

# Building a mindset to adapt to single table design in DynamoDB

A philosophical and retrospective rant

Anurag Kale

Cloud Software Architect @ Polestar Cars

AWS Data Hero

# What's this about?

- Single Table Design in DynamoDB
- A shift of mindset
- Bridge the functional gap
- Diplomatic way/s to navigate with users in an enterprise landscape
- Some opinionated take aways from my personal journey

# \$whoami

Anurag Kale

Cloud Software Architect at Polestar Digital

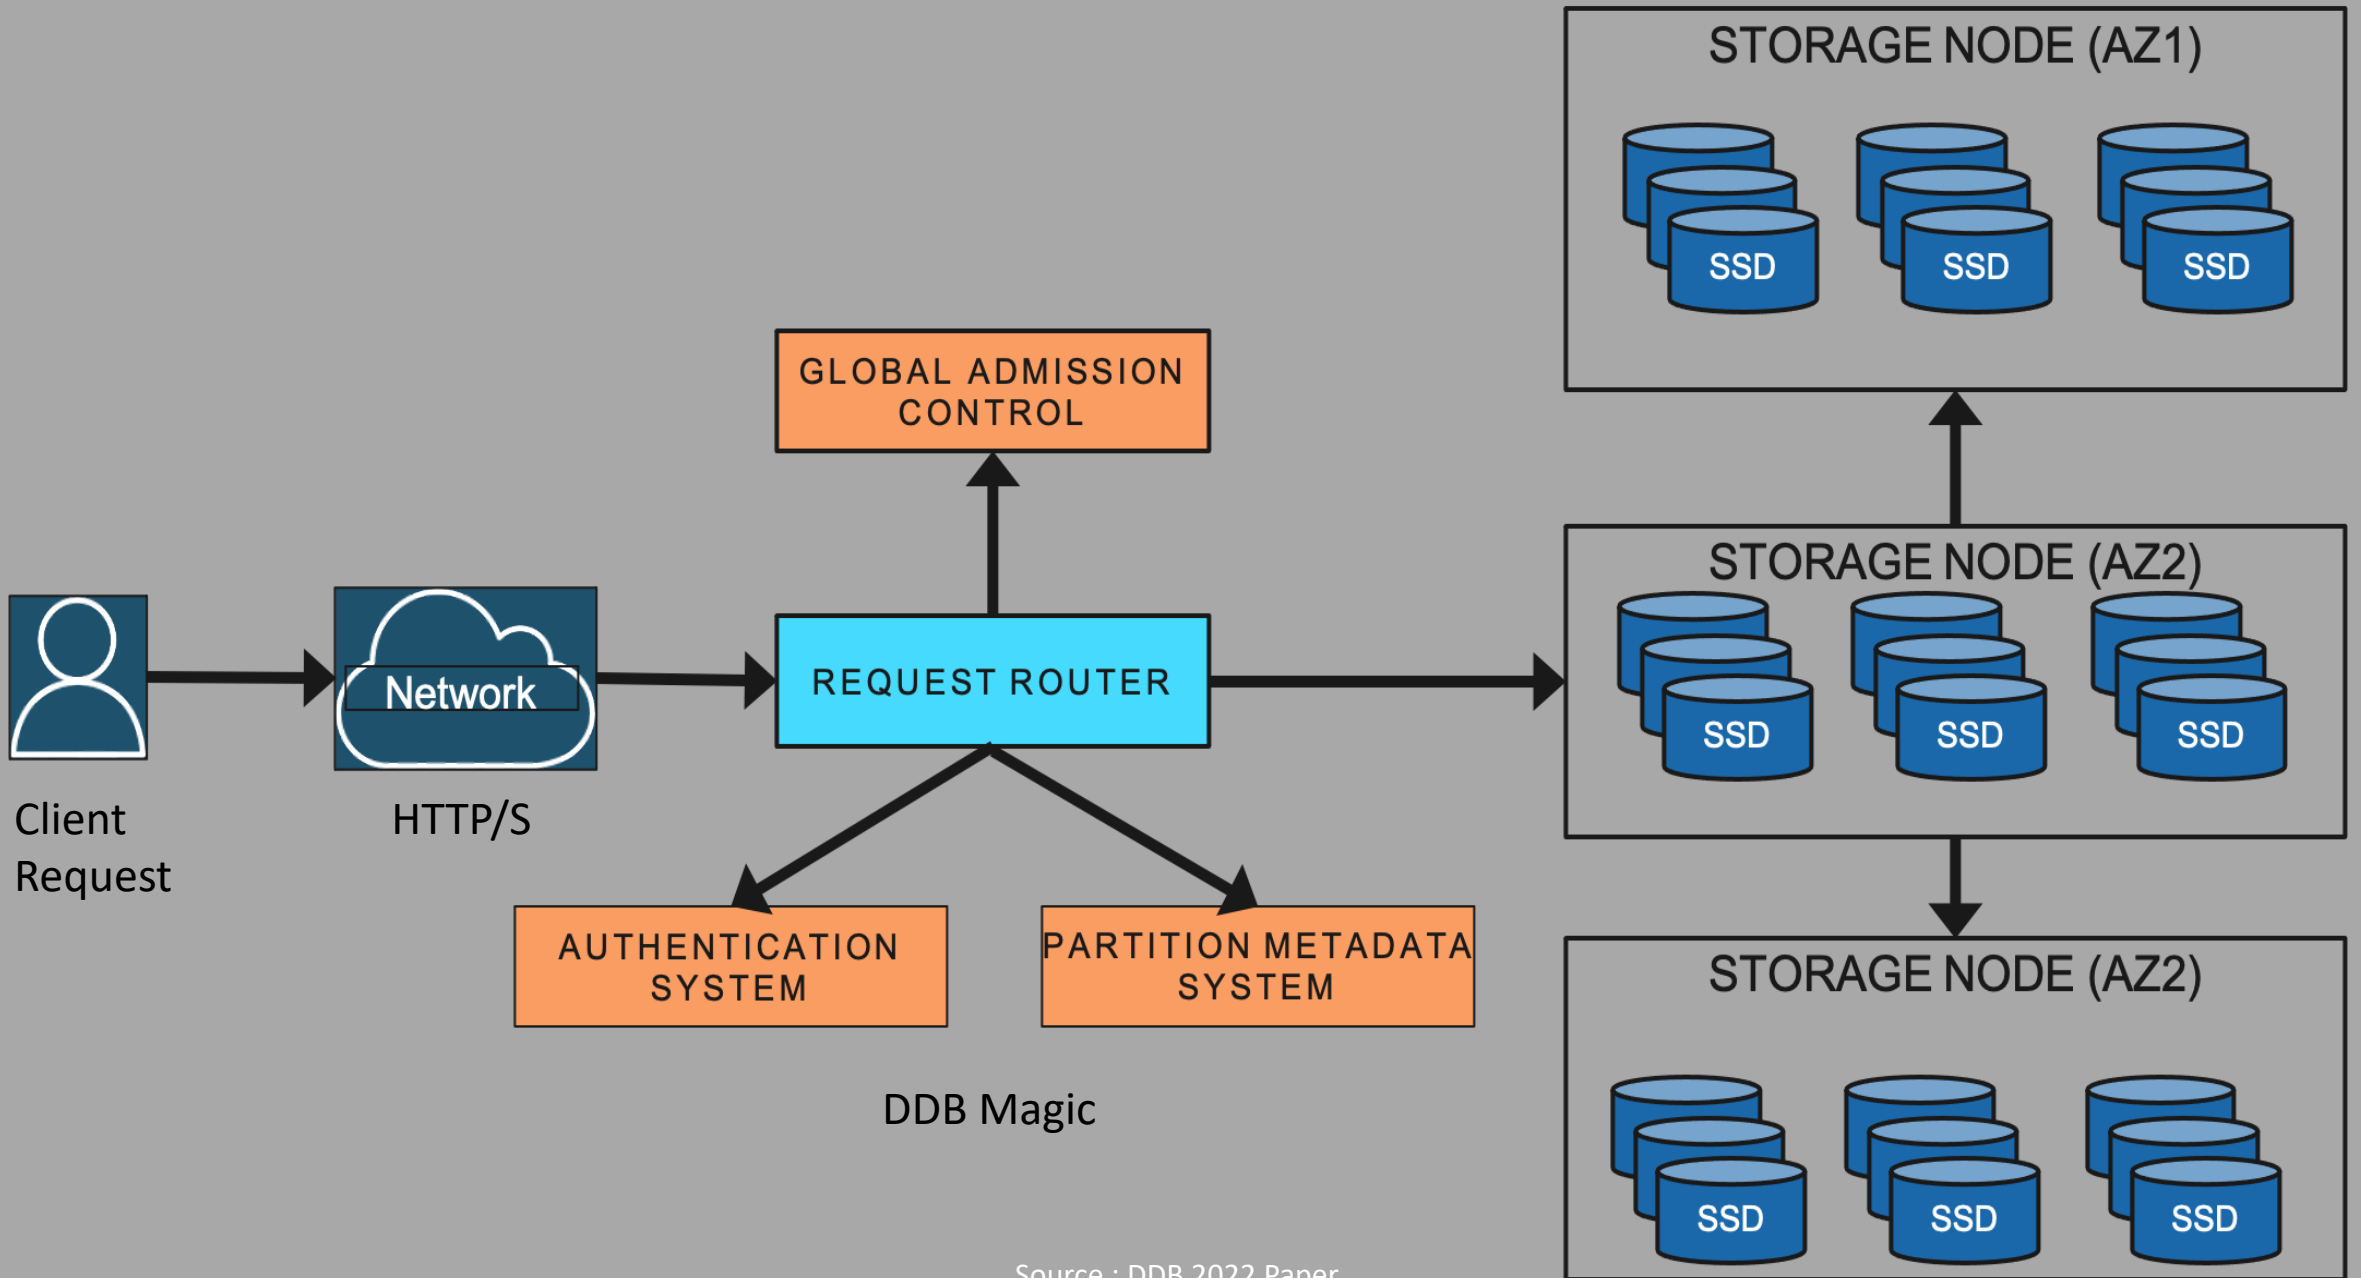
AWS Data Hero



Twitter – @iAnuragKale

# A refresher

DynamoDB Architecture and Single Table Design



# Technical Design Choice

## Specialisation over Generalisation

# Single Table Design

What, why and how

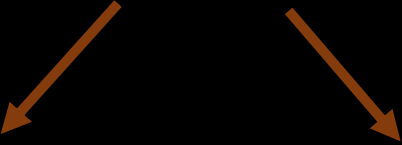
# Table + Primary Key



USERID				
122454678	Full Name	Address	Country	Email
	Anurag Kale	Partille	Sweden	anurag.kale@email.co
216497326	Full Name	Address	Country	Email
	Werner Vogels	Amsterdam	Netherlands	wv@amazon.nl



# Primary Key = Partition Key + Sort Key



UserID	EmailID				
122454678	anurag.kale@email.co	Full Name	Address	Country	Email
		Anurag Kale	Partille	Sweden	anurag.kale@email.co
216497326	wv@amazon.nl	Full Name	Address	Country	Email
		Werner Vogels	Amsterdam	Netherlands	wv@amazon.nl

# Item

UserID	EmailID				
122454678	anurag.kale@email.co	Full Name	Address	Country	Email
		Anurag Kale	Partille	Sweden	anurag.kale@email.co
216497326	wv@amazon.nl	Full Name	Address	Country	Email
		Werner Vogels	Amsterdam	Netherlands	wv@amazon.nl

# Item Collection

PK	SK				
User#122454678	Email#anurag.kale@email.co	Full Name	Address	Country	Email
		Anurag Kale	Partille	Sweden	anurag.kale@email.co
User#122454678	Order#1278903	thing	Price	Warehouse	Make
		iPhone	400 €	Stockholm	Apple
User#122454678	Order#1278904	thing	Price	Reseller	DispatchLocation
		Mouse 187	100 €	MySell LLC	NY, USA

# What is Single Table Design?

- Specialised & opinionated way to store data in a distributed database.
- Data structure is pre-computed to predictable use cases and patterns
- Optimised for predictable performance (at the cost of functionality)
- Specialised platform to store transactional data

Mindshift

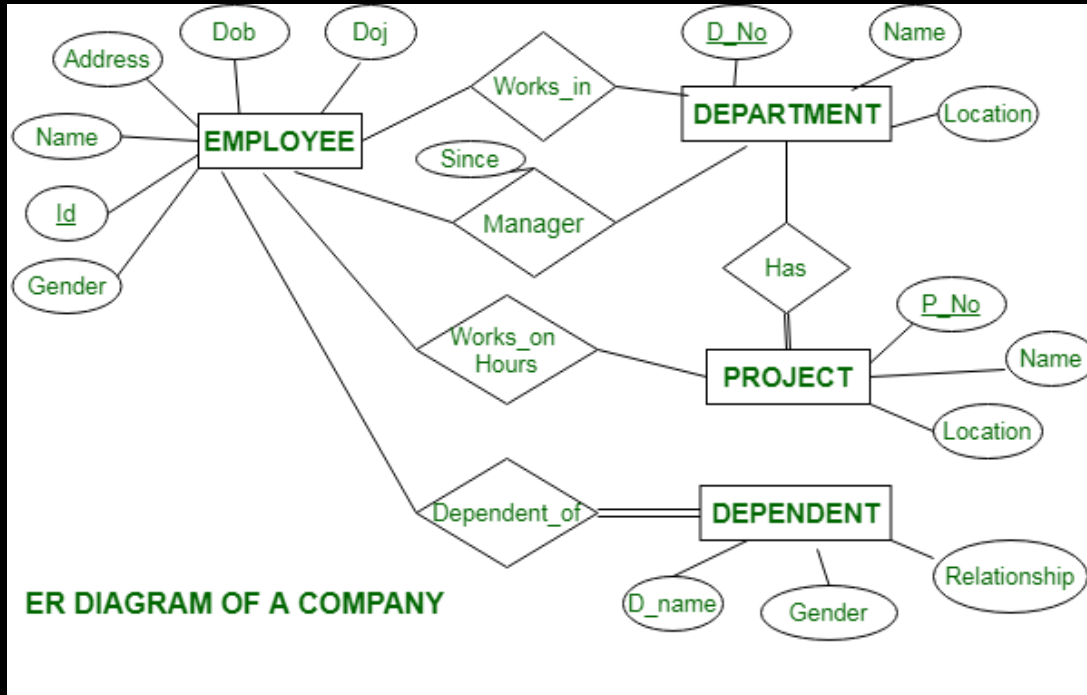
# What should I do?

- Embrace cloud operation model
  - Specilised workloads
  - Distributed responsibility
- Elaborate Design Phase
  - Predict use
  - Pre-empt usage patterns
- Re-train
  - Read - <https://www.dynamodbbook.com>
  - Implement - <https://amazon-dynamodb-labs.workshop.aws>
  - Iterate

# Navigating Enterprise

How to work with people that do not know DynamoDB?

# Isolate Logical and Physical Design



## Movie database

ID (PK)	SK (GSI PK)	Attributes				
MOV#xyz1234	MOV#xyz1234	YEAR	SNAME	TYPE	RUNTIME	...
		2018	Blaze	movie	129	...
	ACT#aac1004	NAME	TITLE	CHARACTER	CO	
		Alia Shawkat	Blaze	Sybil	1	
	...	...	...	...	...	...
MOV#xyz1234	ACT#ab00345	NAME	TITLE	CHARACTER	CO	
		Steve Zahn	Blaze	Oilman	8	
	ACT#bc024311	NAME	TITLE	CHARACTER	CO	
		Ethan Hawke	Blaze	Radio DJ	10	
DIR#bce0557	DIR#bce0557	NAME	TITLE	DIRECTOR		
		Ethan Hawke	Blaze	1		
ACT#aac1004	ACT#aac1004	SNAME	BIO	DOB	KNOIWN FOR	...
		Alia Shawkat	...	1989	Arrested Development	...
#DIR#bce0557	#DIR#bce0557	NAME	TITLE			
		Ethan Hawke	Blaze			
DIR#bce0557	DIR#bce0557	SNAME	BIO	DOB	KNOIWN FOR	...
		Ethan Hawke	...	1970	Training Day	...

Entities are identified by the partition key

Relationships are identified by partition and sort key combo

Start with the most important access patterns...



# Design SQL's, then translate

```
SELECT
```

```
    dept.NAME,
```

```
    Count(emp.id)
```

```
From Employee emp,
```

```
    INNER JOIN Department dept ON dept.id = emp.id
```

```
Where emp.dept_id is not null
```

# Create Item / Item Collection


PK	SK		
Dept#Sales	Count#200	Created_at	Created_by
		2022-03-21	Count_lambda

# Build competence

- Organise Competence Groups
- Invite experts from within organisation OR AWS to deliver fine tuned training sessions
- Use AWS Organisations and/or Control Tower for sandbox accounts
- Build, fail, iterate

Some personal takeaways

# Timeline



# What I learned

- It takes 3-5 iterations to get the model right
- Win some, lose some
- Be a diplomat and meet people in the middle
- Internalise, document and share experiences
- Put it out there, reach out for feedback
- Use AWS if you are on a support plan
- Talk to Hero's, Community Builders and learn from their experience

## In Summary

- Understand the architectural difference
- Get comfortable with specialized workloads
- Reshape and resuse your previous experiences

# Thank You!

Twitter - @iAnuragKale

LinkedIn – [LI.com/in/anuragkale](https://www.linkedin.com/in/anuragkale)