databases

Amazon DynamoDB

Deep dive on key features that drive business impact

Pete Naylor - DynamoDB Specialist SA - AWS
March 25, 2021

Topics of discussion

- Elasticity – throughput and storage
- Caching
- Backup/restore
- Export to Amazon S3
- PartiQL support (SQL statements)
- Multi-Region architectures
- Streaming of change data
- Materializing views in additional indexes
- Building on DynamoDB

aws databases

Elasticity – throughput and storage

aws databases

Throughput modes

Provisioned capacity mode (since service GA in 2012)

- Storage scales automatically.
- Provide throughput requirement as signal to the service.
- Auto scaling read and write capacity since June 2017.

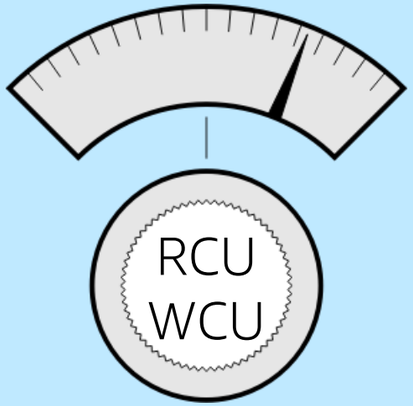

On-demand capacity mode (added November 2018)

- Storage scales automatically.
- Throughput scales instantly within 2x prior maximum consumption level – pay per request.

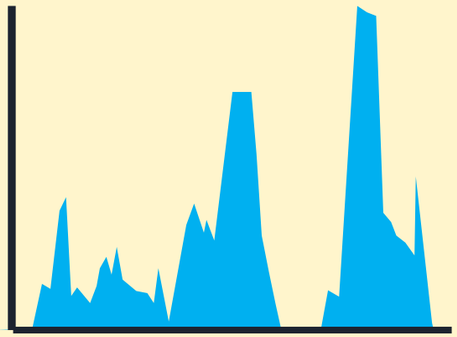

aws databases

Partitions split to double your capability

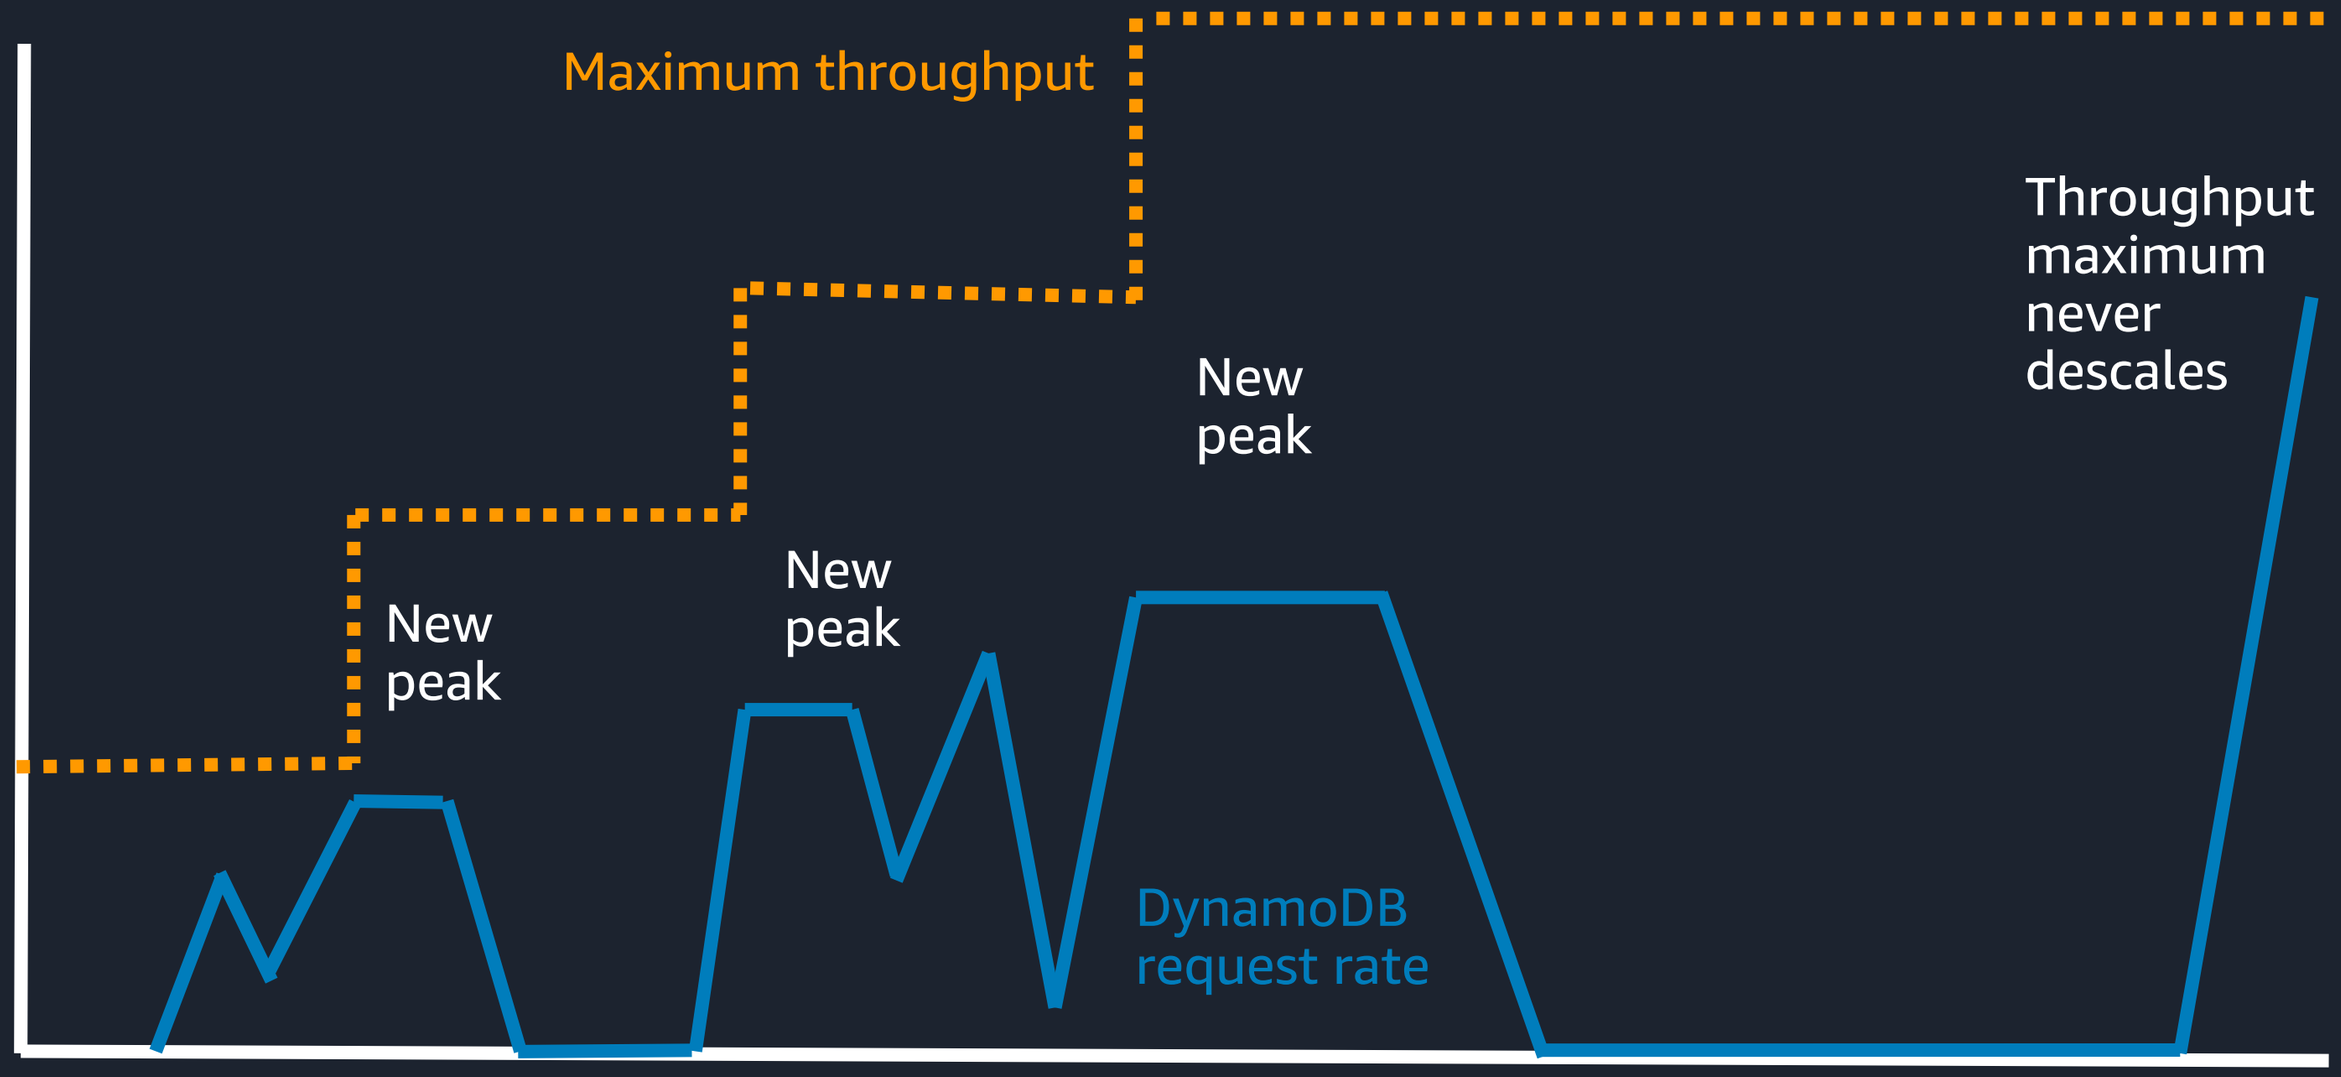

aws databases

Choosing the right capacity mode

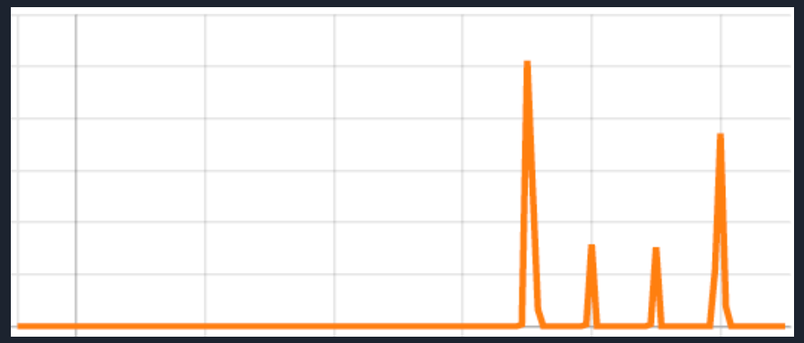

Cost versus price

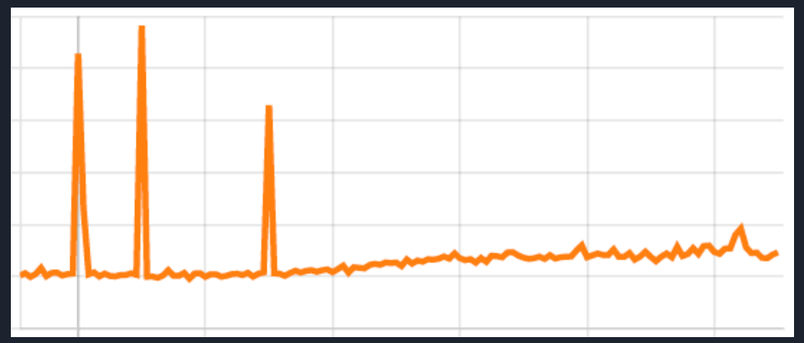

Engineering time and opportunity cost

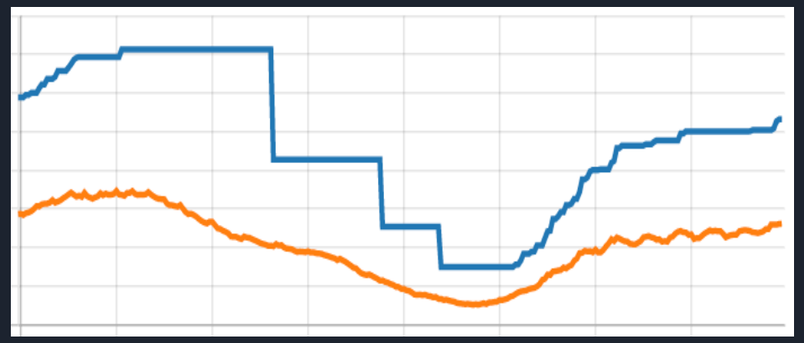

Operational load and risk

Resource lifecycle and utilization

aws databases

Caching

aws databases

DynamoDB Accelerator (DAX)

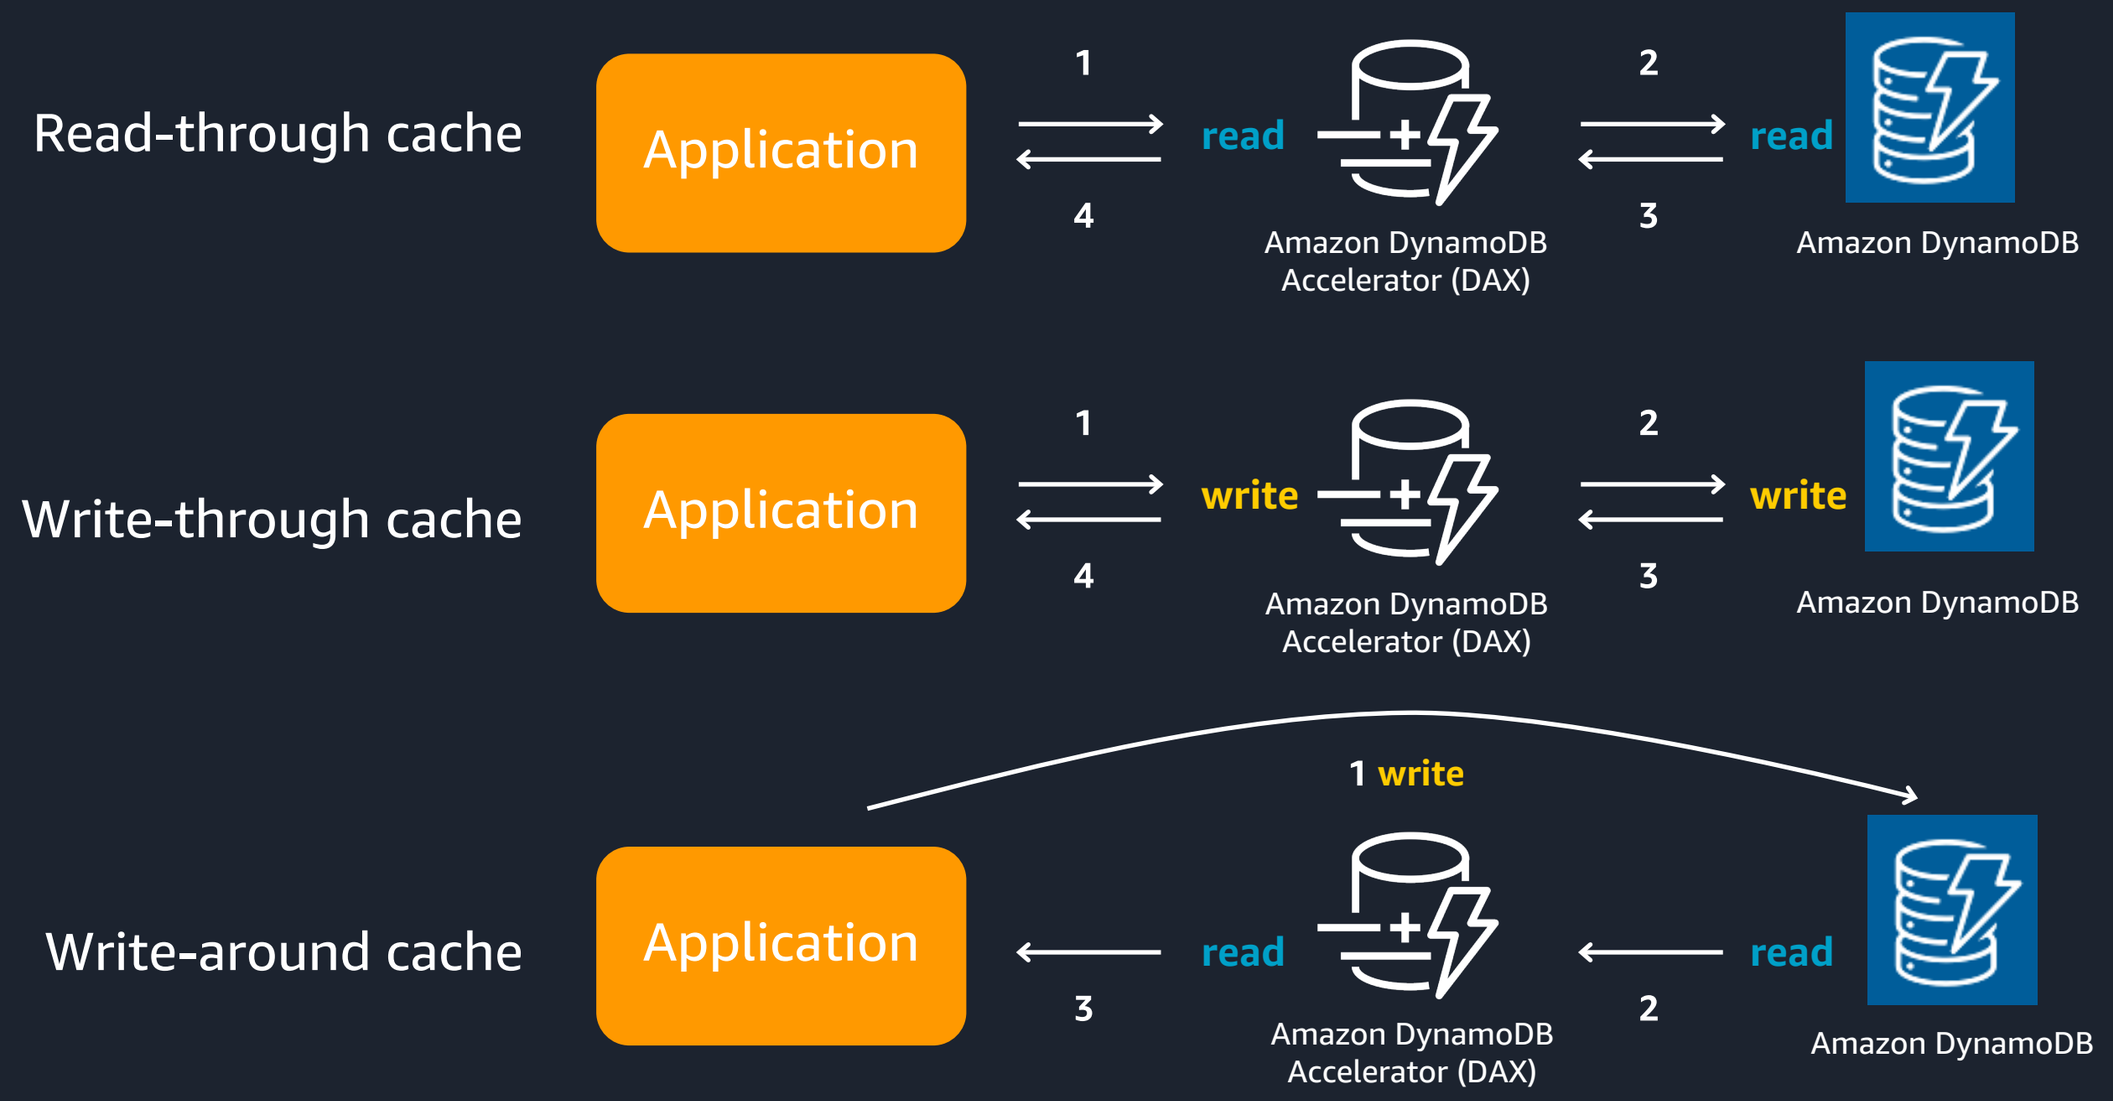

aws databases

DAX: Reduced read unit consumption for DynamoDB

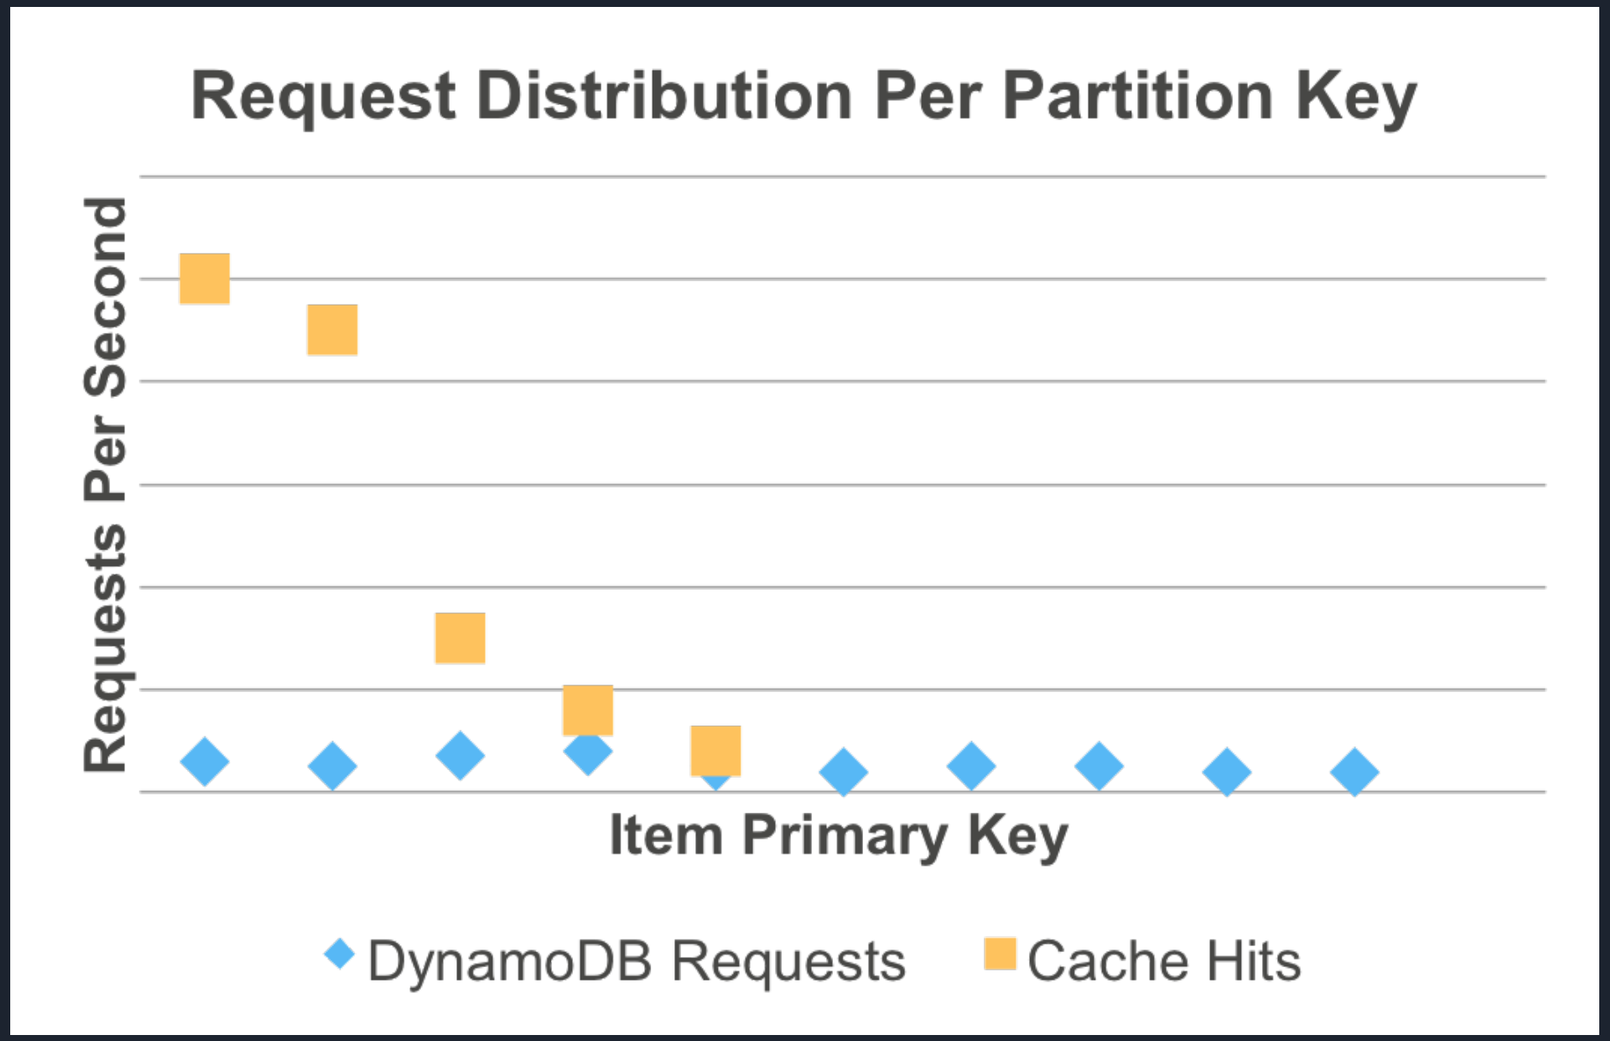

aws databases

DAX: Business benefits

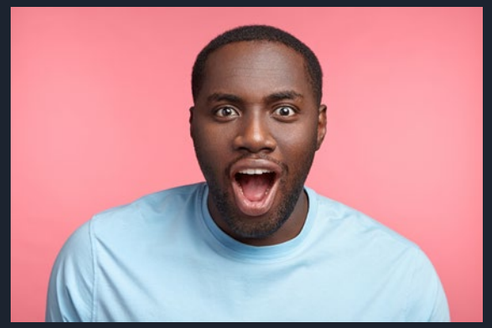

Lower latency – more delightful customer experience

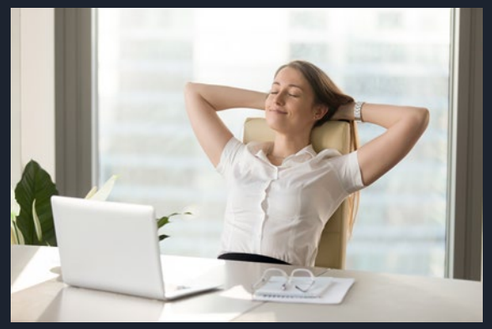

Smoothing of unanticipated skew in data pressure

Less read units consumed – improved cost efficiency

aws databases

Backup/restore

aws databases

Point-in-time recovery (PITR): properties

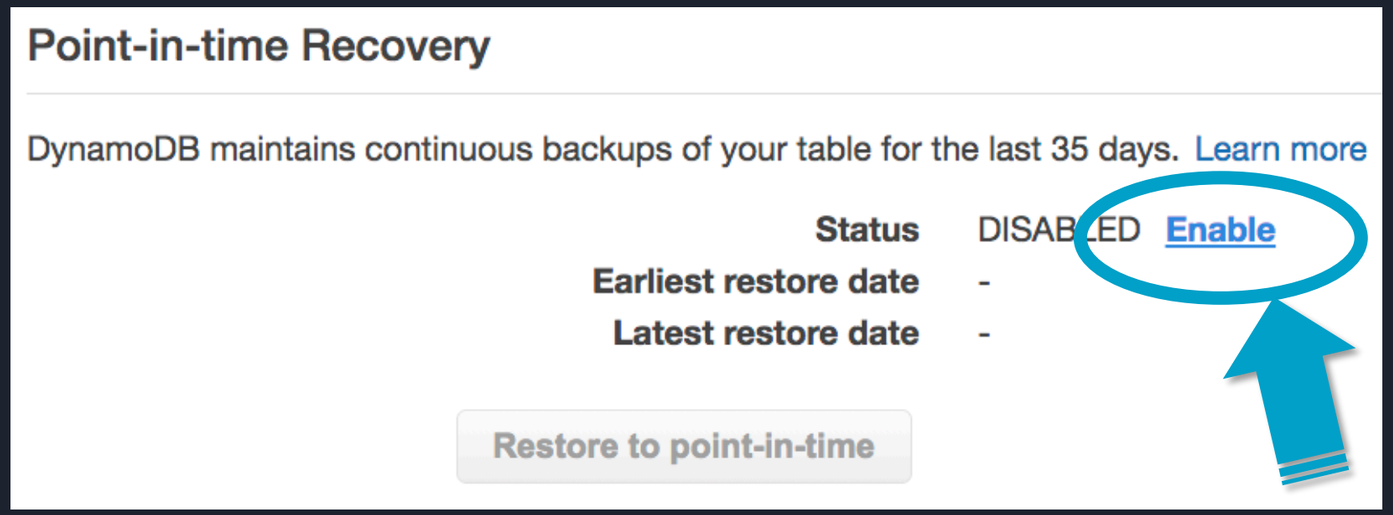

- 35-day rolling window of protection
- Restore to a new table with second granularity
- No impact to ongoing operations
- Covers table deletion and corruption risks

aws databases

Point-in-time recovery (PITR): restore scenarios

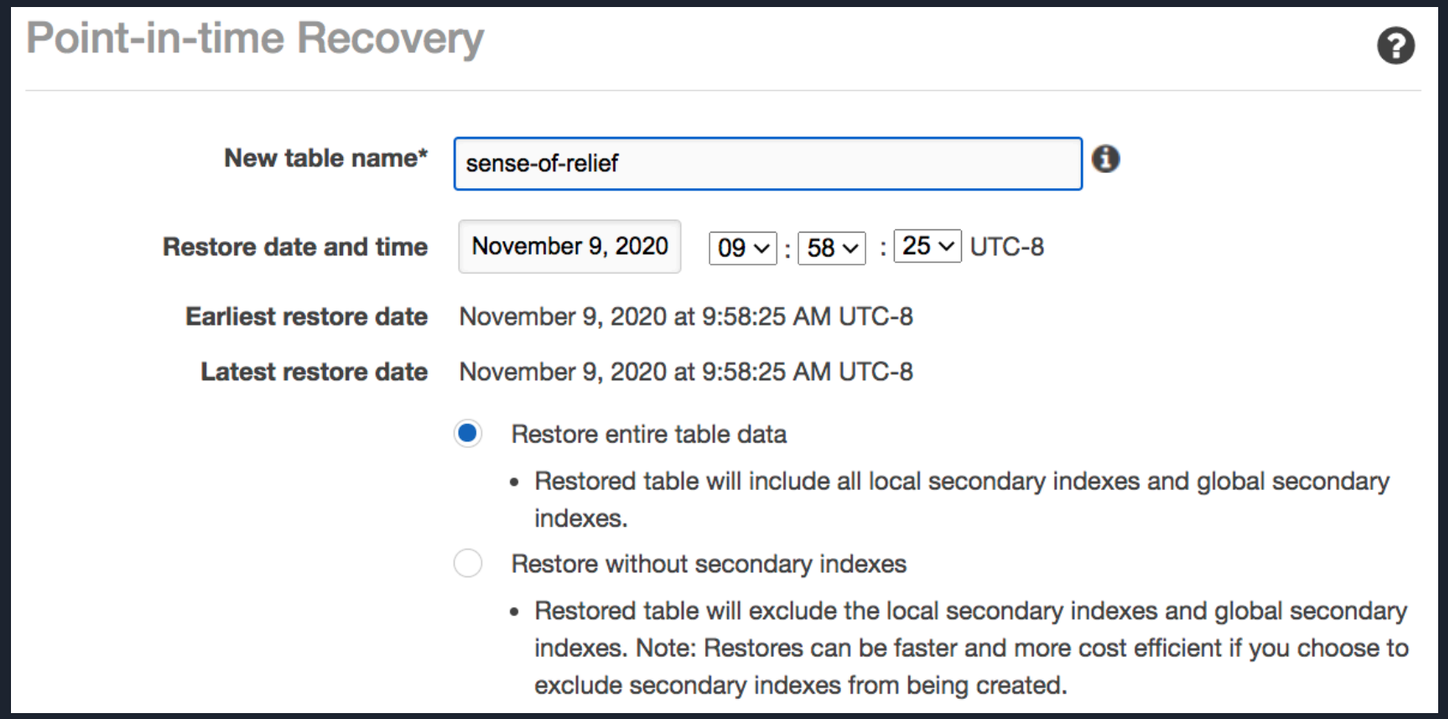

- Accidental (or malicious) table delete or data change
- Restore from known-good timestamp to original table name
- Restore from known-good timestamp to a new table
- Comparison and selective reversion of any unwanted changes?

aws databases

Point-in-time recovery (PITR): peace of mind

- DynamoDB often stores critical business data
- Business continuity and disaster recover are hard and often not fully considered or tested
- Having PITR as a continuous backup is an easy first step for production tables that will give you a starting point for recovery

aws databases

Export to Amazon S3

aws databases

Exporting table data to S3

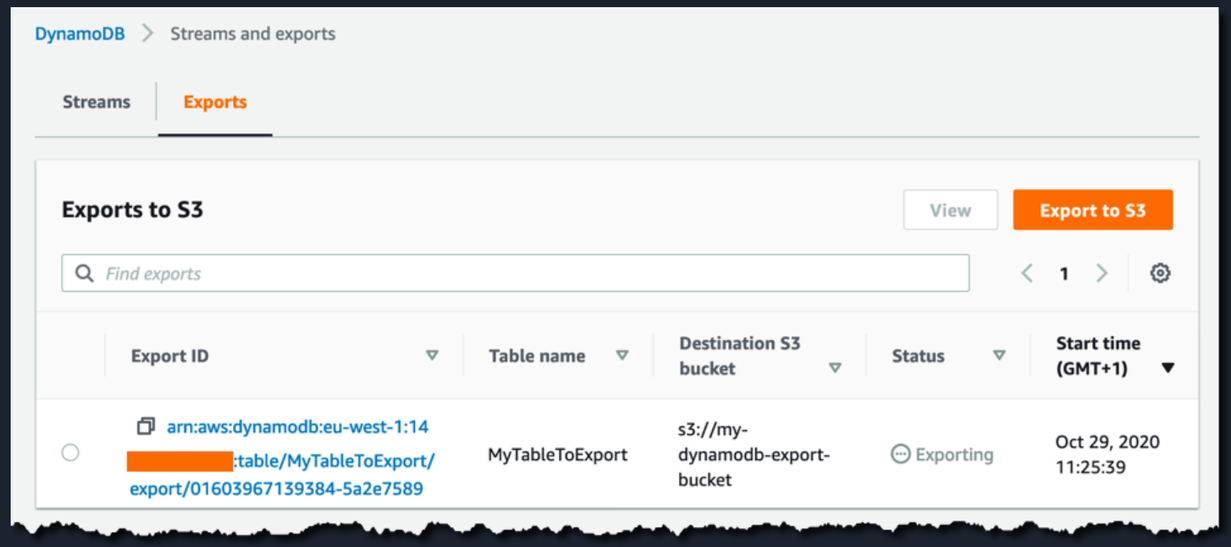

You now can make a recovery from your PITR backup into S3 – simple API or use the console

- Consumes no read units – no load on your table
- Output to your S3 bucket in (compressed) DynamoDB JSON format or Amazon Ion format
- Restore to any second in your PITR window

aws databases

PartiQL support (SQL statements)

aws databases

SQL-compatible statements for DynamoDB operations

DynamoDB > PartiQL editor

PartiQL editor

Operations performed using the PartiQL editor may incur charges. Learn more

Resources

Find tables

Tables (5) 1

- composite
- dynacount
- partiqltest
- simple
- usertable

Query 9 | Query 10

```
1 SELECT * FROM "composite" WHERE "pk" = 'abc123' AND begins_with("sk", 'inv#');
```

Run Clear

aws databases

SQL-compatible statements for DynamoDB operations

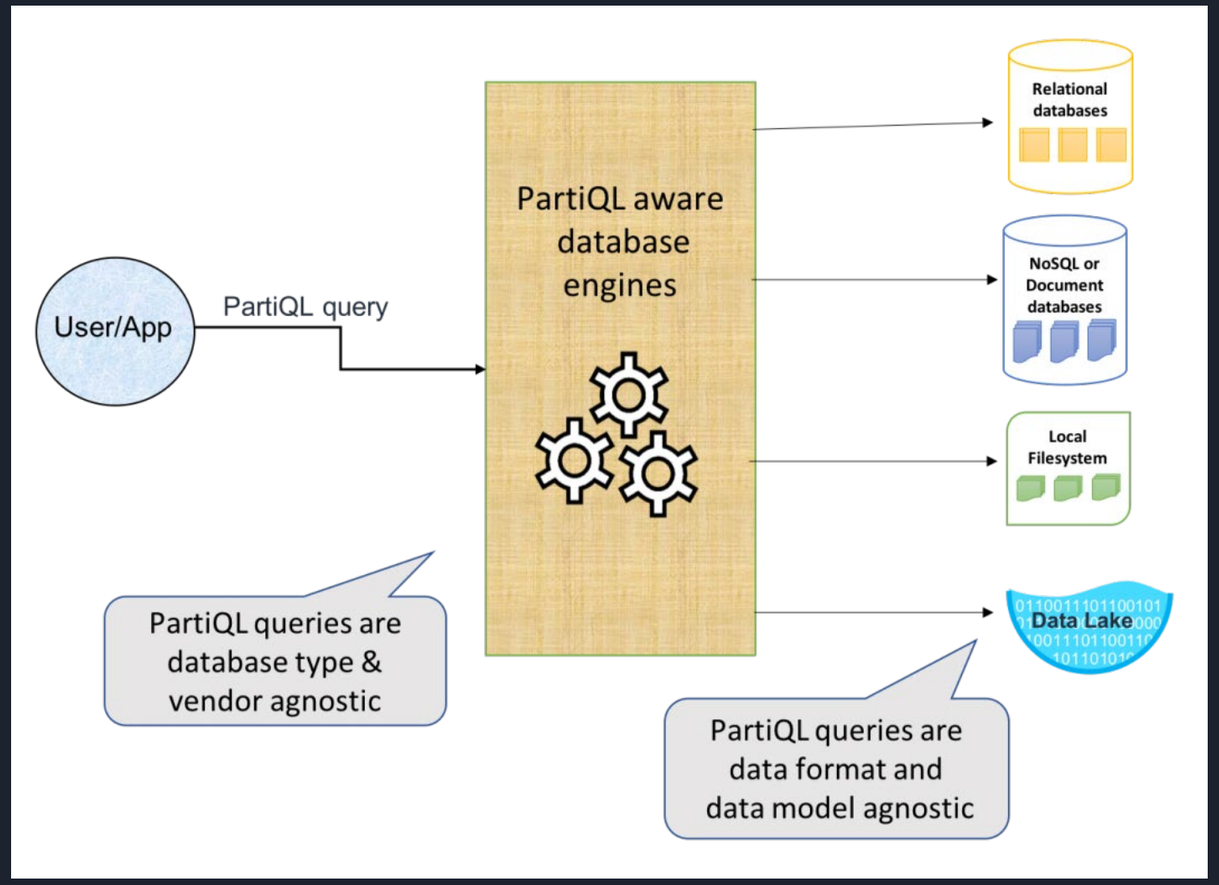

PartiQL is an open source specification for a SQL-compatible query language supporting any level of data structure. Three new DynamoDB API actions:

`ExecuteStatement`

`BatchExecuteStatement`

`ExecuteTransaction`

SQL verbs:

`INSERT, UPDATE, SELECT, DELETE`

aws databases

Multi-Region architectures

aws databases

Global tables: Multi-Region architectures

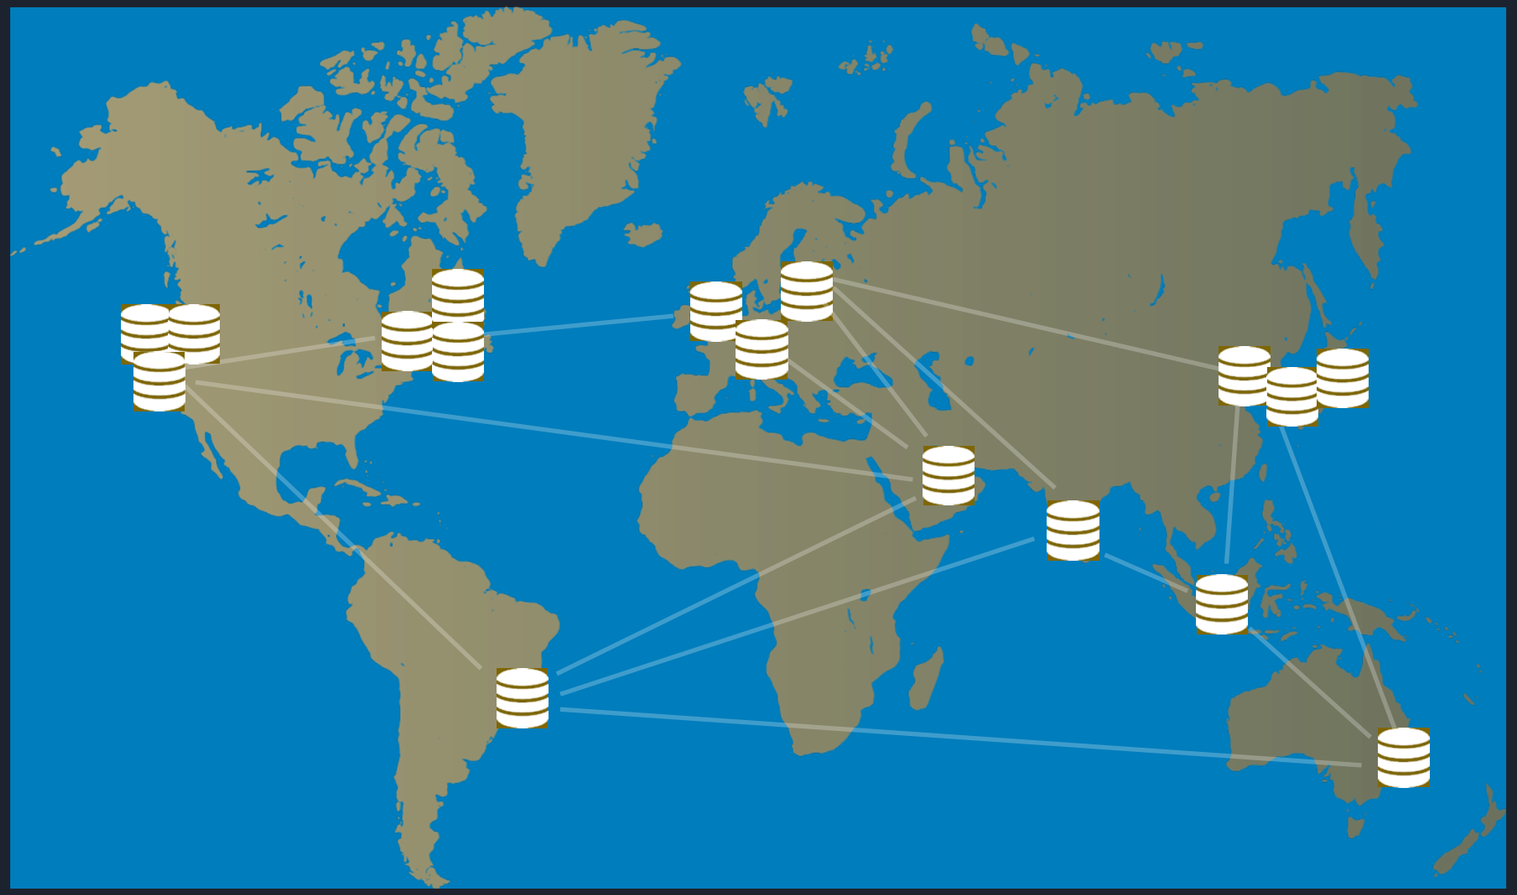

1. Geographically distributed customer base
   - Low-latency data operations
   - Meet legal and data regulatory compliance
2. Business continuity / disaster recovery

aws databases

Global tables: architecture beyond the data

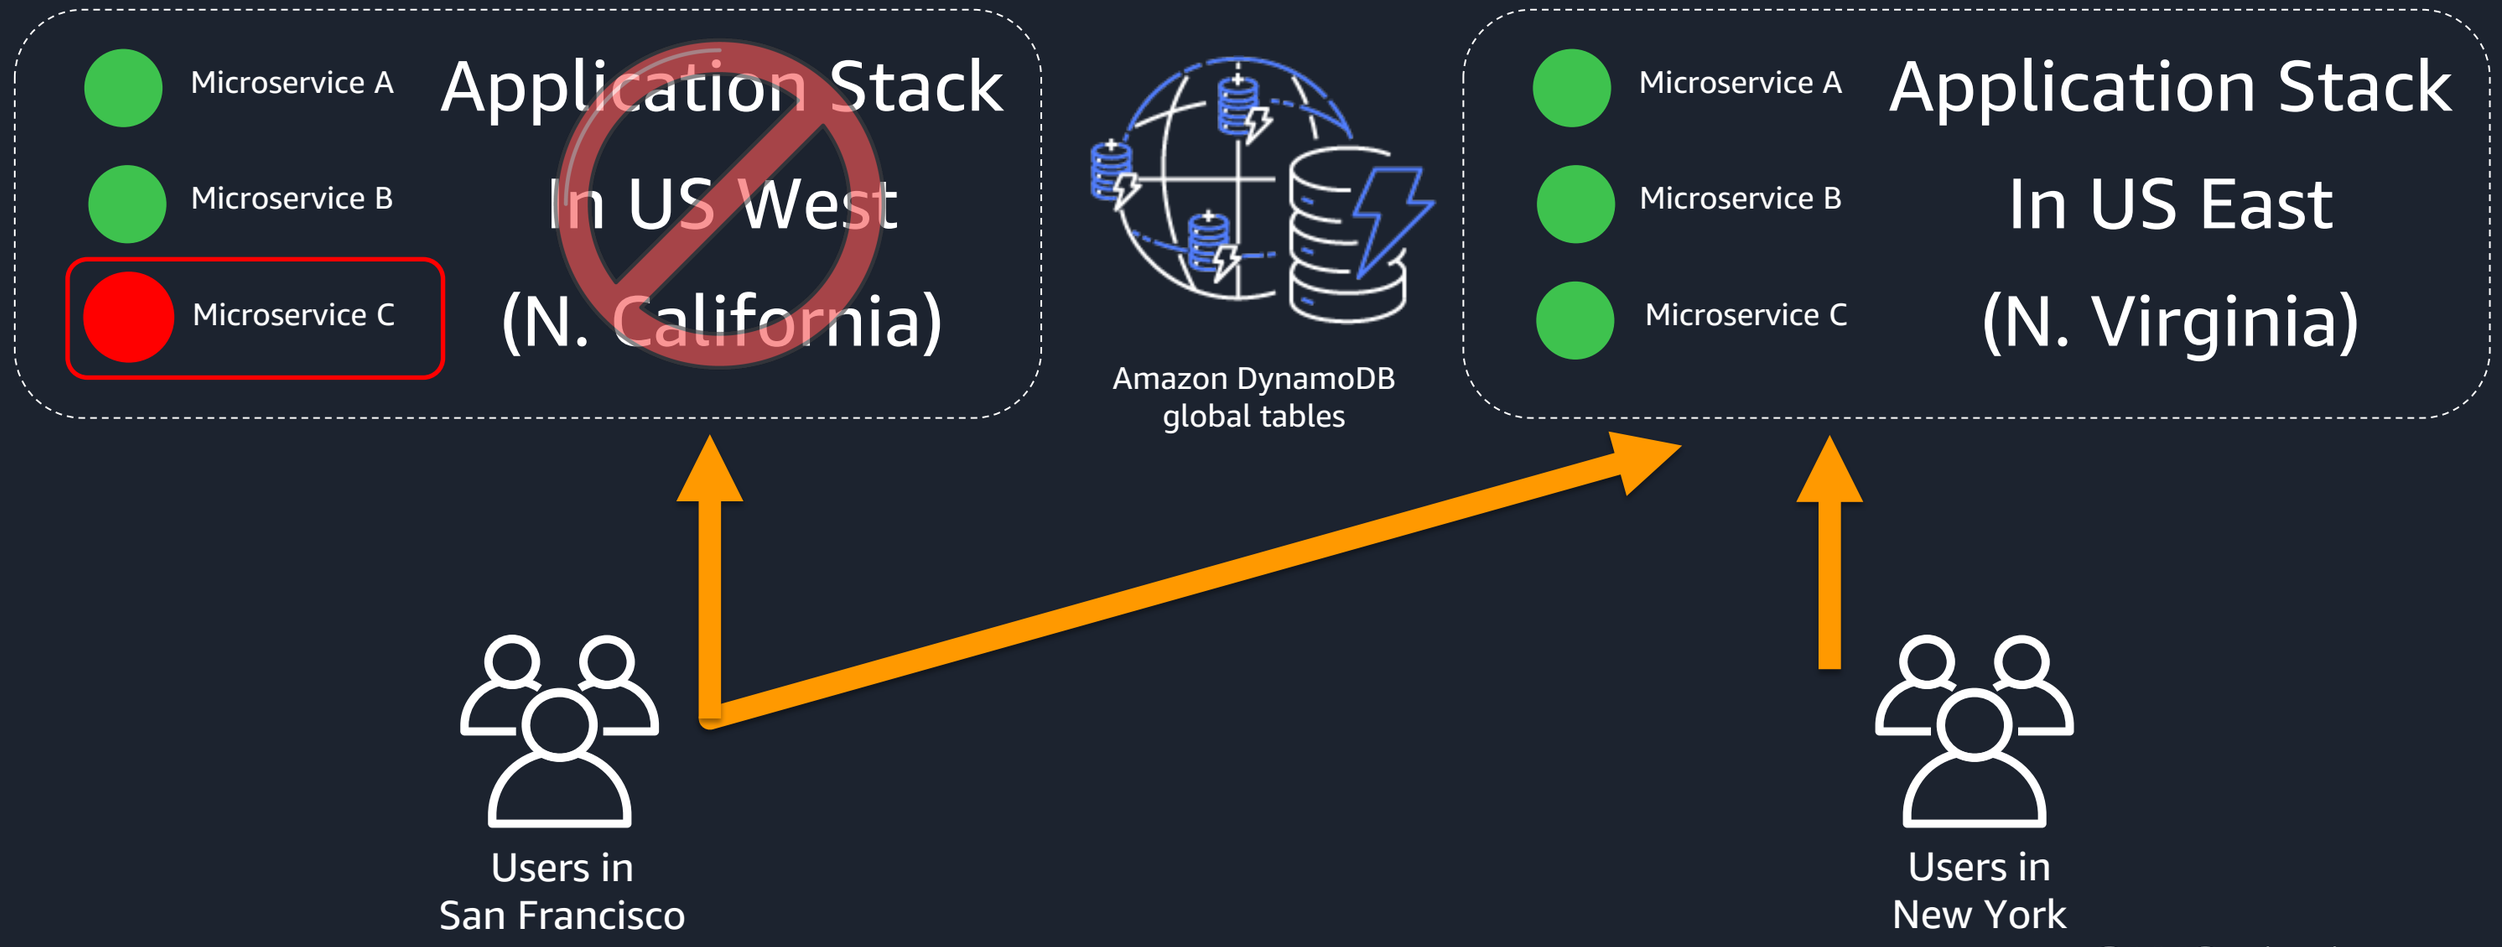

aws databases

Global tables: traffic distribution choices

- Reads and writes to a single Region (active-passive)
- Local reads, with writes to a single “primary” Region
- Local reads and writes (active-active) – consider item ”homing”
- Manual and/or automatic failover
- Amazon Route 53 health checks and routing policy – latency, geolocation
- Consider AWS Global Accelerator for fastest failover, lowest latency
- Amazon API Gateway can provide rate limiting and caching

aws databases

Streaming of change data

aws databases

DynamoDB and change data capture

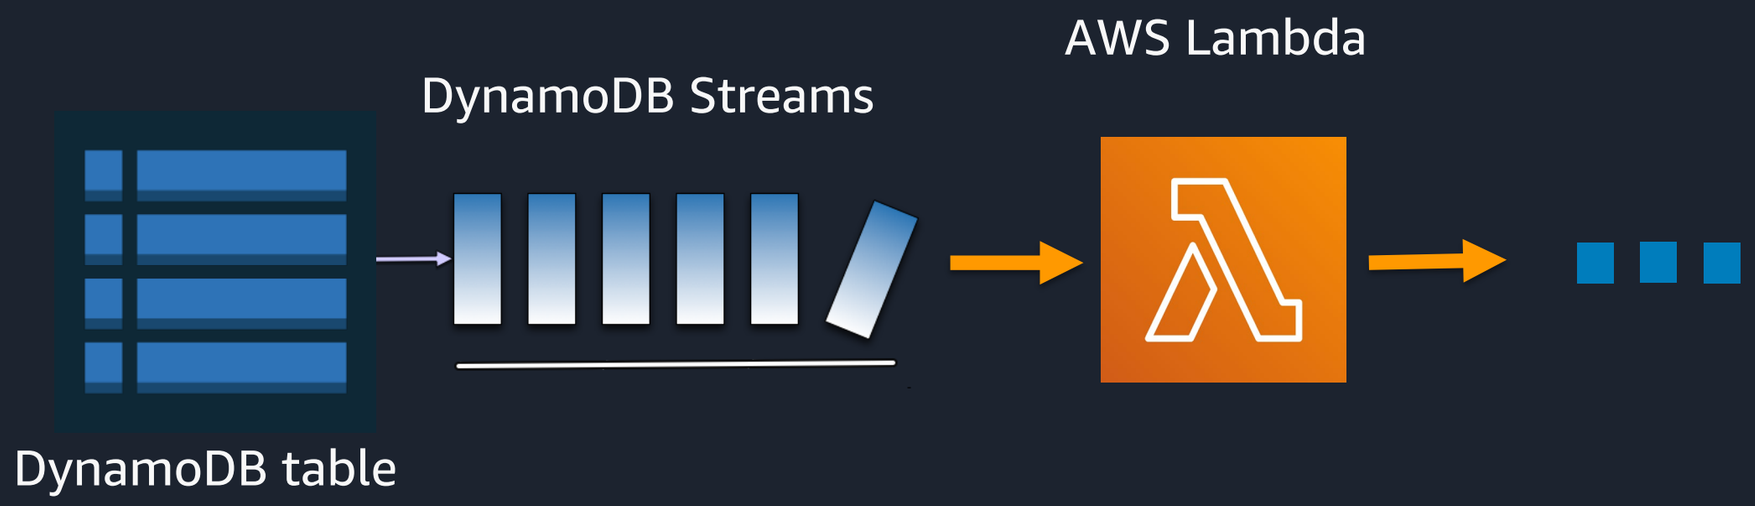

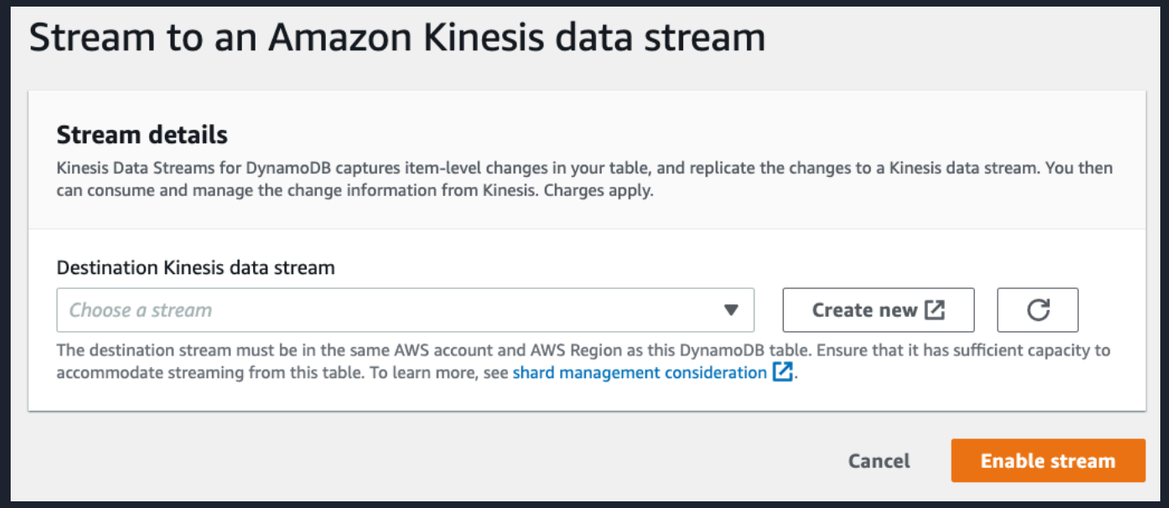

You now can capture change events direct to Amazon Kinesis Data Streams

Amazon Kinesis Data Streams

aws databases

Materializing views in additional indexes

aws databases

Propagating data to support different query needs

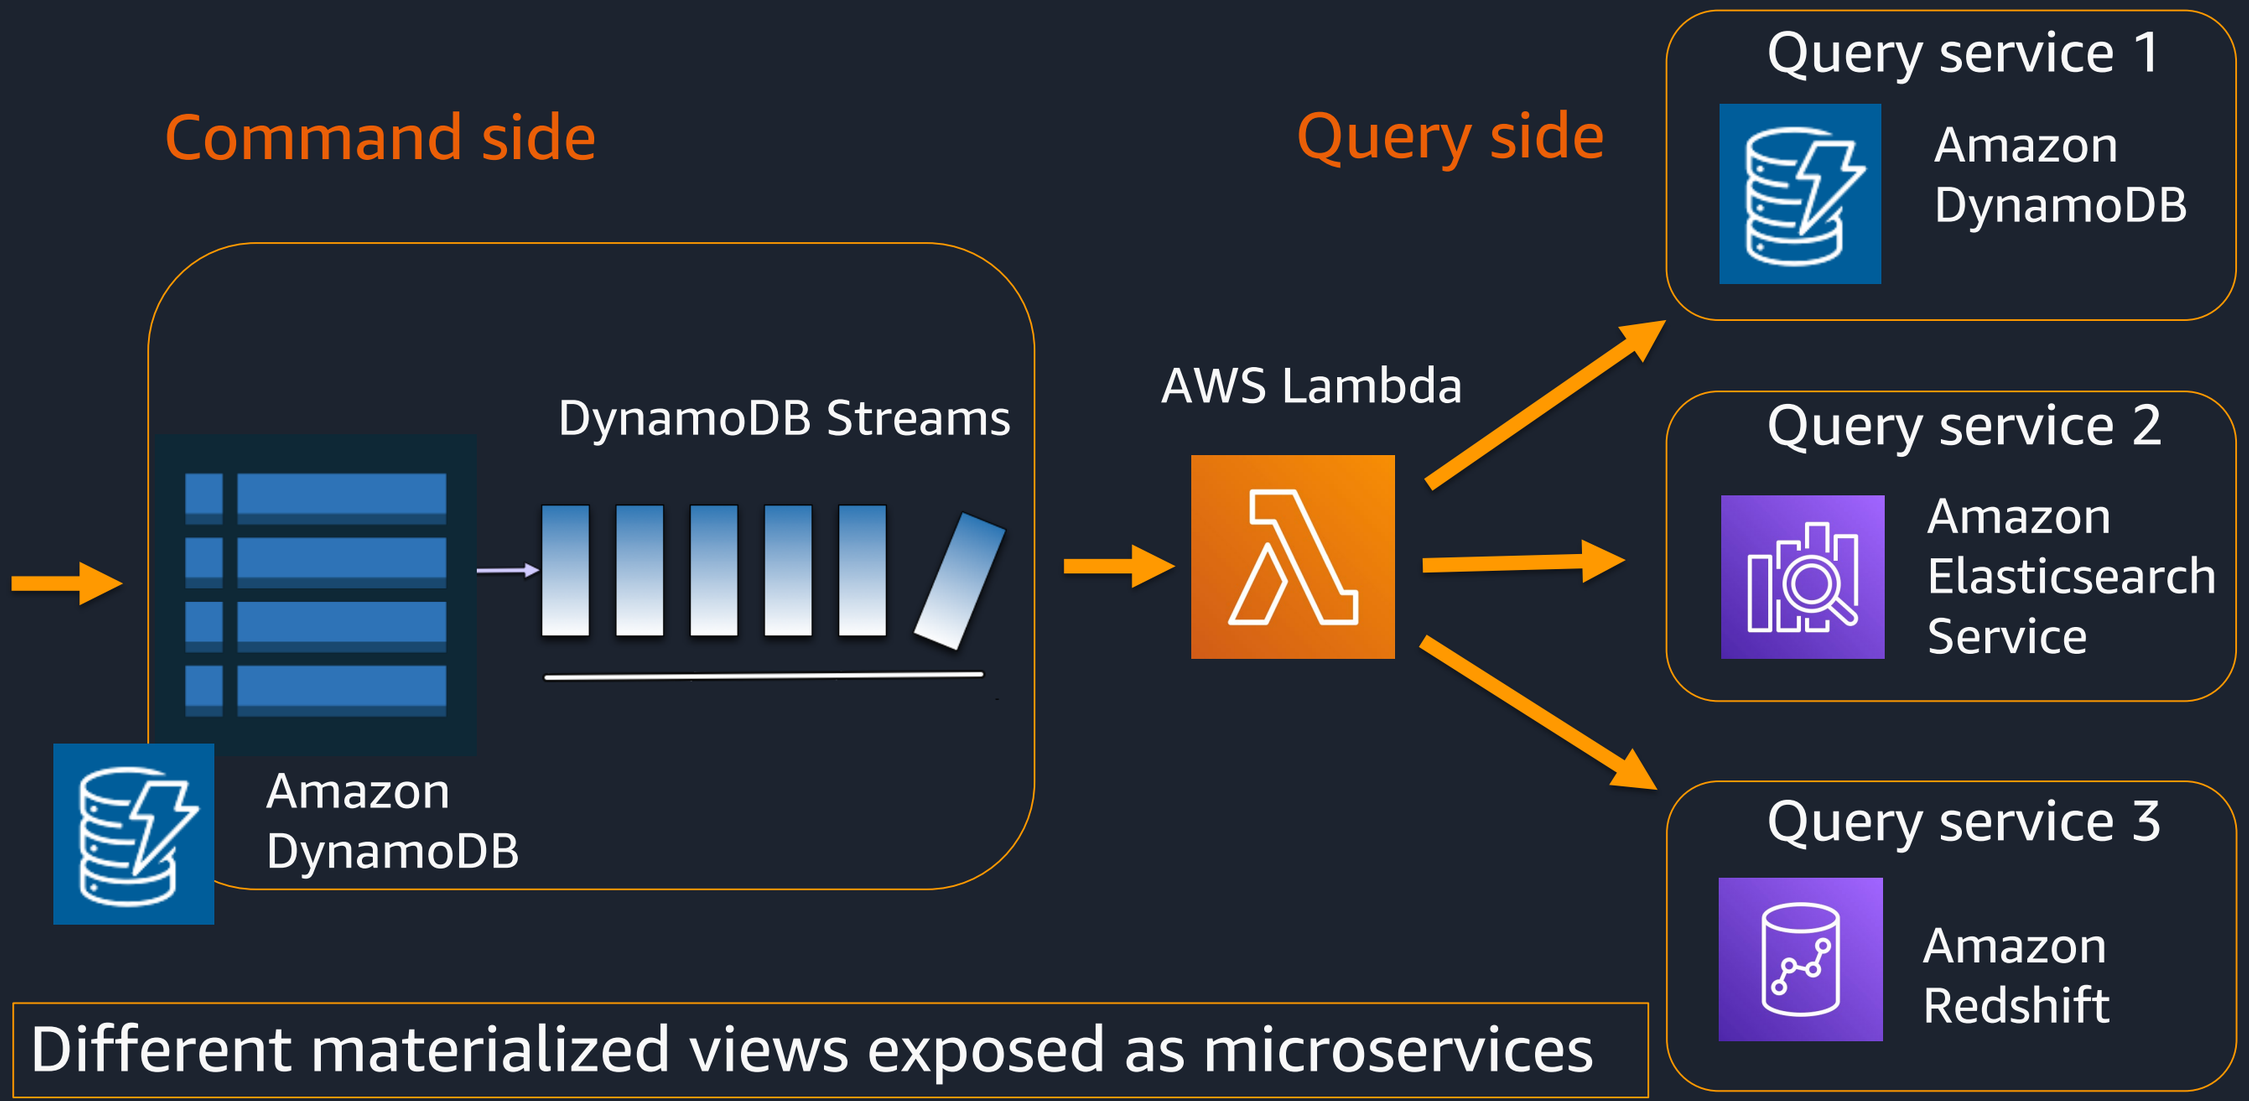

Different materialized views exposed as microservices

aws databases

Building on DynamoDB

aws databases

NoSQL Workbench: data modeling flow

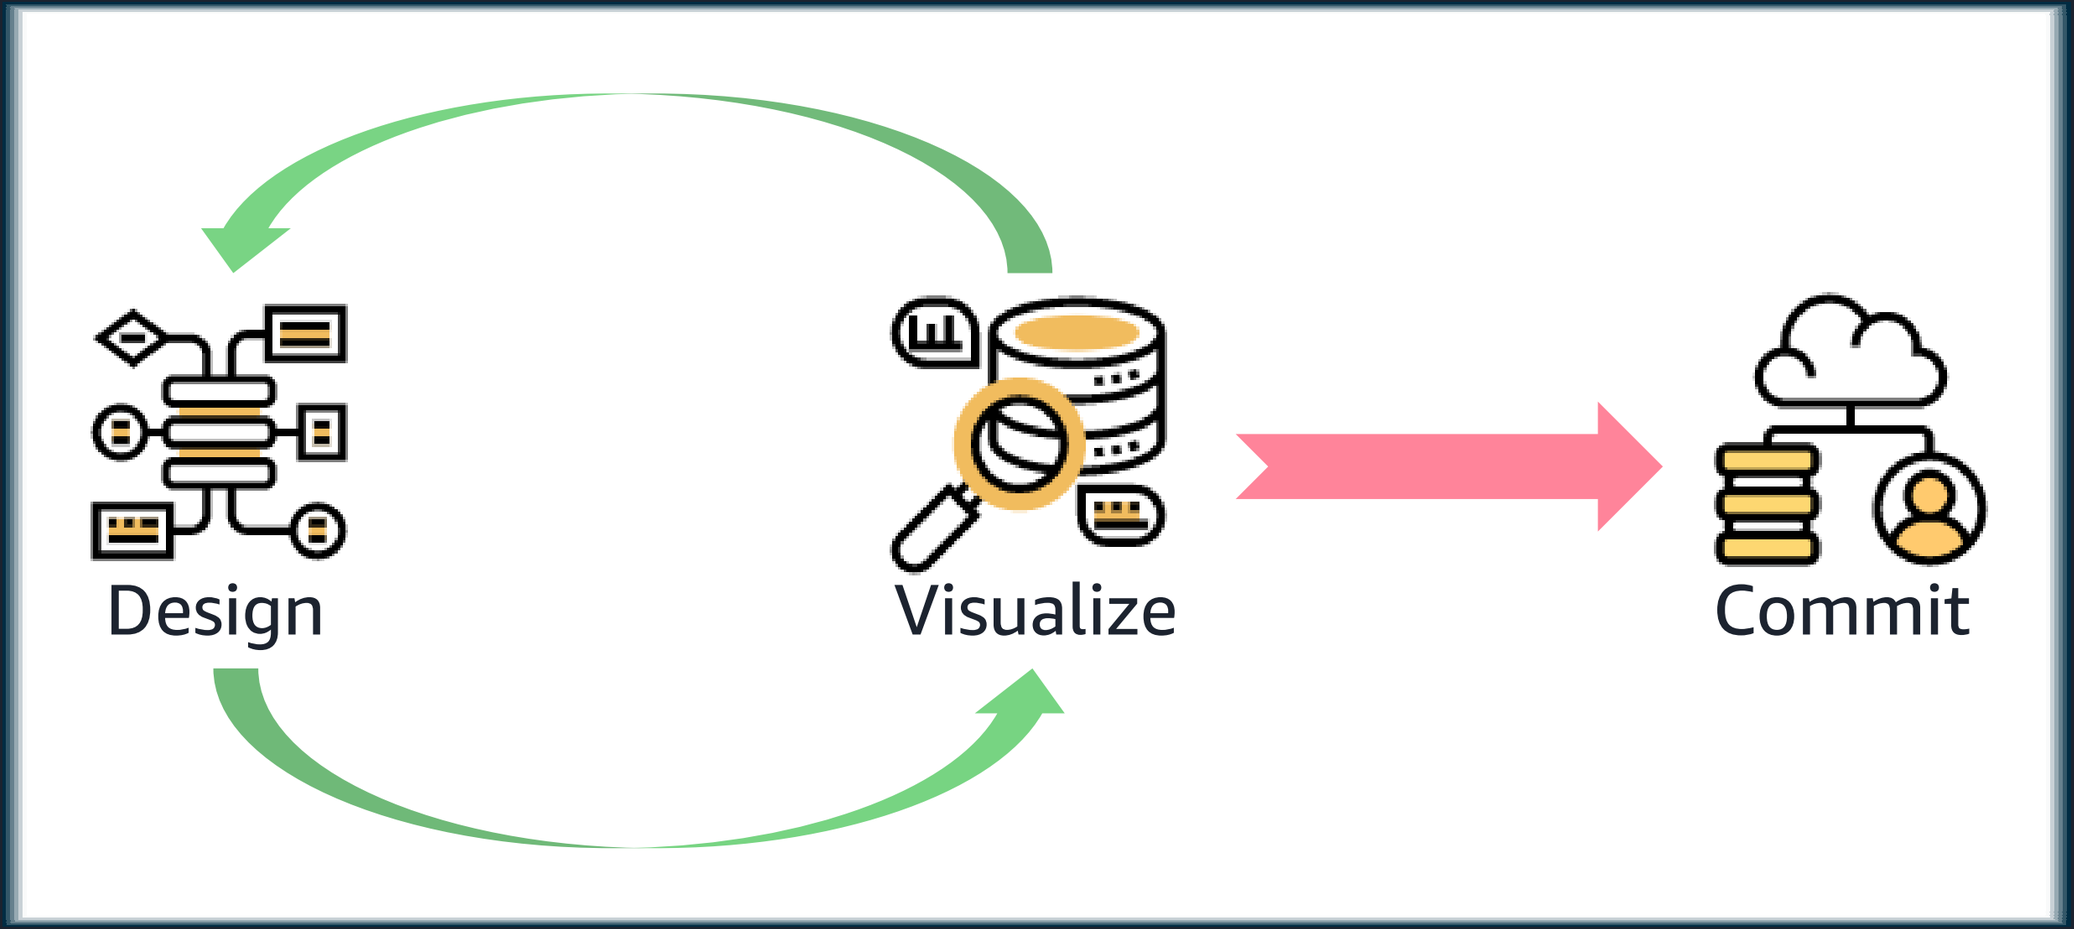

aws databases

NoSQL Workbench: collaborate and document

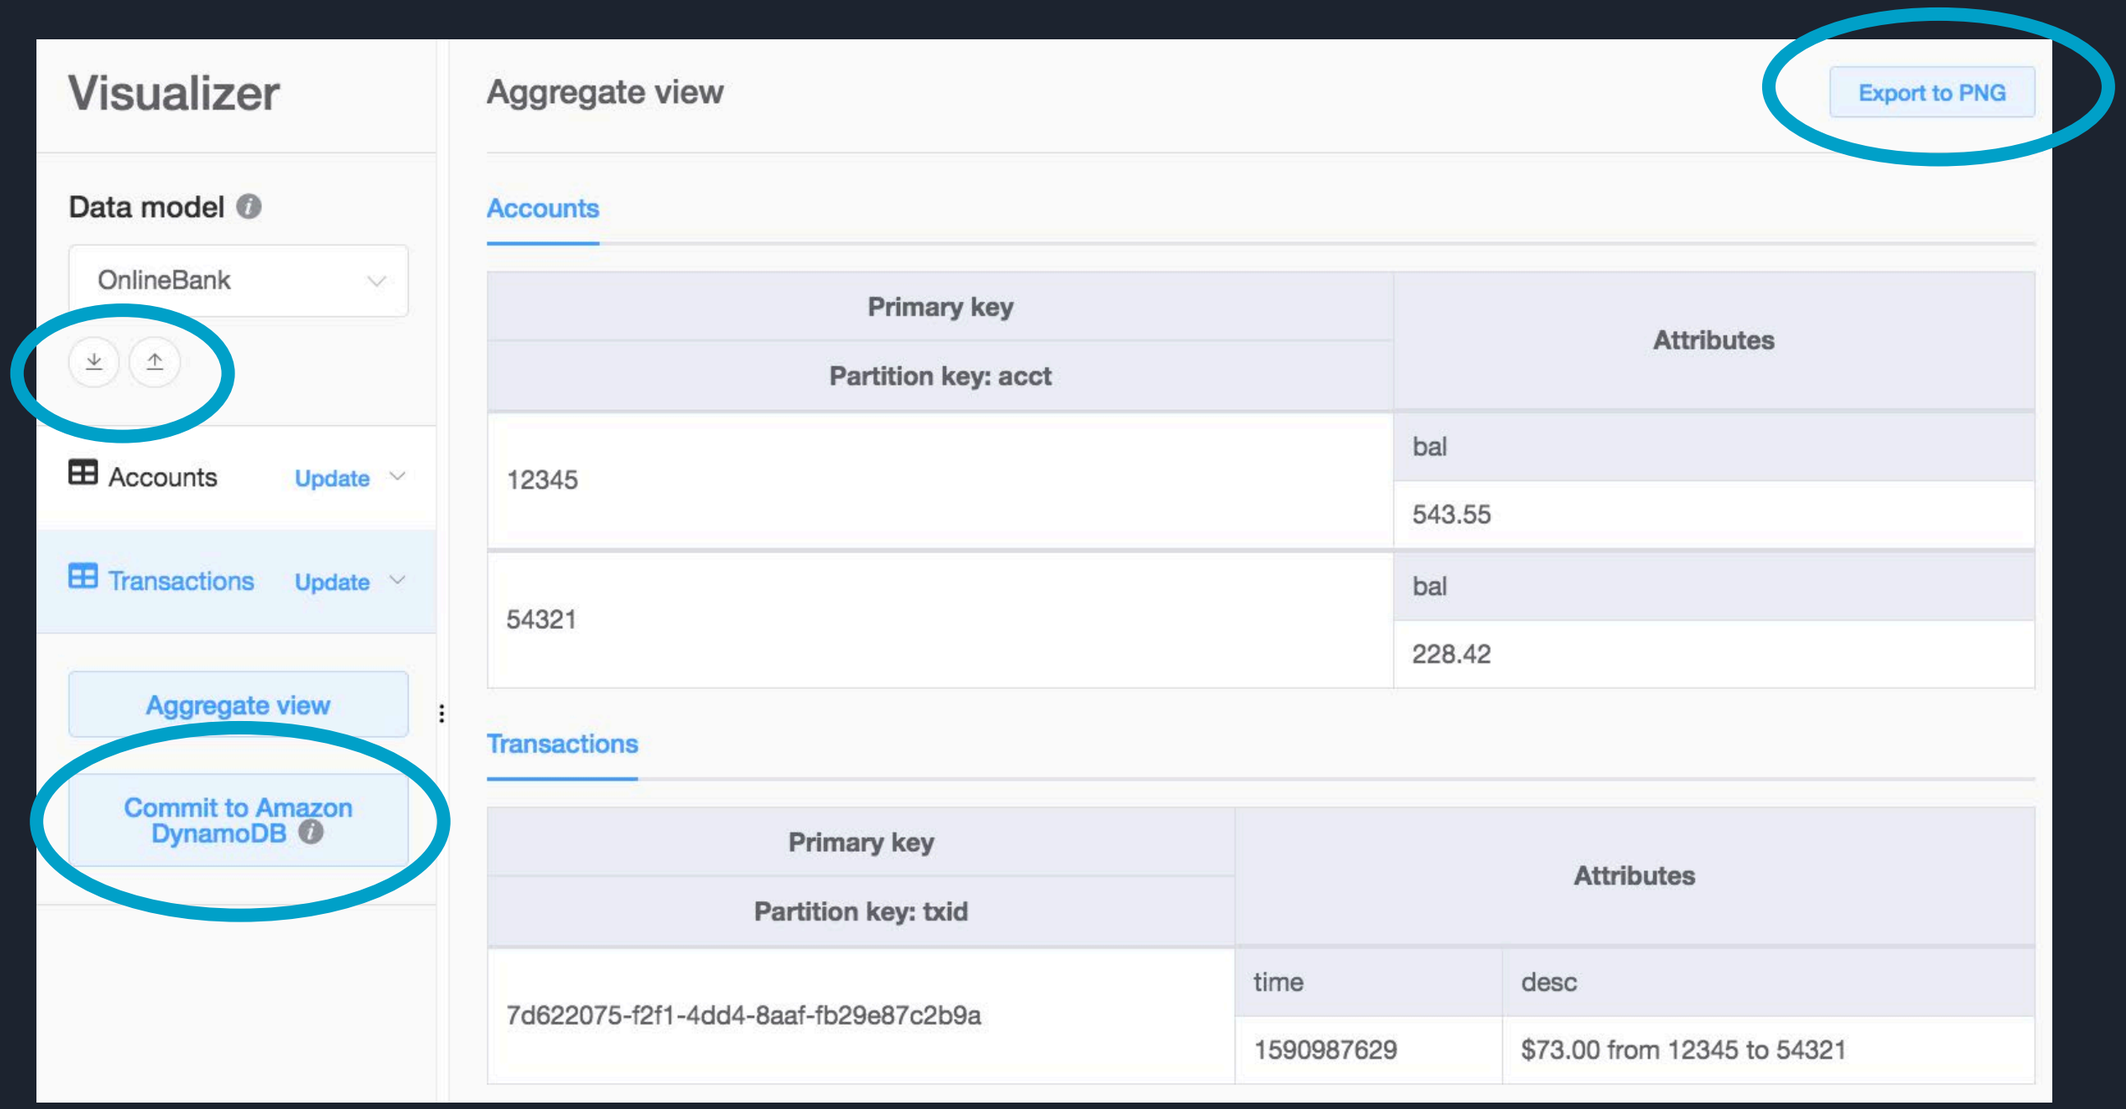

aws databases

Q&A

Pete Naylor
DynamoDB Specialist SA

aws databases

Thank you!

@DynamoDB – Twitter

aws.amazon.com/dynamodb

aws databases

Implementing application workflows

- Workbench operation builder generates sample code for you (Java, JavaScript, Python)
- AWS Management Console, AWS CLI, and Amazon CloudWatch
- ReturnConsumedCapacity
- SDKs: Java, Python, .NET, Go, JavaScript, Node.js, C++, Ruby, and PHP

AWS Free Tier

25 GB storage

25 provisioned RCU

25 provisioned WCU

Resources

On-demand capacity

DynamoDB Accelerator (DAX)

Point-in-time recovery

Global tables

NoSQL Workbench

Advanced data modeling (YouTube)

What's new with Amazon DynamoDB

Amazon's purpose-built database journey

Caching challenges and strategies (Amazon Builders' Library)

Demos for Devs (Twitch)