

Core Concepts

Basic Definitions

<b>DynamoDB:</b> A fully managed, serverless, key-value and document database offered by Amazon Web Services (AWS).
<b>Table:</b> A collection of items, similar to a table in a relational database.
<b>Item:</b> A collection of attributes, which is analogous to a row in a relational database.
<b>Attribute:</b> A key-value pair that describes a property of an item.
<b>Primary Key:</b> A unique identifier for each item in a table, composed of either a partition key or a partition key and sort key.
<b>Partition Key (Hash Key):</b> Used to distribute data across partitions for scalability.
<b>Sort Key (Range Key):</b> Used to sort items within a partition.
<b>Secondary Index:</b> A data structure that allows you to query the table using attributes other than the primary key.

Data Types

<b>Scalar Types</b>	String, Number, Binary, Boolean, Null
<b>Document Types</b>	List, Map
<b>Set Types</b>	String Set, Number Set, Binary Set
<b>Note</b>	DynamoDB is schemaless, meaning each item in a table can have different attributes.

Provisioned vs. On-Demand Capacity

<b>Provisioned Capacity</b>	You specify the number of read and write capacity units (RCUs/WCUs) your application requires. Good for predictable workloads.
<b>On-Demand Capacity</b>	DynamoDB automatically scales capacity based on your application's traffic patterns. Good for unpredictable workloads.

Basic Operations

CRUD Operations

<b>PutItem</b> : Creates a new item or replaces an existing item.  <b>Example (AWS CLI):</b> <pre>aws dynamodb put-item --table-name MyTable --item '{"id": {"N": "1"}, "name": {"S": "Example"}}'</pre>
<b>GetItem</b> : Retrieves an item by its primary key.  <b>Example (AWS CLI):</b> <pre>aws dynamodb get-item --table-name MyTable --key '{"id": {"N": "1"}}'</pre>
<b>UpdateItem</b> : Modifies an existing item.  <b>Example (AWS CLI):</b> <pre>aws dynamodb update-item --table-name MyTable --key '{"id": {"N": "1"}}' --update-expression 'SET #n = :val' --expression-attribute-names '{"#n": "name"}' --expression-attribute-values '{ ":val": {"S": "Updated Example"} }'</pre>
<b>DeleteItem</b> : Deletes an item by its primary key.  <b>Example (AWS CLI):</b> <pre>aws dynamodb delete-item --table-name MyTable --key '{"id": {"N": "1"}}'</pre>

Query and Scan

<b>Query</b>	Retrieves items based on primary key attributes. Requires the partition key and optionally a condition on the sort key. More efficient than <b>Scan</b> .
<b>Scan</b>	Retrieves all items in a table (or a subset based on filter expressions). Less efficient than <b>Query</b> , especially for large tables, as it reads every item.
<b>Example Query (AWS CLI)</b>	<pre>aws dynamodb query --table-name MyTable --key-condition-expression 'id = :id' --expression-attribute-values '{ ":id": {"N": "1"} }'</pre>
<b>Example Scan (AWS CLI)</b>	<pre>aws dynamodb scan --table-name MyTable</pre>

Batch Operations

<b>BatchWriteItem</b>	Performs multiple <b>PutItem</b> and <b>DeleteItem</b> operations in a single request, improving efficiency for bulk data operations.
<b>BatchGetItem</b>	Retrieves multiple items from one or more tables in a single request, reducing the number of API calls.

## Indexes

### Global Secondary Index (GSI)

An index that allows queries on attributes other than the primary key. Can have a different partition and sort key than the base table.
<b>Key characteristics:</b>
<ul style="list-style-type: none"><li>Can be created or deleted at any time.</li><li>Queries can span all items in the table.</li><li>Has its own provisioned throughput capacity.</li></ul>

### Local Secondary Index (LSI)

An index that has the same partition key as the base table but a different sort key. Must be created when the table is created.
<b>Key characteristics:</b>
<ul style="list-style-type: none"><li>Shares the provisioned throughput capacity of the base table.</li><li>Limited to 5 LSIs per table.</li><li>Offers strong consistency reads.</li></ul>

### Choosing an Index

<b>Use GSI when:</b>	<ul style="list-style-type: none"><li>You need to query on attributes other than the primary key.</li><li>Your query patterns are diverse and don't align with the base table's primary key.</li><li>You need to project only a subset of attributes to improve query performance and reduce costs.</li></ul>
<b>Use LSI when:</b>	<ul style="list-style-type: none"><li>You need to query using an alternate sort key but the same partition key as the base table.</li><li>You require strongly consistent reads.</li></ul>

## Best Practices

### Data Modeling

<b>Understand Access Patterns:</b> Before designing your table, carefully analyze your application's read and write access patterns to optimize your schema for performance and cost.
<b>Avoid Hot Partitions:</b> Ensure even distribution of data across partitions by choosing appropriate partition keys. Avoid keys with low cardinality or that lead to uneven distribution of writes.
<b>Denormalization:</b> Consider denormalizing your data by embedding related data within a single item to reduce the need for multiple queries.

### Performance Optimization

<b>Use Projections:</b> When querying indexes, project only the attributes you need to reduce the amount of data read and improve performance.
<b>Batch Operations:</b> Use <code>BatchGetItem</code> and <code>BatchWriteItem</code> to perform multiple read and write operations in a single request, reducing latency and improving throughput.
<b>Parallel Scans:</b> For large tables, use parallel scans to divide the scan operation into multiple segments, improving the overall scan time. (Use with caution as it can consume significant RCUs).

### Security

<b>IAM Roles:</b> Use IAM roles to grant fine-grained permissions to your application to access DynamoDB tables, following the principle of least privilege.
<b>Encryption:</b> Enable encryption at rest and in transit to protect sensitive data. DynamoDB supports encryption using AWS KMS.