Exhibit 13

Amazon DynamoDB

Developer Guide

API Version 2012-08-10

Amazon DynamoDB, you will encounter many similarities, but also many things that are different. This section describes common database tasks, comparing and contrasting SQL statements with their equivalent DynamoDB operations.

NoSQL is a term used to describe non-relational database systems that are highly available, scalable, and optimized for high performance. Instead of the relational model, NoSQL databases (like DynamoDB) use alternate models for data management, such as key-value pairs or document storage. For more information, see http://aws.amazon.com/nosql.

> **Note**
> The SQL examples in this section are compatible with the MySQL relational database management system.
> The DynamoDB examples in this section show the name of the DynamoDB operation, along with the parameters for that operation in JSON format. For code samples that use these operations, see Getting Started with DynamoDB (p. 53).

Topics

SQL or NoSQL?

Today's applications have more demanding requirements than ever before. For example, an online game might start out with just a few users and a very small amount of data. However, if the game becomes successful, it can easily outstrip the resources of the underlying database management system. It is not uncommon for web-based applications to have hundreds, thousands, or millions of concurrent users, with terabytes or more of new data generated per day. Databases for such applications must handle tens (or hundreds) of thousands of reads and writes per second.

Amazon DynamoDB is well-suited for these kinds of workloads. As a developer, you can start with a small amount of provisioned throughput and gradually increase it as your application becomes more popular. DynamoDB scales seamlessly to handle very large amounts of data and very large numbers of users.

The following table shows some high-level differences between a relational database management system (RDBMS) and DynamoDB:

Characteristic	Relational Database Management System (RDBMS)	Amazon DynamoDB
Optimal Workloads	Ad hoc queries; data warehousing; OLAP (online analytical processing).	Web-scale applications, including social networks, gaming, media sharing, and IoT (Internet of Things).

Characteristic	Relational Database Management System (RDBMS)	Amazon DynamoDB
Data Model	The relational model requires a well-defined schema, where data is normalized into tables, rows and columns. In addition, all of the relationships are defined among tables, columns, indexes, and other database elements.	DynamoDB is schemaless. Every table must have a primary key to uniquely identify each data item, but there are no similar constraints on other non-key attributes. DynamoDB can manage structured or semi-structured data, including JSON documents.
Data Access	SQL (Structured Query Language) is the standard for storing and retrieving data. Relational databases offer a rich set of tools for simplifying the development of database-driven applications, but all of these tools use SQL.	You can use the AWS Management Console or the AWS CLI to work with DynamoDB and perform ad hoc tasks. Applications can leverage the AWS software development kits (SDKs) to work with DynamoDB using object-based, document-centric, or low-level interfaces.
Performance	Relational databases are optimized for storage, so performance generally depends on the disk subsystem. Developers and database administrators must optimize queries, indexes, and table structures in order to achieve peak performance.	DynamoDB is optimized for compute, so performance is mainly a function of the underlying hardware and network latency. As a managed service, DynamoDB insulates you and your applications from these implementation details, so that you can focus on designing and building robust, high-performance applications.
Scaling	It is easiest to scale up with faster hardware. It is also possible for database tables to span across multiple hosts in a distributed system, but this requires additional investment. Relational databases have maximum sizes for the number and size of files, which imposes upper limits on scalability.	DynamoDB is designed to scale out using distributed clusters of hardware. This design allows increased throughput without increased latency. Customers specify their throughput requirements, and DynamoDB allocates sufficient resources to meet those requirements. There are no upper limits on the number of items per table, nor the total size of that table.

Accessing the Database

Before your application can access a database, it must be *authenticated* to ensure that the application is allowed to use the database, and *authorized* so that the application can only perform actions for which it has permissions.

The following diagram shows client interaction with a relational database, and with DynamoDB.