ONE PAGE SUMMARY

IT & Software › NoSQL › DynamoDB

Introduction to DynamoDB

Unlock the power of scalable data with DynamoDB!

DEFINITION

DynamoDB is a fully managed NoSQL database service by AWS that provides fast and predictable performance with seamless scalability. It is designed for high availability and can handle large amounts of data across multiple regions.

KEY CONCEPTS

Tables

DynamoDB stores data in tables, which are similar to tables in relational databases but are schema-less.

Items

An item is a single record in a table, consisting of key-value pairs.

Attributes

Attributes are the individual data elements within an item, akin to columns in a relational database.

Primary Key

Each table must have a primary key that uniquely identifies each item, which can be a single or composite key.

Provisioned Throughput

DynamoDB allows you to specify the read and write capacity for your tables, ensuring performance meets your application's needs.

EXAMPLES

→ Creating a table for user profiles with unique user IDs as primary keys.

→ Storing product information in a table with attributes like name, price, and category.

→ Querying for items based on specific attributes to retrieve user data.

MEMORY TIPS

★ Remember 'Dynamo' as dynamic - it scales with your needs.

★ Think of tables as containers, items as boxes inside, and attributes as the contents of those boxes.

★ Use the acronym 'PITA' for Primary key, Item, Table, Attributes to recall core components.

COMMON MISTAKES

✗ Confusing primary keys with secondary indexes.

✗ Neglecting to set appropriate read/write capacity leading to throttling.

✗ Assuming DynamoDB is a relational database and trying to enforce strict schemas.

QUICK RECAP

DynamoDB is a scalable NoSQL database service by AWS that uses tables to store items with unique primary keys. Understanding its structure, including attributes and throughput settings, is crucial for effective database management.