

# DynamoDB

*Daniel Charvát, Denisa Šrámková, Viliam Juríček, Šimon Berka*



Amazon DynamoDB

# Introduction to DynamoDB

- **Key-value** document database developed by Amazon
  - They discovered that 90 % of their operations query a single table
  - SQL database tables was thus mostly redundant
- Fully managed, multimaster, durable database with in-memory caching
- Partition through **consistent hashing** to spread data across instance nodes
- Size is defined through **read and write capacity units**
  - Allowed number of operations per second
  - Generally cheaper with less frequent usage

# Companies that are using DynamoDB



# Ranking

Rank			DBMS	Database Model	Score		
Dec 2019	Nov 2019	Dec 2018			Dec 2019	Nov 2019	Dec 2018
1.	1.	1.	Oracle +	Relational, Multi-model	1346.39	+10.33	+63.17
2.	2.	2.	MySQL +	Relational, Multi-model	1275.67	+9.38	+114.42
3.	3.	3.	Microsoft SQL Server +	Relational, Multi-model	1096.20	+14.29	+55.86
4.	4.	4.	PostgreSQL +	Relational, Multi-model	503.37	+12.30	+42.74
5.	5.	5.	MongoDB +	Document, Multi-model	421.12	+7.94	+42.50
6.	6.	6.	IBM Db2 +	Relational, Multi-model	171.35	-1.25	-9.40
7.	7.	8.	Elasticsearch +	Search engine, Multi-model	150.25	+1.85	+5.55
8.	8.	7.	Redis +	Key-value, Multi-model	146.23	+1.00	-0.59
9.	9.	9.	Microsoft Access	Relational	129.47	-0.60	-10.04
10.	10.	11.	Cassandra +	Wide column	120.71	-2.52	-1.10
11.	11.	10.	SQLite +	Relational	120.36	-0.66	-2.65
12.	12.	12.	Splunk	Search engine	90.53	+1.46	+8.34
13.	13.	14.	MariaDB +	Relational, Multi-model	86.79	+1.22	+9.53
14.	14.	15.	Hive +	Relational	86.05	+1.83	+18.67
15.	15.	13.	Teradata +	Relational, Multi-model	78.49	-1.86	-0.67
16.	16.	21.	Amazon DynamoDB +	Multi-model	61.63	+0.26	+7.33
17.	17.	16.	Solr	Search engine	57.22	-0.56	-4.13
18.	19.	20.	SAP Adaptive Server	Relational	55.55	+0.25	-0.27

Source: <https://db-engines.com/en/ranking/>

# Features

## Strengths

- Seamless scalability through automatic instance expansion
- Data backed up to Amazon S3
- Ease of integration with other AWS services

## Drawbacks

- No ACID transactions
  - Although eventual consistency is almost guaranteed
- Not suitable for large binary objects.
- Cross-region replicability not available

# Setting up DynamoDB

There are several options:

- Local
  - Windows, Linux, Mac OS (*downloadable version*)
  - Apache Maven (*POM file*)
  - Docker
- Web service
  1. Sign up to AWS
  2. Get an AWS access key
  3. Configure credentials

# Accessing DynamoDB

Again there are several options:

```
denka@Denka-PC:~$ aws dynamodb list-tables
{
  "TableNames": [
    "Certificates",
    "User"
  ]
}
```

- [AWS Management Console](#)
- [AWS Command Line Interface](#)
- [DynamoDB API](#) - supports Java, JavaScript, .NET, Node.js, PHP, Python (*AWS SDK called Boto 3*), Ruby, C++, Go, Android and iOS

The screenshot shows the AWS Management Console for DynamoDB. The left sidebar contains navigation options: Dashboard, Tables, Backups, Reserved capacity, Preferences, DAX, and Subnet groups. The main content area shows the 'Tables' view with a search filter and a table listing two tables: 'Certificates' and 'User'. The table has columns for Name, Status, Partition key, Sort key, Indexes, Total read capacity, Total write capacity, Auto Scaling, and Encryption.

Name	Status	Partition key	Sort key	Indexes	Total read capacity	Total write capacity	Auto Scaling	Encryption
Certificates	Active	id (Number)	cert_common_name (String)	5	1	1	-	DEFAULT
User	Active	id (Number)	-	0	1	1	-	DEFAULT

# Core components

- a **table** is a collection of **items**
- each **item** is a collection of **attributes**
- **primary keys** are used to uniquely identify each item in a table
- **secondary indexes** provide more querying flexibility

*Other than the primary key, the 'People' table is **schemaless**, which means that neither the attributes nor their data types need to be defined beforehand.*

People

<pre>{   "PersonID": 101,   "LastName": "Smith",   "FirstName": "Fred",   "Phone": "555-4321" }</pre>
<pre>{   "PersonID": 102,   "LastName": "Jones",   "FirstName": "Mary",   "Address": {     "Street": "123 Main",     "City": "Anytown",     "State": "OH",     "ZIPCode": 12345   } }</pre>
<pre>{   "PersonID": 103,   "LastName": "Stephens",   "FirstName": "Howard",   "Address": {     "Street": "123 Main",     "City": "London",     "PostalCode": "E8 5K8"   },   "FavoriteColor": "Blue" }</pre>



# Core components

The **primary key** for table 'Music' **consists of two attributes** (Artist and SongTitle).

Each item in the table must have these two attributes.

The **combination** of Artist and SongTitle **distinguishes each item** in the table from all of the others.

```
{  
  "Artist": "No One You Know",  
  "SongTitle": "My Dog Spot",  
  "AlbumTitle": "Hey Now",  
  "Price": 1.98,  
  "Genre": "Country",  
  "CriticRating": 8.4  
}
```

```
{  
  "Artist": "No One You Know",  
  "SongTitle": "Somewhere Down The Road",  
  "AlbumTitle": "Somewhat Famous",  
  "Genre": "Country",  
  "CriticRating": 8.4,  
  "Year": 1984  
}
```

```
{  
  "Artist": "The Acme Band",  
  "SongTitle": "Still in Love",  
  "AlbumTitle": "The Buck Starts Here",  
  "Price": 2.47,  
  "Genre": "Rock",  
  "PromotionInfo": {  
    "RadioStationsPlaying": [  
      "KHCR",  
      "KQBX",  
      "WTNR",  
      "NJJH"  
    ],  
    "TourDates": {  
      "Seattle": "20150625",  
      "Cleveland": "20150630"  
    },  
    "Rotation": "Heavy"  
  }  
}
```

```
{  
  "Artist": "The Acme Band",  
  "SongTitle": "Look Out, World",  
  "AlbumTitle": "The Buck Starts Here",  
  "Price": 0.99,  
  "Genre": "Rock"  
}
```

# Primary key

- must be specified when creating a table
- DynamoDB support two kinds of primary keys:
  - **partition key**: composed of one attribute (*PersonID*)
  - **partition key and sort key**: composed of two attributes (*Artist, SongTitle*)
- *each primary key attribute must be a scalar (of a data type: string, number, or binary)*

# Secondary indexes

- we can **query the data** in the table **using an alternate key** (in addition to queries against the primary key)
- two kinds of indexes:
  - **global**: both partition and sort key can be different from those on the table
  - **local**: same partition key as the table, but a different sort key
  
- *default indexes limit per table: 20 global, 5 local*

# Secondary indexes

We can query data items by *Artist* (partition key) or by *Artist* and *SongTitle* (partition key and sort key).

If we also wanted to query the data by *Genre* and *AlbumTitle*:

1. Create an index on *Genre* and *AlbumTitle*
2. Query the index



# RDBMS vs. DynamoDB

Characteristics	<i>Relational Database Management System (RDBMS)</i>	<i>Amazon DynamoDB</i>
<i>Optimal Workloads</i>	Ad hoc queries; data warehousing; OLAP	Web-scale applications
<i>Data Model</i>	<b>Requires</b> a well-defined <b>schema</b> (data is normalized into tables, rows, and columns).	<b>Schemaless</b> - Can manage structured or semistructured data.
<i>Performance</i>	<b>optimized for storage</b>	<b>optimized for compute</b>
<i>Scaling</i>	Scale up through faster hardware tables can be span across multiple hosts in a distributed system ( <b>upper limits on scalability</b> ).	Designed to scale out using distributed clusters of hardware ( <b>No upper limit</b> ).

# Creating a table - schema example

cert_info_link	not_valid_before	not_valid_after	cert_common_name	cert_authority	log_type
<a href="http://ct.google...">http://ct.google...</a>	1573257600	1581119999	thesmartlocal0...	cPanel, Inc.	X509LogEntry

```
{
  "TableName": "TestCertificates",
  "KeySchema": [
    { "AttributeName": "cert_info_link", "KeyType": "HASH" },
    { "AttributeName": "cert_authority", "KeyType": "RANGE" }
  ],
  "GlobalSecondaryIndexes": [
    { "IndexName": "log_type_index",
      "KeySchema": [
        {
          "AttributeName": "log_type",
          "KeyType": "HASH"
        },
        {
          "AttributeName": "cert_authority",
          "KeyType": "RANGE"
        }
      ]
    }
  ],
}
```

# Creating a table - schema example

```
    "Projection": {
      "ProjectionType": "ALL"
    },
    "ProvisionedThroughput" : {
      "ReadCapacityUnits": 5,
      "WriteCapacityUnits": 5
    }
  ],
  "AttributeDefinitions": [
    { "AttributeName": "cert_info_link", "AttributeType": "S" },
    { "AttributeName": "cert_authority", "AttributeType": "S" },
    { "AttributeName": "log_type", "AttributeType": "S" }
  ],
  "ProvisionedThroughput": {
    "ReadCapacityUnits": 5,
    "WriteCapacityUnits": 5
  }
}
```

**LIVE DEMO**