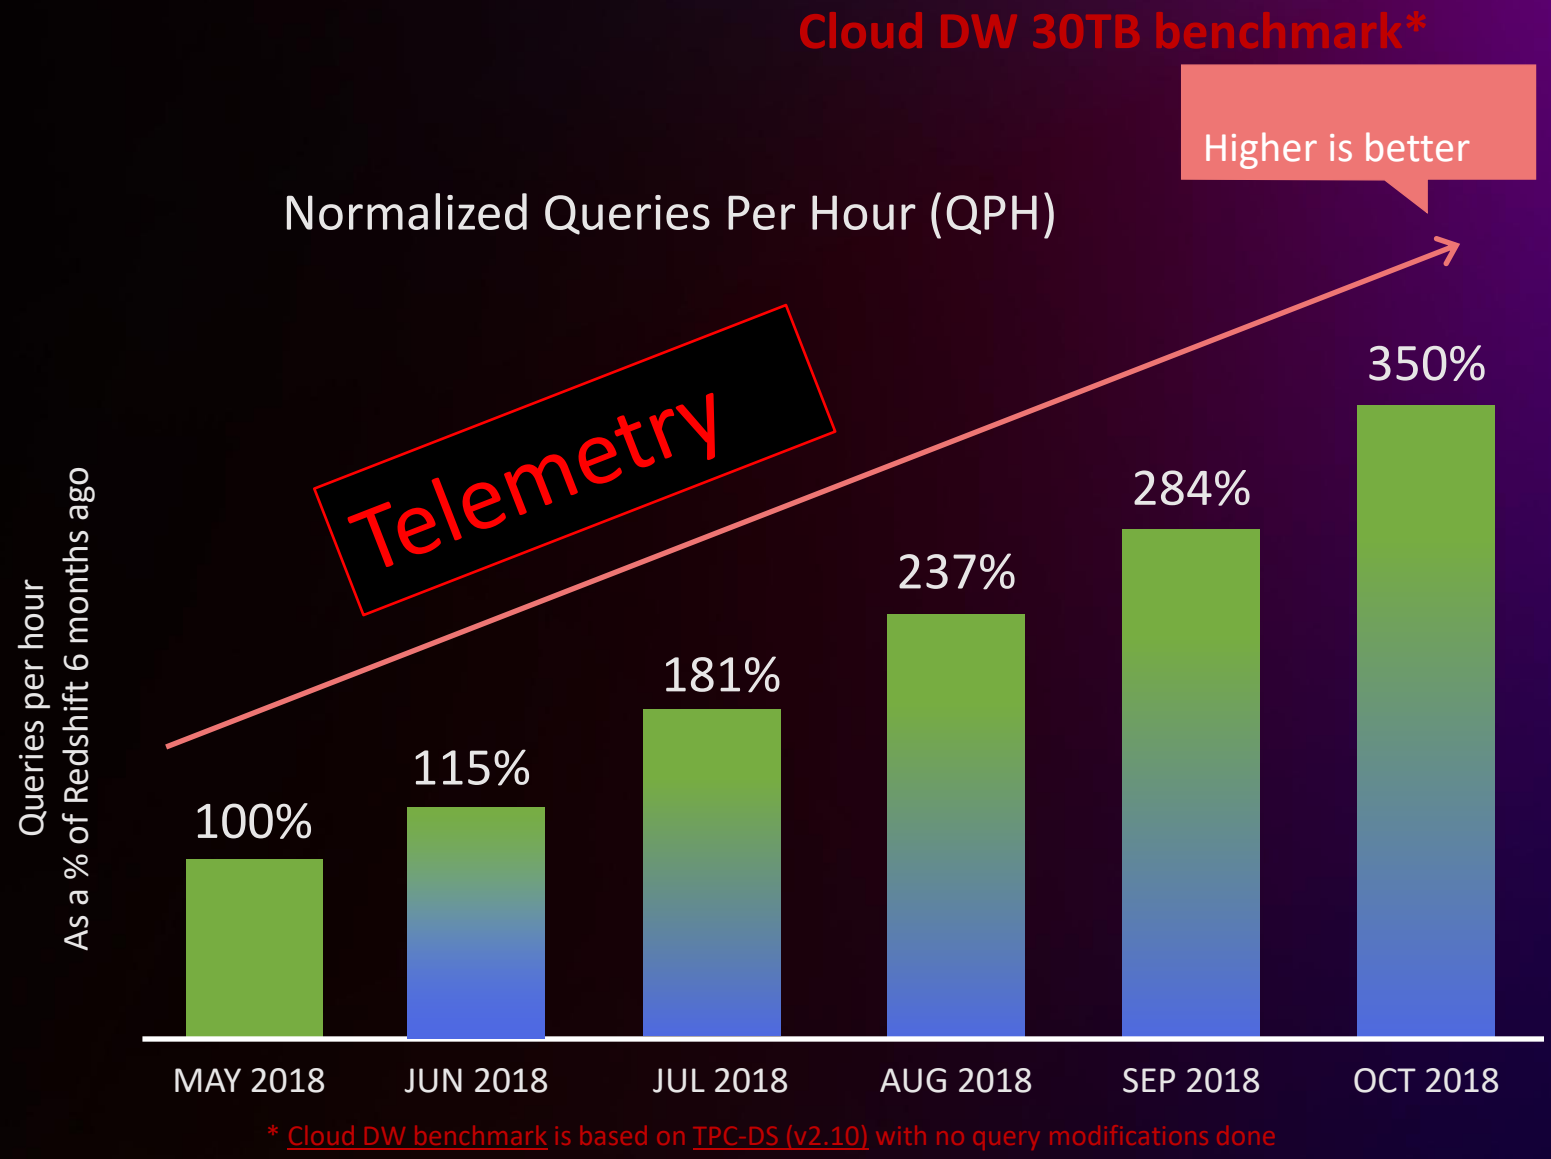
+ Min/Max Pruning, Late Materialization, . . .

aws

Continuous telemetry and benchmarking

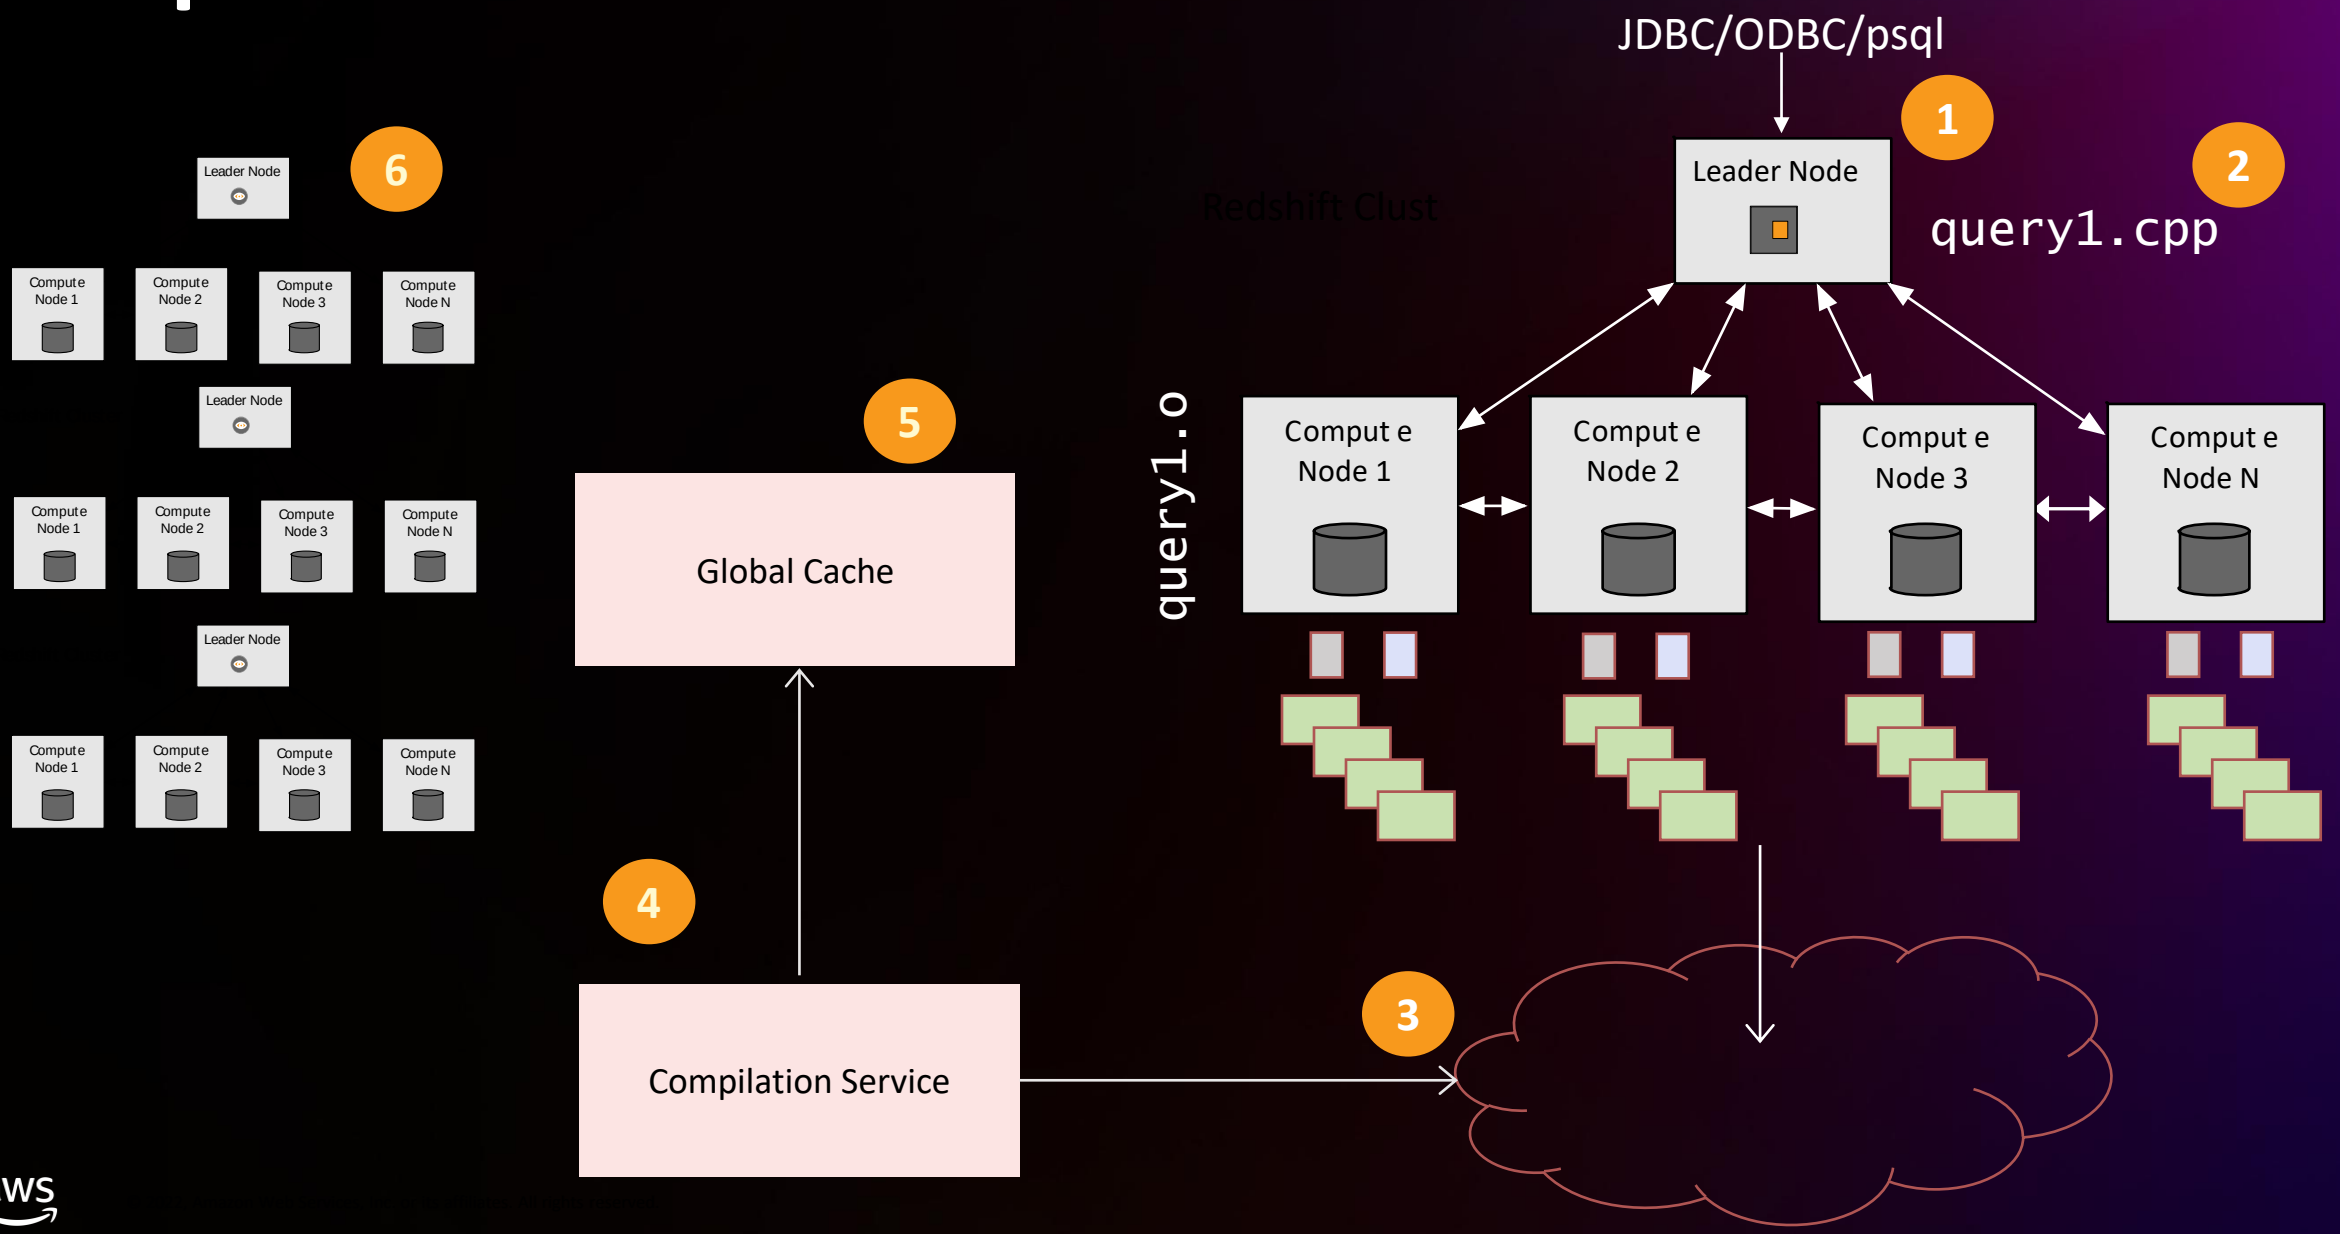
>3x faster between May 2018 and Oct 2018

- Short query acceleration
- Query result caching
- Compiled code cache
- Query rewrites for selective joins
- Hash join memory utilization
- Faster string manipulation
- Cache line prefetching for HashJoins and HashAggregations
- Late materialization



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Compilation as a service



Amazon Redshift String Processing

ADVANCED STRING COMPRESSION ENCODINGS FOR UPTO 64X SPEED-UP OF STRING-HEAVY QUERIES

LoCardCol1 (Boolean value as a CHAR string)	MedCardCol3 (Country name as a VARCHAR string)	HiCardCol2 (Txn_Date+Timestamp 'MMDDYYYY_Hr:Min:Sec as a VARCHAR string)
'Y'	'United States'	'10262022_15:21:55'
'Y'	'Russia'	'10262022_15:22:08'
'N'	'Lichtenstein'	'10272022_01:06:59'
'Y'	'Fiji'	'10272022_03:14:21'
'Y'	'Burkina Faso'	'10282022_13:03:39'
'Y'	'United States'	'10282022_14:09:11'
'N'	'Fiji'	'10282022_15:57:42'

Up to 50x speed-up with low cardinality column with small number of unique values.

5x - 64x speed-up with medium cardinality column with limited set of unique values (country names, only 195 values)

High cardinality column.

New! BYTEDICT encoding is optimal for low-cardinality string columns.

Automatically enabled in Automatic Table Optimization (ATO) and recommendations are viewable in the Redshift Recommendation Engine in the Redshift Console.

aws

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

- Best-in-class price performance at any scale
- Looking Amazon Redshift under the hood
- Storage and compute elasticity
- Autonomics and Serverless
- Q&A



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Amazon Redshift managed storage

SCALE COMPUTE AND STORAGE INDEPENDENTLY



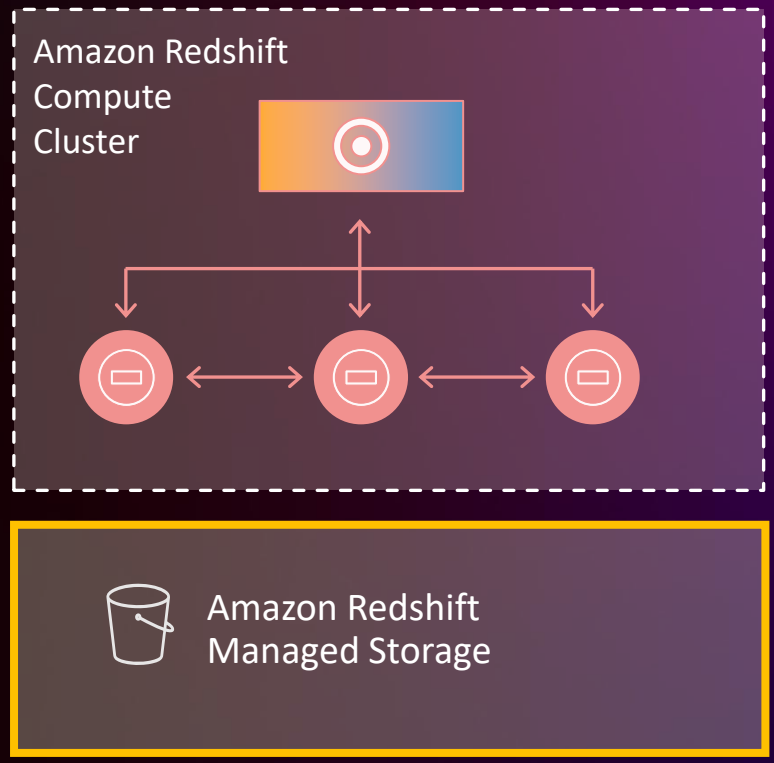
Managed storage



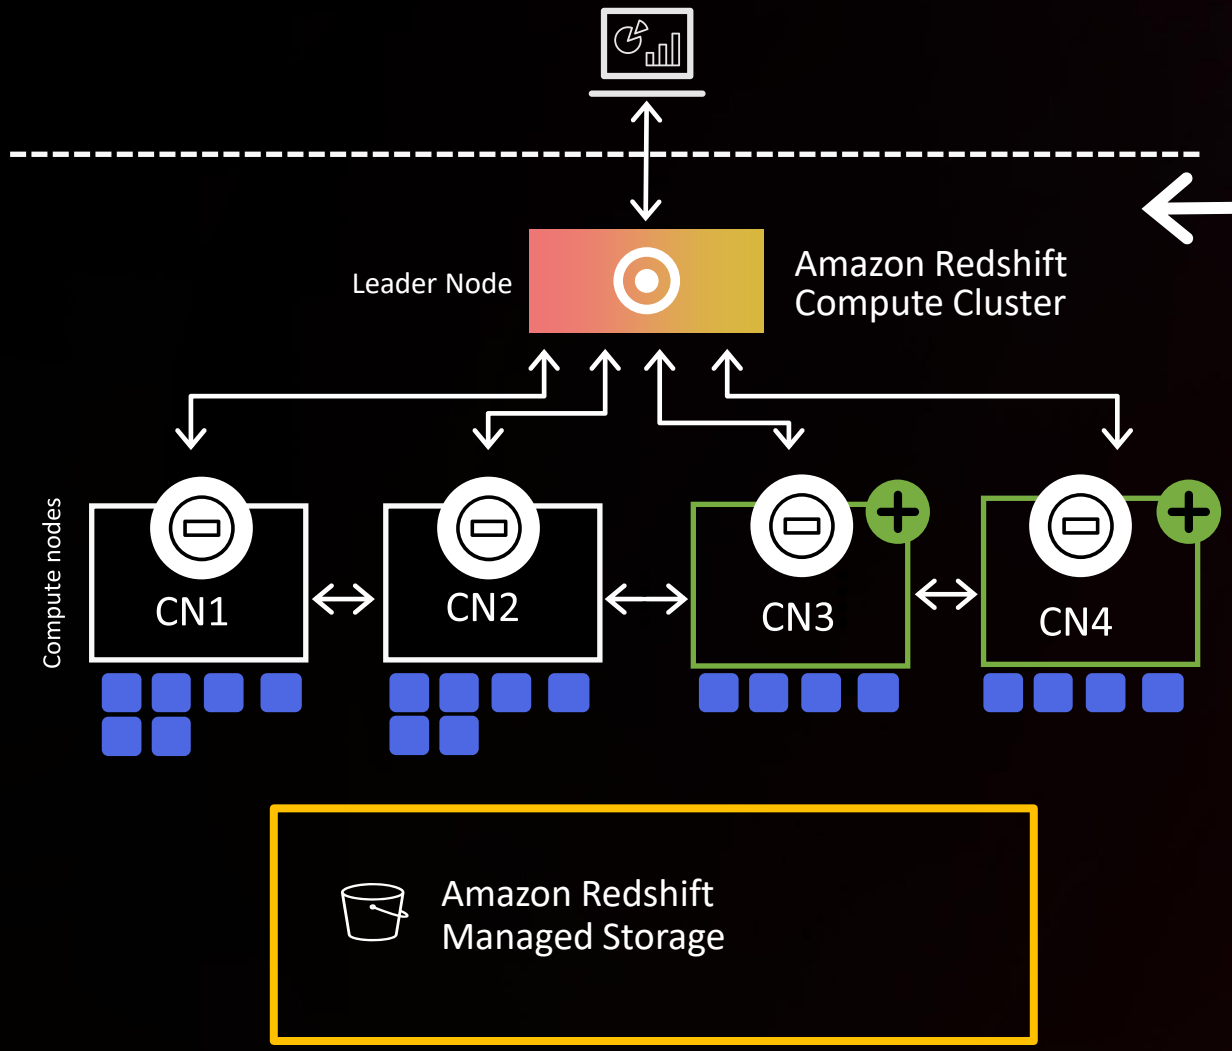
Large high-speed cache

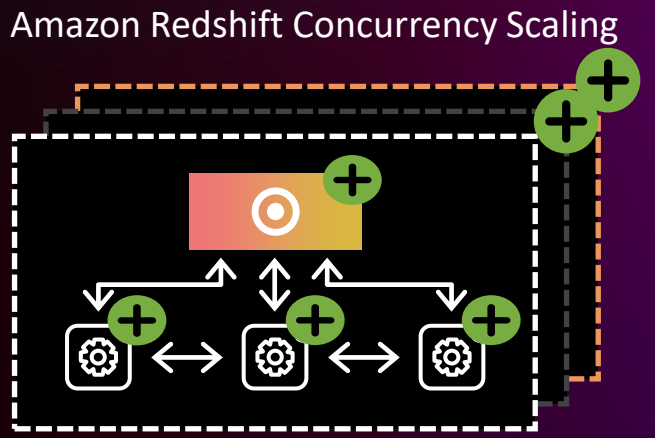


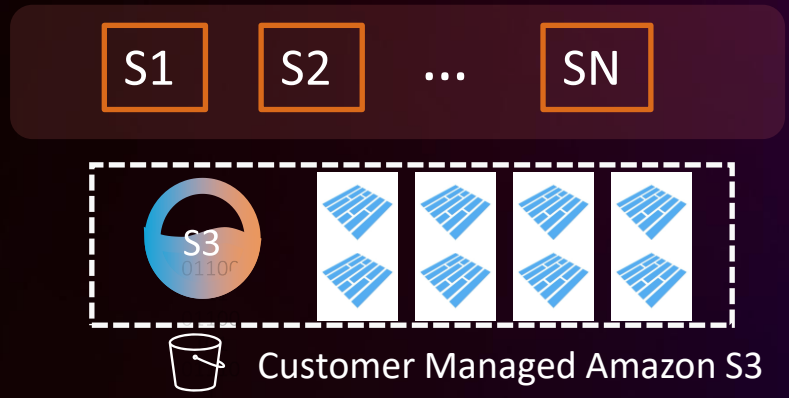
High-bandwidth networking



Compute elasticity

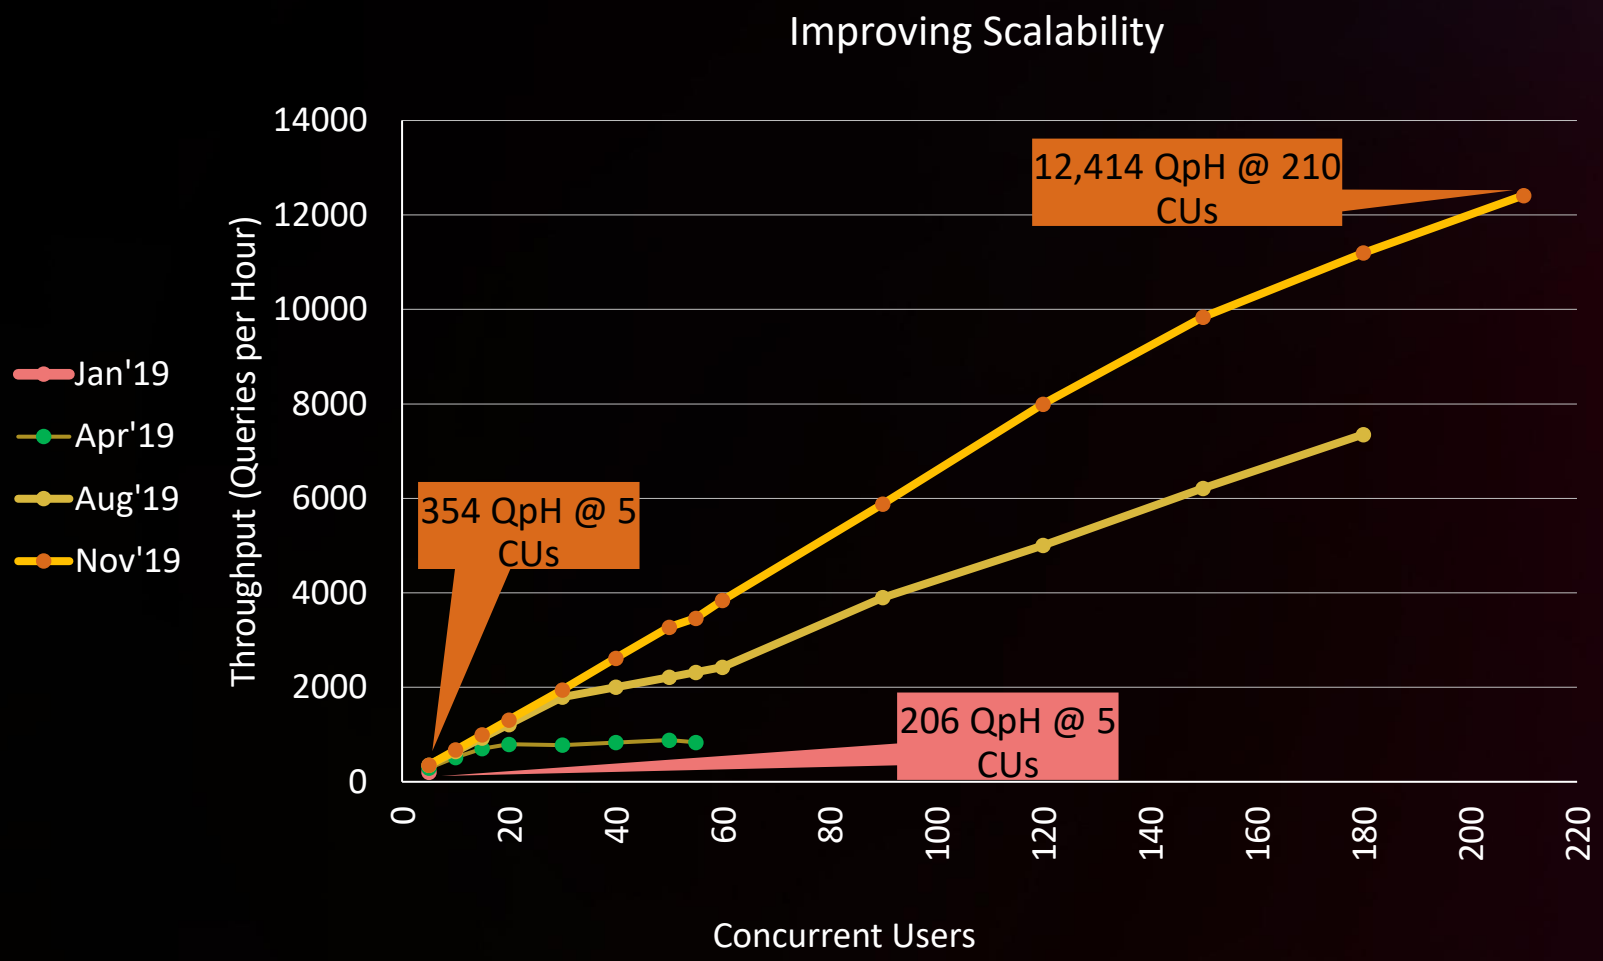






© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Compute elasticity improvements



Cloud DW derived from TPC-DS 3 TB

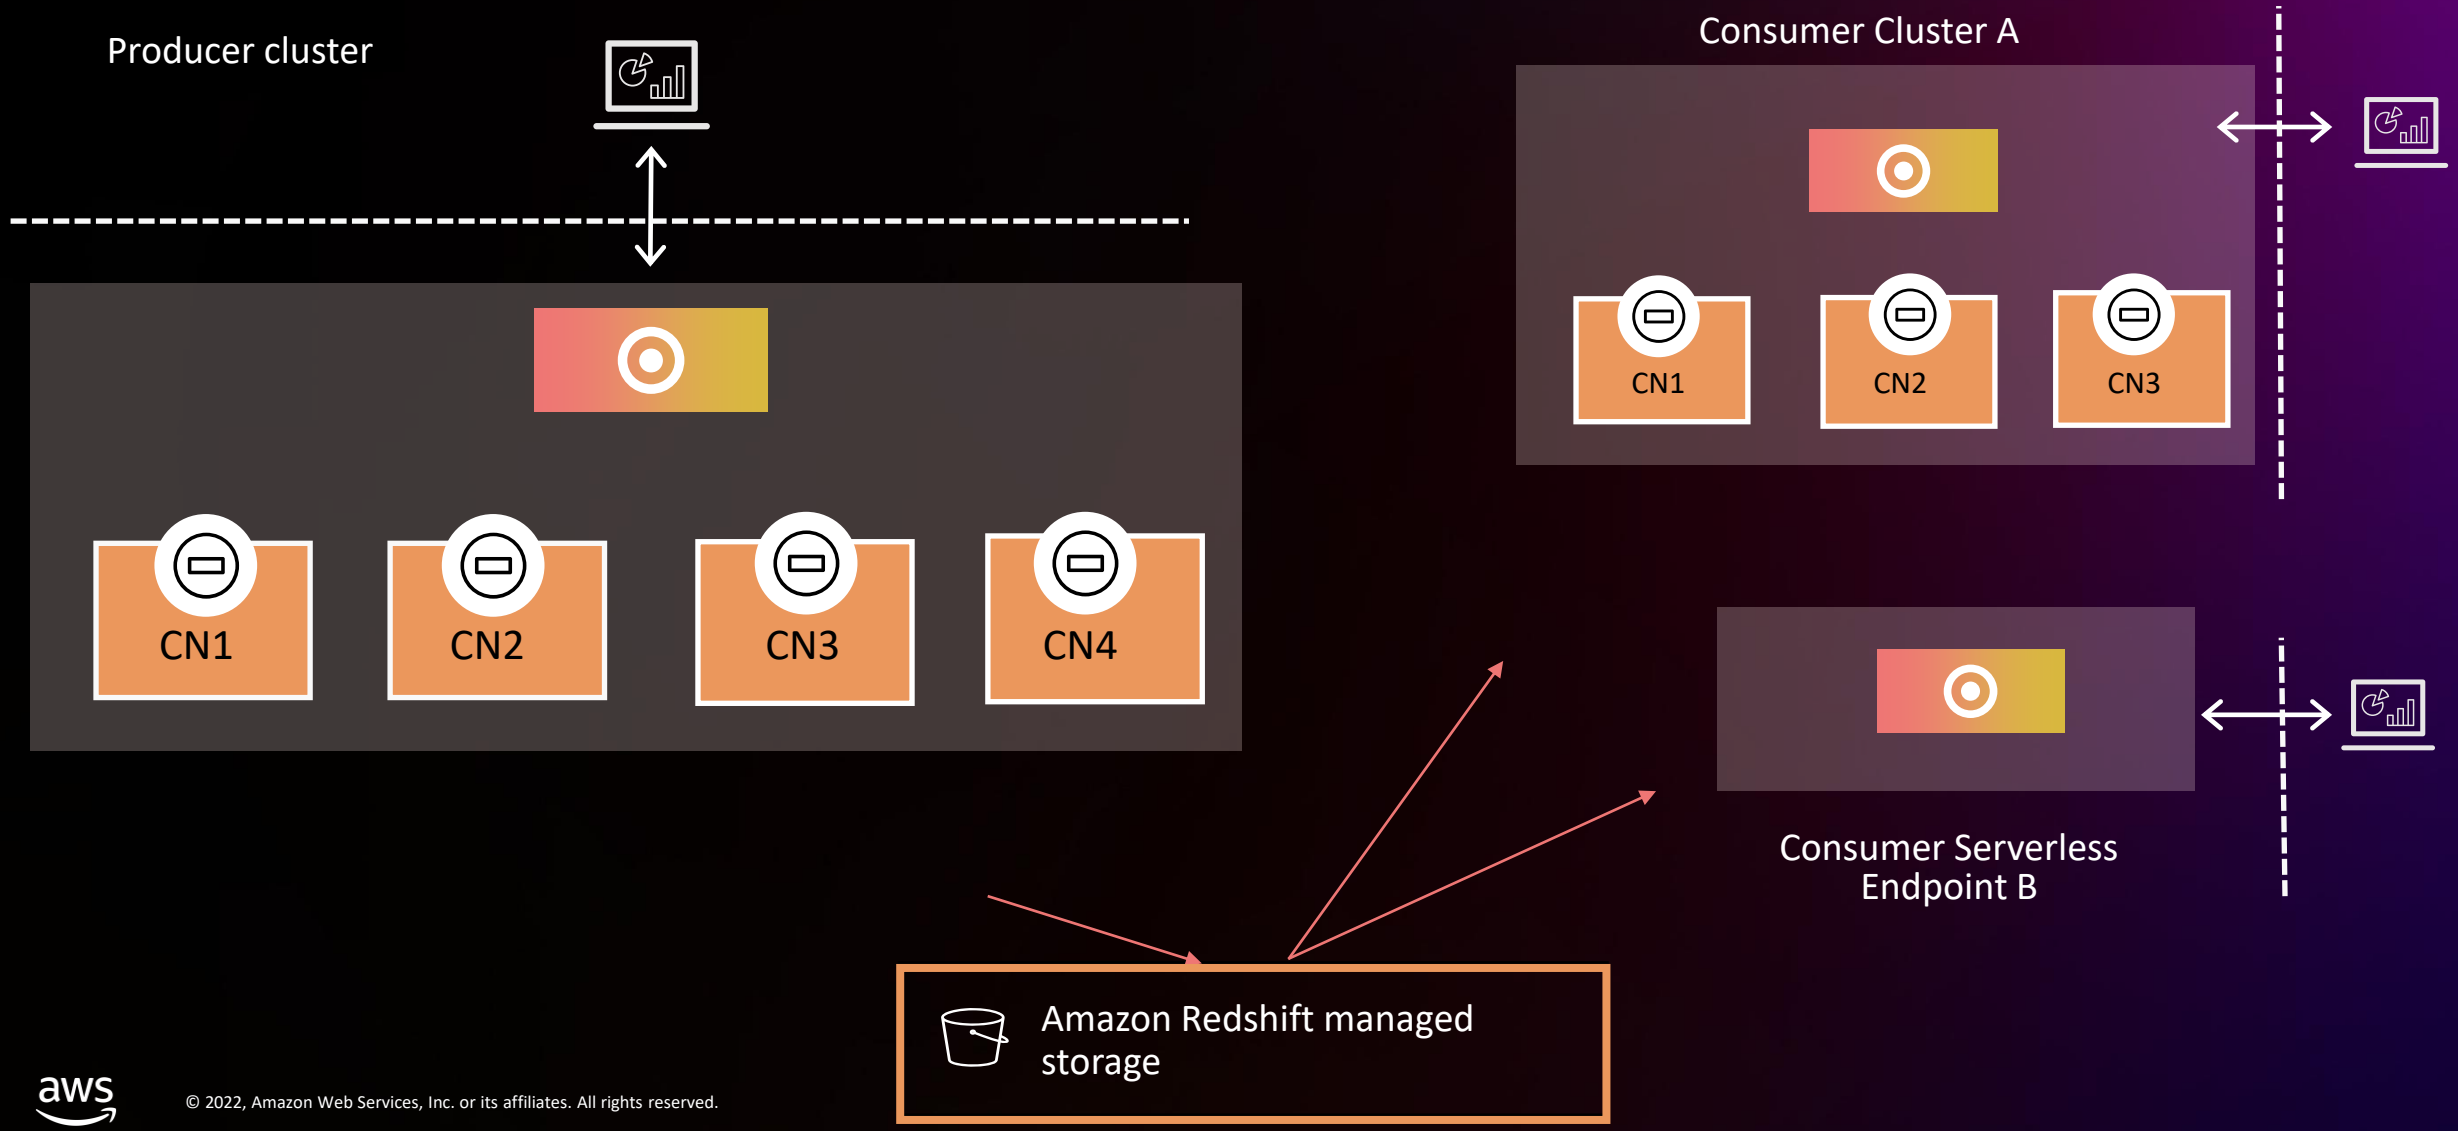
4-node DC2.8XL

>60x improvement in concurrency

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data sharing

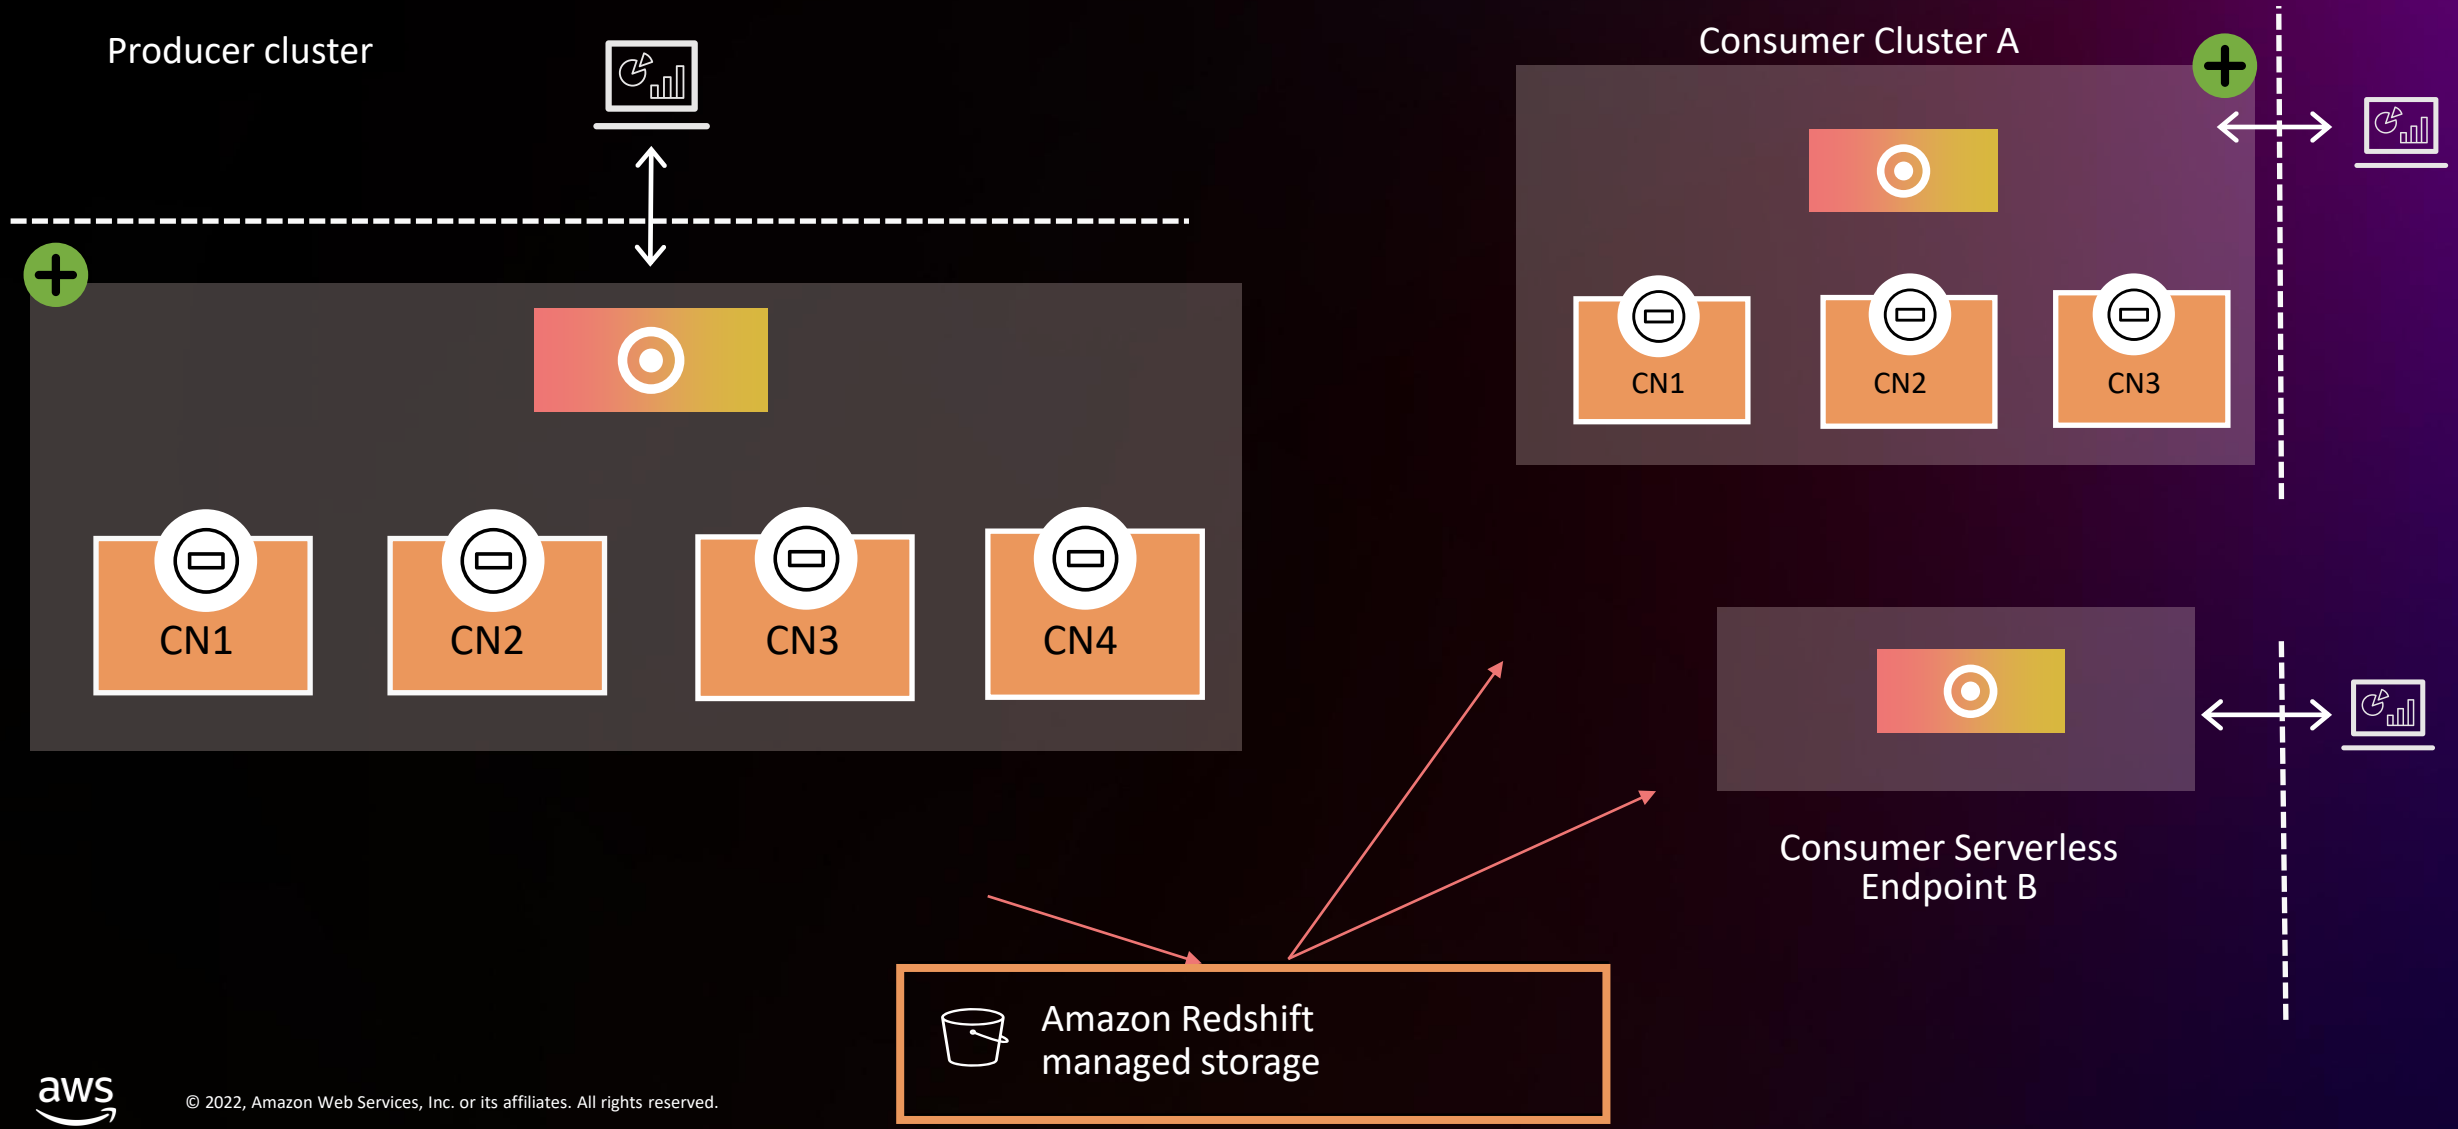
ISOLATE AND SCALE COMPUTE



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data sharing with auto-scaling

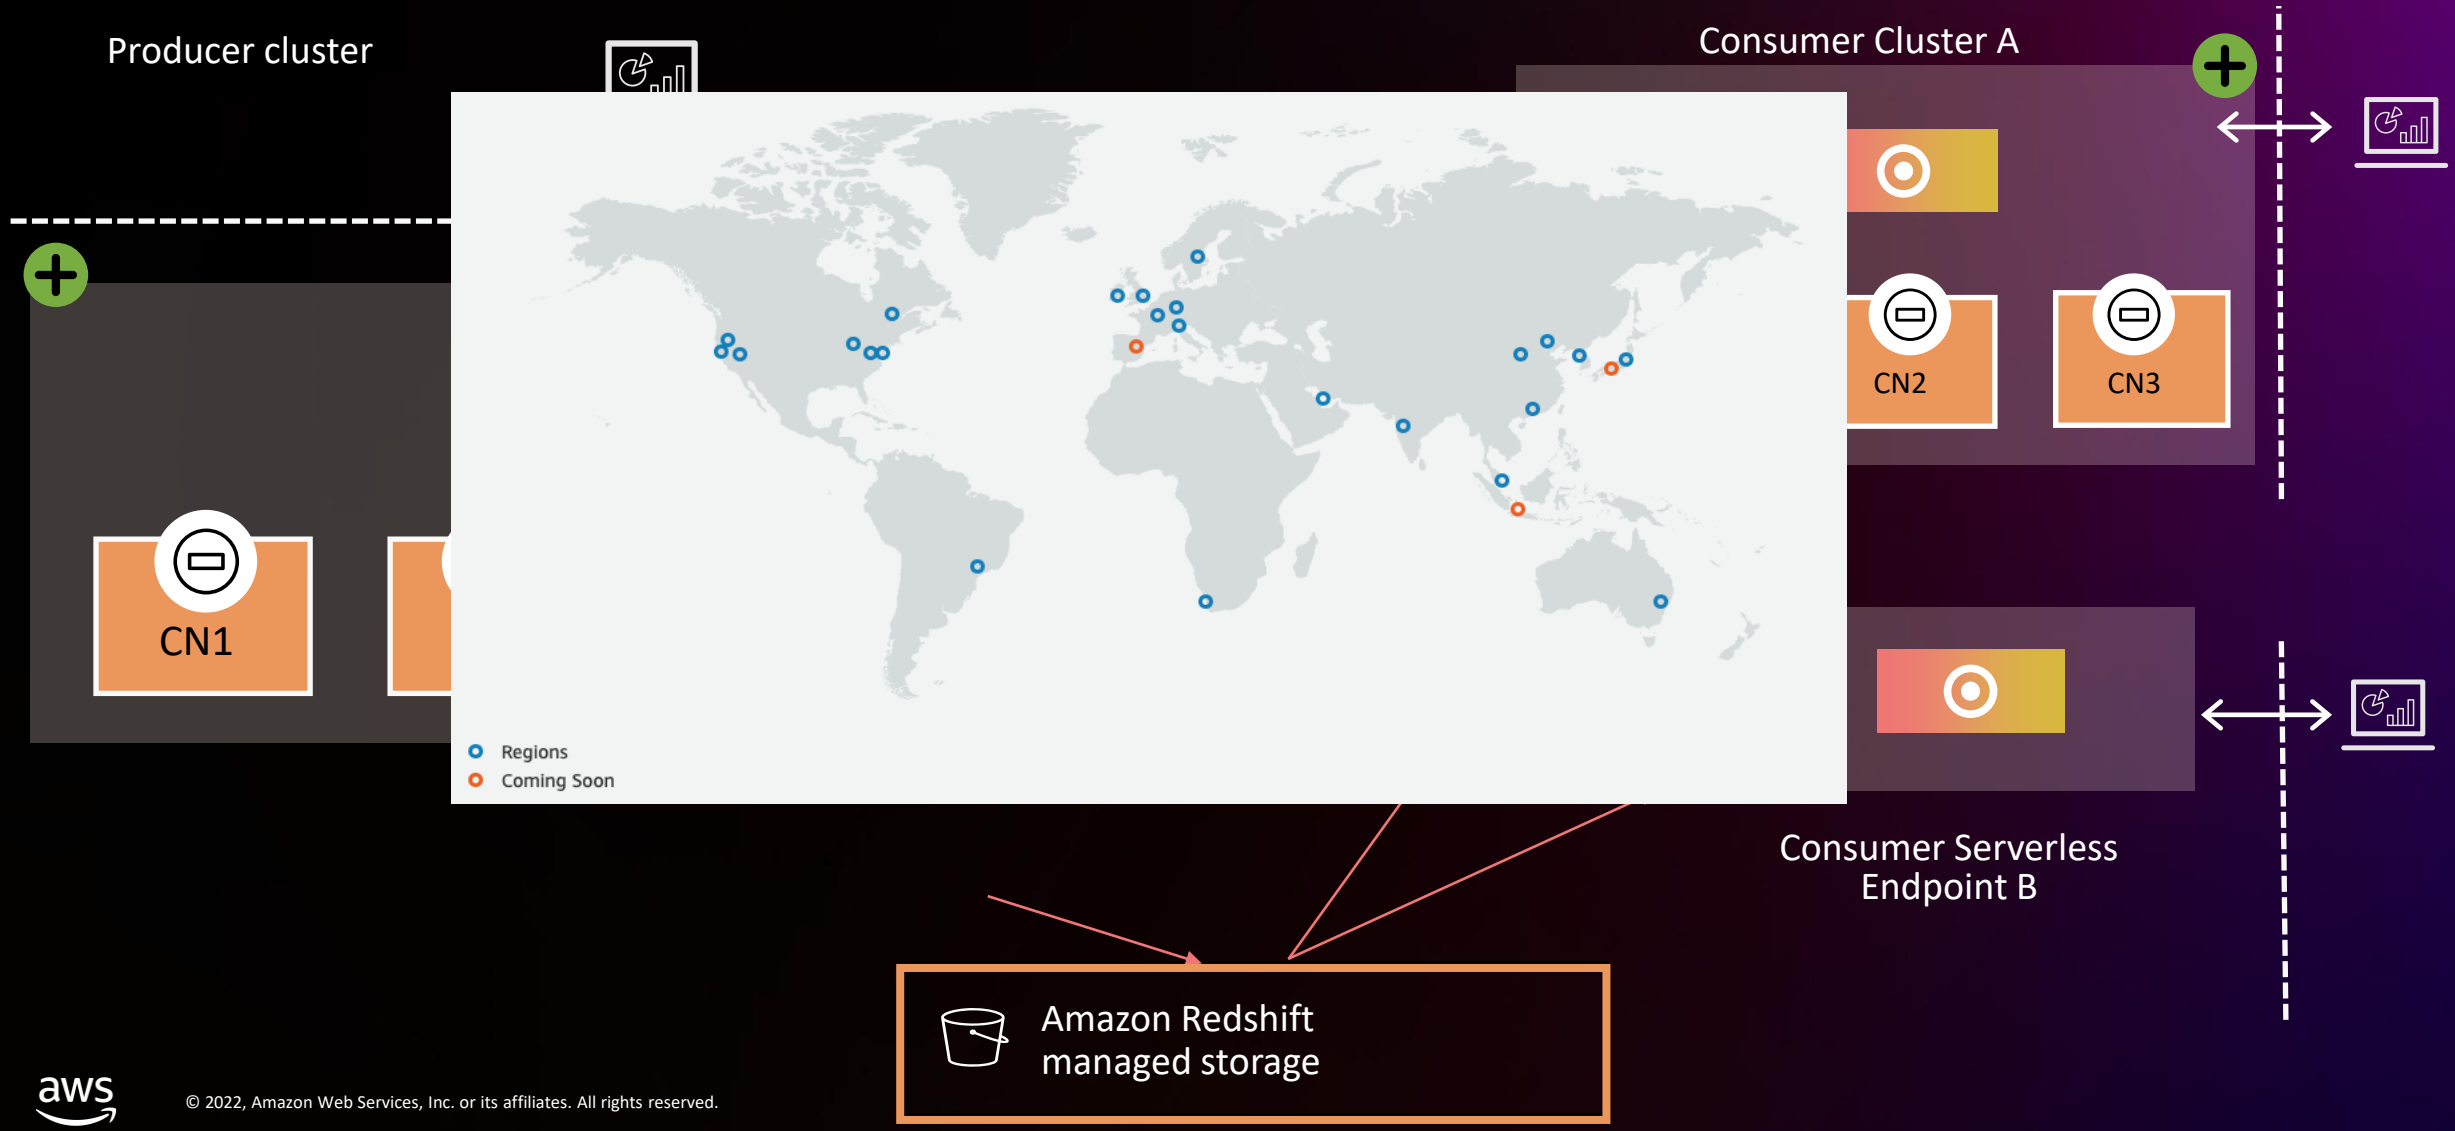
ISOLATE AND SCALE COMPUTE



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Data sharing with auto-scaling, across the globe

ISOLATE AND SCALE COMPUTE



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

- Best-in-class price performance at any scale
- Looking Amazon Redshift under the hood
- Storage and compute elasticity
- Autonomics and Serverless
- Q&A



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Auto Maintenance Operations

ML-BASED OPTIMIZATIONS TO GET STARTED EASILY AND GET THE FASTEST PERFORMANCE QUICKLY

At times of no or low activity, kick off maintenance operations

Automatic Table Optimization

- Self-tuning
- Optimize the physical design without DBA intervention
- Apply sort and distribution keys

Auto-Analyze

- Start with the tables that are more outdated and used more

Auto-Vacuum {Delete, Full}

- Priority-based and incremental
- Start with the most frequently used tables
- Perform vacuum if many rows have been added or deleted

Auto-Materialized View Refresh

- Refresh the MVs that are more likely to be used next

. . . And keep on adding more optimizations



Automatic Analyze



Automatic Vacuum Delete



Distribution/Sort Key Advisors



Automatic Table Distribution Keys



Automatic Table Sort Keys



Automatic Table Column Encoding



Auto Workload Manager



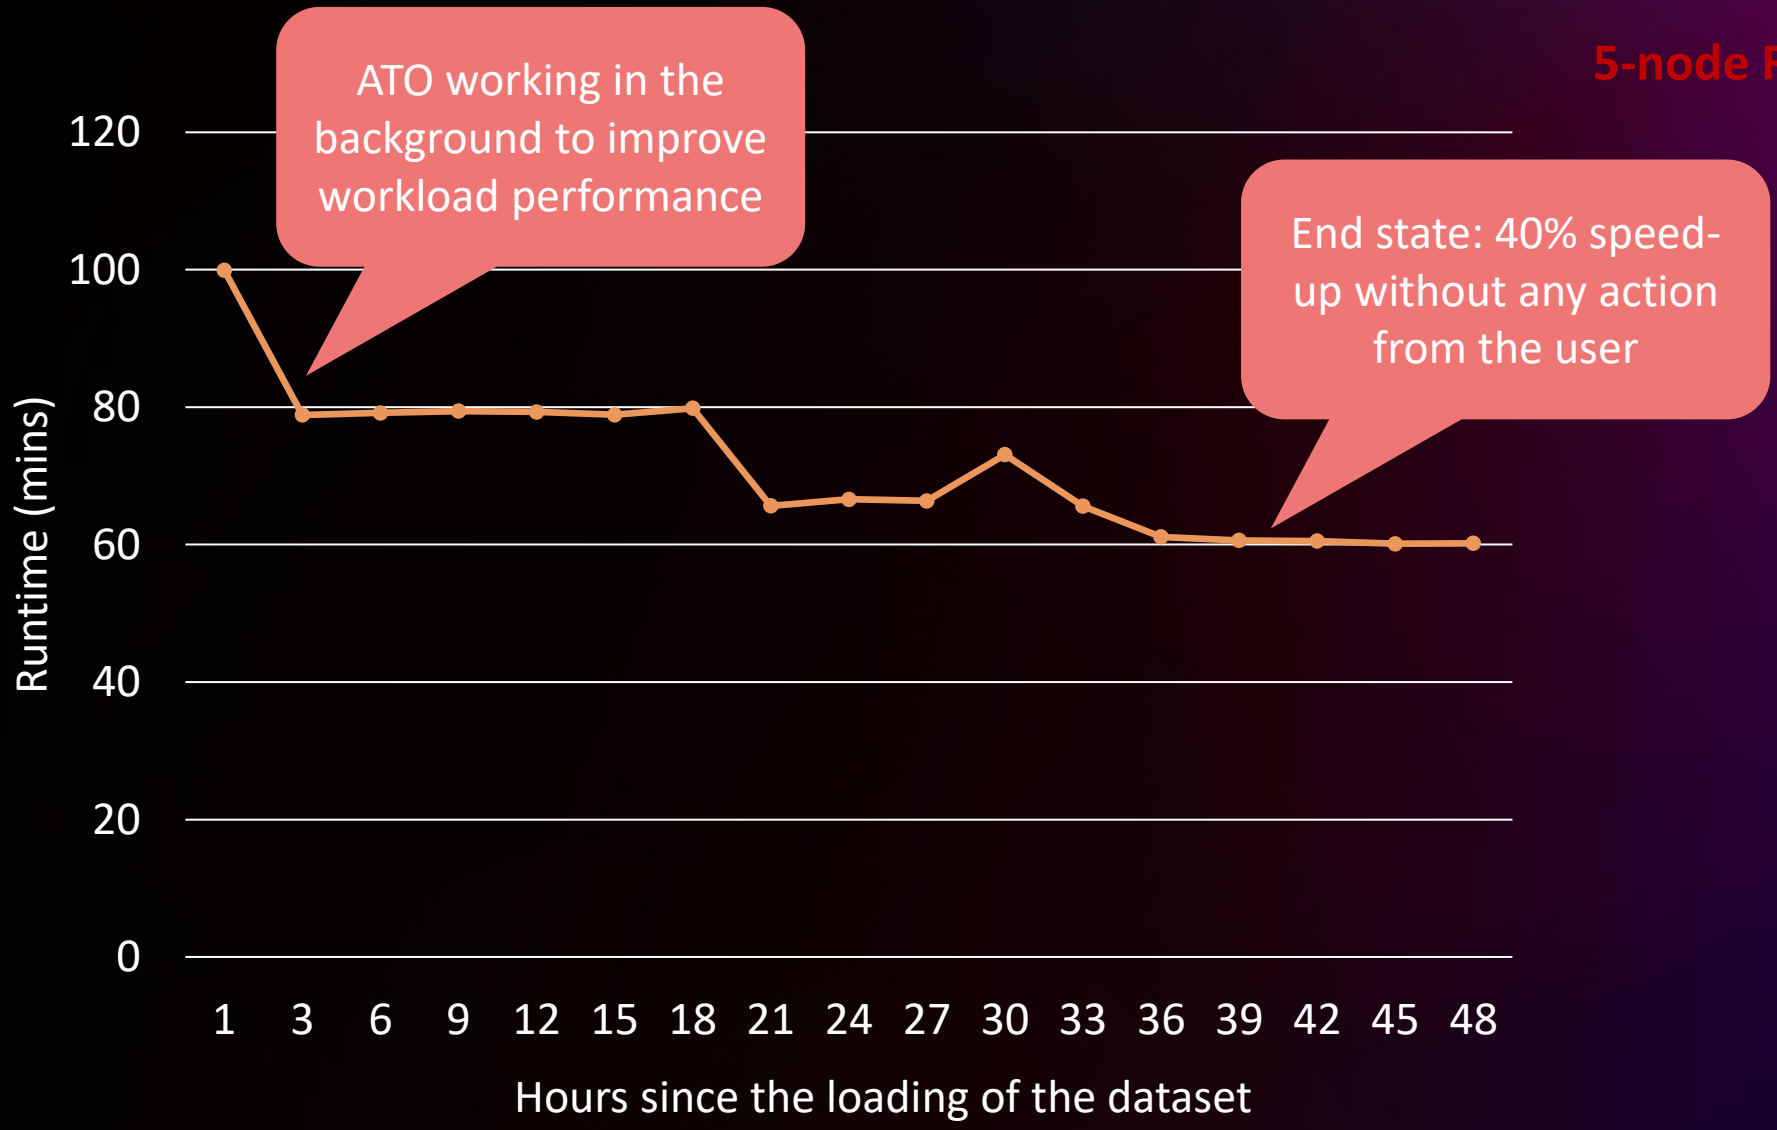
Automatic Table Sort



Auto refresh & rewrite of MVs

Automatic Table Optimization

Cloud DW derived from TPC-H 30TB 5-node RA3.16XL



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Automated materialized views

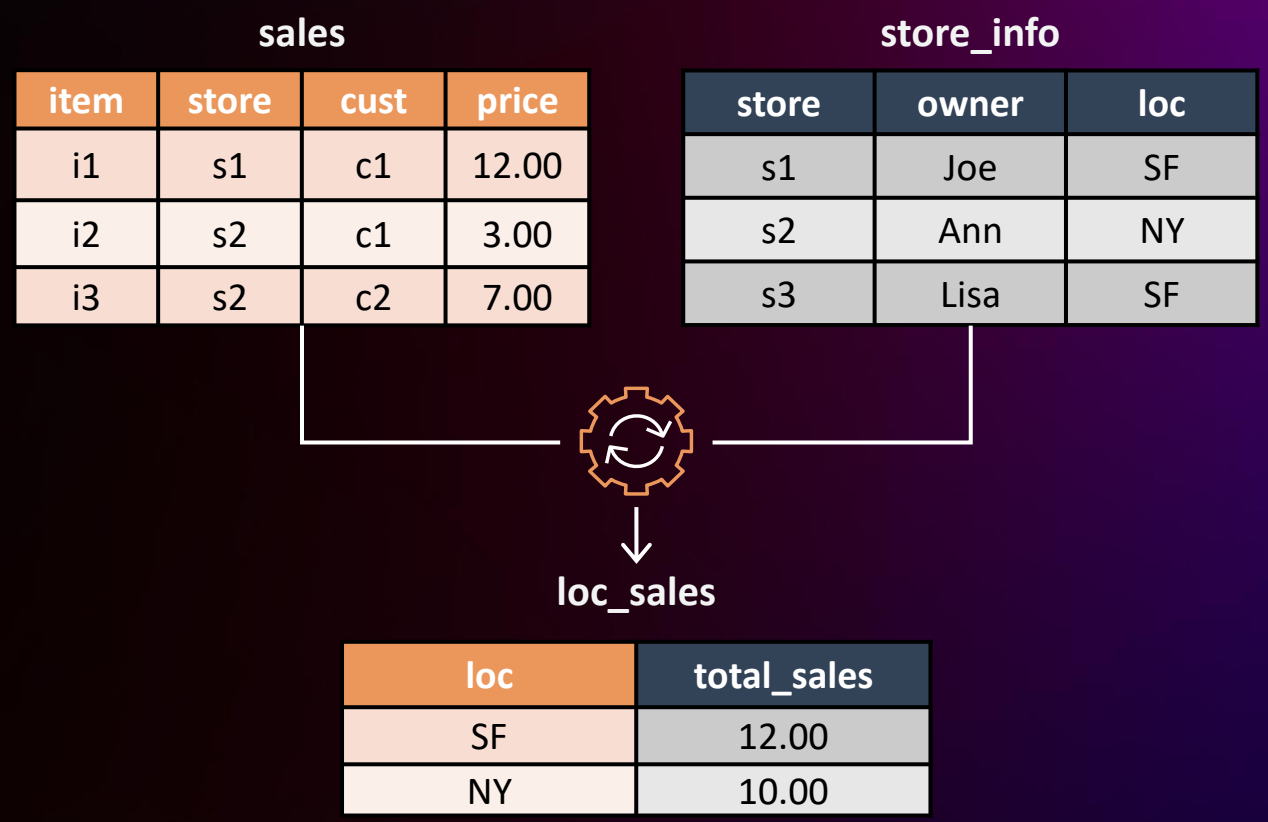
Automated creation and deletion of materialized views, with incremental refresh

Up to orders of magnitude query performance improvement

Continual monitoring to identify queries that will benefit from having a MV and automatically create

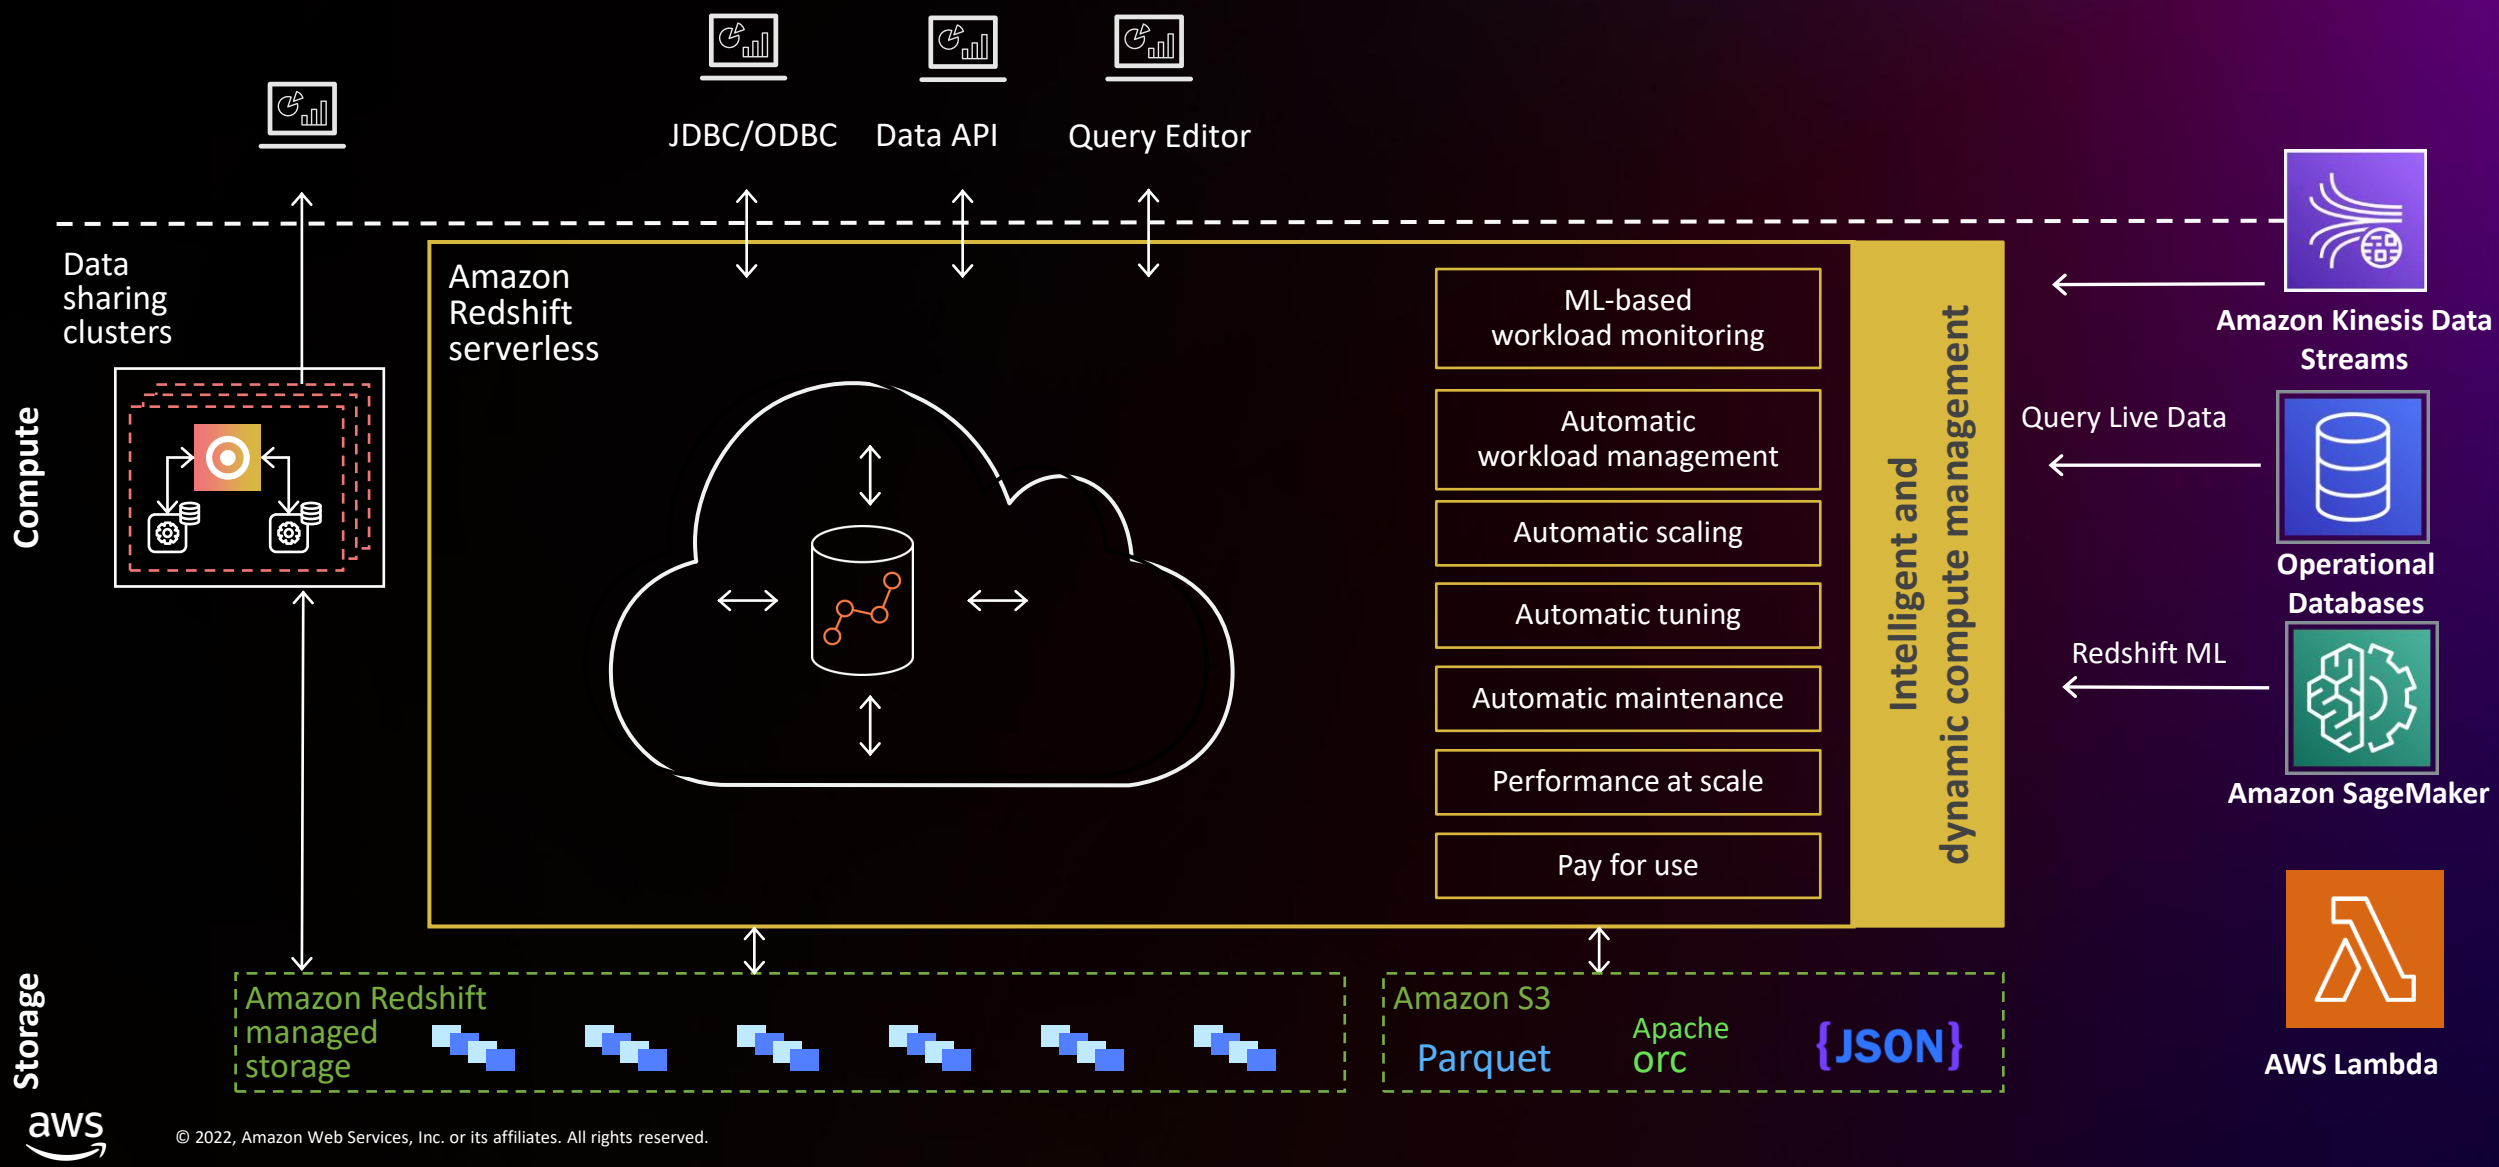
Automated query rewrite to leverage MV

Automatically deletes MVs that are no longer useful



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Intelligent and dynamic compute management evolution



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

The evolution of Amazon Redshift



- Tens of thousands of customers process exabytes of data daily with Amazon Redshift
- Continuous focus on performance and scalability
- Ability to elastically support tens of PBs of data and thousands of users
- Autonomics that make Amazon Redshift easy to use
- Serverless experience with intelligent compute management

And it is still Day 1. Thank you!

Agenda

- Best-in-class price performance at any scale
- Looking Amazon Redshift under the hood
- Storage and compute elasticity
- Autonomics and Serverless
- Q&A



© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Thank you!

Ippokratis Pandis
ippo@amazon.com

Stefan Gromoll
sgromoll@amazon.com



Please complete the session survey in the **mobile app**

© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.