

Comparing the Use of Amazon DynamoDB and Apache HBase for NoSQL

Wangechi Doble

September 2014



Contents

Contents	2
Abstract	3
Introduction	3
Overview	4
Amazon DynamoDB Overview	4
Apache HBase Overview	5
Feature Summary	7
Use Cases	9
Data Models	9
Data Types	13
Indexing	15
Data Processing	18
Throughput Model	18
Consistency Model	19
Transaction Model	20
Table Operations	21
Architecture	22
Amazon DynamoDB Architecture Overview	22
Apache HBase Architecture Overview	22
Apache HBase on Amazon EMR Architecture Overview	24
Partitioning	24
Performance Optimizations	25
Amazon DynamoDB Performance Considerations	25
Apache HBase Performance Considerations	27
Apache HBase on Amazon EMR	29
Conclusion	31
Further Reading	31

Abstract

One challenge that architects and developers face today is how to process large volumes of data in a timely, cost effective, and reliable manner. There are several NoSQL solutions in the market today, and choosing the right one for your use case can be difficult. This paper compares two popular NoSQL data stores, [Amazon DynamoDB](#), a fully managed NoSQL cloud database service,¹ and [Apache HBase](#), an open-source, column-oriented, distributed big data store.² Both Amazon DynamoDB and Apache HBase are available in the [Amazon Web Services \(AWS\)](#) cloud.³

Introduction

The Amazon Web Services (AWS) cloud accelerates big data analytics. With access to instant scalability and elasticity on AWS, you can focus on analytics instead of infrastructure. Whether you are indexing large data sets, analyzing massive amounts of scientific data, or processing clickstream logs, AWS provides a range of big data products and services that you can leverage for virtually any data-intensive project.

There is a wide adoption of NoSQL databases in the growing industry of big data and real-time web applications. Amazon DynamoDB and Apache HBase are examples of NoSQL databases, which are highly optimized to yield significant performance benefits over a traditional relational database management system (RDBMS). Both Amazon DynamoDB and Apache HBase can process large volumes of data with high performance and throughput.

Amazon DynamoDB provides a fast, fully managed NoSQL database service. It lets you offload operating and scaling a highly available, distributed database cluster. Apache HBase is an open-source, column-oriented, distributed big data store that runs on the [Apache Hadoop](#) framework.⁴

In the AWS cloud, you can choose to deploy Apache HBase on [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) and manage it yourself.⁵ Alternatively, you can leverage Apache HBase as a managed service on [Amazon Elastic MapReduce \(Amazon EMR\)](#), a fully managed, hosted⁶ Hadoop framework on top of Amazon EC2. The following figure

¹ <http://aws.amazon.com/dynamodb/>

² <http://hadoop.apache.org/>

³ <http://aws.amazon.com/>

⁴ <http://hadoop.apache.org/>

⁵ <http://aws.amazon.com/ec2/>

⁶ <http://aws.amazon.com/elasticmapreduce/>

shows the relationship of AWS, Amazon DynamoDB, Amazon EC2, Amazon EMR, and Apache HBase.




Figure 1: Relation of AWS, Amazon DynamoDB, Amazon EC2, Amazon EMR, and Apache HBase

Overview

This section summarizes the main benefits, features, and use cases of Amazon DynamoDB and Apache HBase.

Amazon DynamoDB Overview

Amazon DynamoDB is a fully managed NoSQL database service that provides fast and predictable performance with seamless scalability. Amazon DynamoDB offers the following benefits:

- **No administrative overhead**—Amazon DynamoDB manages the burdens of hardware provisioning, setup and configuration, replication, cluster scaling, hardware and software updates, and monitoring and handling of hardware failures.
- **Virtually unlimited throughput and scale**—The provisioned throughput model of Amazon DynamoDB allows you to specify throughput capacity to serve nearly any level of request traffic. With Amazon DynamoDB, there is virtually no limit to the amount of data that can be stored and retrieved.
- **Elasticity and flexibility**—Amazon DynamoDB can handle unpredictable workloads with predictable performance and still maintain a stable latency profile that shows no latency increase or throughput decrease as the data volume rises with increased

usage. Amazon DynamoDB lets you dial up and down capacity as needed to handle variable workloads. These qualities render Amazon DynamoDB a suitable choice for online applications with the potential to go viral anytime.

- **Integration with other AWS services**—Amazon DynamoDB integrates seamlessly with other AWS services such as [Amazon Identity and Access Management \(Amazon IAM\)](#) to control access⁷ to Amazon DynamoDB resources, [Amazon CloudWatch](#) to monitor⁸ a variety of Amazon DynamoDB performance metrics, [Amazon Kinesis](#) for real-time data ingestion,⁹ [Amazon Simple Storage Service \(Amazon S3\)](#) for Internet storage,¹⁰ Amazon EMR to provide enhanced advanced analytics capabilities, [Amazon Redshift](#) to provide business intelligence¹¹ capabilities, and [AWS Data Pipeline](#) to automate data-driven workflows.¹²

Apache HBase Overview

Apache HBase, a Hadoop database, offers the following benefits:

- **Efficient storage of sparse data**—Apache HBase provides fault-tolerant storage for large quantities of sparse data using column-based compression. Apache HBase is capable of storing and processing billions of rows and millions of columns per row.
- **Store for high frequency counters**—Apache HBase is suitable for tasks such as high-speed counter aggregation because of its consistent reads and writes.
- **High write throughput and update rates**—Apache HBase supports low latency lookups and range scans, efficient updates and deletions of individual records, and high write throughput.
- **Support for multiple Hadoop jobs**—The Apache HBase data store allows data to be used by one or more Hadoop jobs on a single cluster or across multiple Hadoop clusters.

Two of the most common deployment models for Apache HBase are the self-managed model and the managed Apache HBase solution hosted on Amazon EMR. The following sections provide a description of these deployment options.

Self-managed Apache HBase Deployment Model

The Apache HBase self-managed model offers the most flexibility in terms of cluster management, but also presents the following challenges:

⁷ <http://aws.amazon.com/iam/>

⁸ <http://aws.amazon.com/cloudwatch/>

⁹ <http://aws.amazon.com/kinesis/>

¹⁰ <http://aws.amazon.com/s3/>

¹¹ <http://aws.amazon.com/redshift/>

¹² <http://aws.amazon.com/datapipeline/>

- **Administrative overhead**—You must deal with the administrative burden of provisioning and managing your Apache HBase clusters.
- **Capacity planning**—As with any traditional infrastructure, capacity planning is difficult and often prone to significant costly error. For example, you could over-invest and end up paying for unused capacity or under-invest and risk performance or availability issues.
- **Memory management**—Apache HBase is mainly memory-driven. Memory can become a limiting factor as the cluster grows. It is important to determine how much memory is needed to run diverse applications on your Apache HBase cluster to prevent nodes from swapping data too often to the disk. The number of Apache HBase nodes and memory requirements should be planned well in advance.
- **Compute, storage, and network planning**—Other key considerations for effectively operating an Apache HBase cluster include compute, storage, and network. These infrastructure components often require dedicated Apache Hadoop/Apache HBase administrators with specialized skills.

Managed Apache HBase on Amazon EMR

Apache HBase on Amazon EMR is optimized to run on AWS and offers the following benefits:

- **Minimal administrative overhead**—Amazon EMR handles provisioning of Amazon EC2 instances, security settings, Apache HBase configuration, log collection, health monitoring, and replacement of faulty instances. You still have the flexibility to access the underlying infrastructure and customize Apache HBase further, if desired.
- **Easy and flexible deployment options**—You can deploy Apache HBase on Amazon EMR using the AWS Management Console or leverage the [AWS Command Line Interface \(AWS CLI\)](#) tools.¹³ Once launched, resizing an Apache HBase cluster is easily accomplished with a single API call. Activities such as modifying the Apache HBase configuration at launch time or installing third-party tools such as [Ganglia](#) for monitoring performance metrics¹⁴ are feasible with custom or predefined scripts.
- **Unlimited scale**—With Apache HBase running on Amazon EMR, you can gain significant cloud benefits such as easy scaling, low cost, pay only for what you use, and ease of use as opposed to the self-managed deployment model on Amazon EC2.
- **Integration with other AWS services**—Amazon EMR is designed to seamlessly integrate with other AWS services such as Amazon S3, Amazon DynamoDB, Amazon EC2, and Amazon CloudWatch.

¹³ <http://docs.aws.amazon.com/general/latest/gr/GetTheTools.html>

¹⁴ <http://ganglia.sourceforge.net/>

- Built-in backup feature**—A key benefit of Apache HBase running on Amazon EMR is the built-in mechanism available for backing up Apache HBase data durably in Amazon S3. Using this feature, you can schedule full or incremental backups, and rollback or even restore backups to existing or newly launched clusters anytime.

Feature Summary

Amazon DynamoDB and Apache HBase both possess characteristics that are critical for successfully processing massive amounts of data. The following table provides a summary of key features of Amazon DynamoDB and Apache HBase that can help you understand key similarities and differences between the two databases. These features are discussed in later sections.

Feature	Amazon DynamoDB	Apache HBase
Description	Hosted, scalable database service by Amazon	Column store based on Apache Hadoop and on concepts of BigTable
Website	aws.amazon.com/dynamodb	hbase.apache.org
Documentation	aws.amazon.com/documentation/dynamodb	hbase.apache.org
Developer	Amazon	Apache Software Foundation
Initial Release	2012	2008
License	N/A	Open Source
Implementation Language	-	Java
Server Operating Systems	Hosted	Linux, Unix, Windows
Database Model	Key-value store	Wide column store
Data Scheme	Schema free	Schema free
Typing	Yes	No
APIs and Other Access Methods	RESTful HTTP API, .NET, ColdFusion, Erlang, Groovy	Java API, RESTful HTTP API, Thrift, C, C#, C++
Supported Programming Languages	Java, JavaScript, Perl, PHP, Python, Ruby	Groovy, Java, PHP, Python, Scala
Server-side Scripts	No	Yes
Triggers	No	Yes
Partitioning Methods	Sharding	Sharding

Feature	Amazon DynamoDB	Apache HBase
Throughput Model	User provisions throughput	Limited to hardware configuration
Partitioning	Automatic partitioning	Automatic sharding
Replication	Multiple times within a region	Multiple times across many servers
Durability	Yes	Yes
Administration	No administration overhead	High administration overhead in self-managed and minimal on Amazon EMR
User Concepts	Access rights for users and roles can be defined via AWS Identity and Access Management (IAM)	Access Control Lists (ACL)
Data Model		
Row	An item can have any number of attributes	A row can have any number of columns that can further be grouped into column families
Row Size	Item size restriction	No row size restrictions
Primary Key	Hash or composite hash-range key	Row key
Foreign Key	No	No
Indexes	Optional local secondary indexes and global secondary indexes can be created for tables with hash-range primary keys	No built-in index model, but indexes can be implemented as secondary tables or coprocessors
Transactions		
Row Transactions	Item-level transactions	Single-row transactions
Multi-row Transactions	No	No
Cross-table Transactions	No	No
Consistency Model	Eventually consistent and strongly consistent reads	Strongly consistent reads and writes
Concurrency	Yes	Yes
Updates	Conditional updates	Atomic read-modify-write

Table 1: Amazon DynamoDB and Apache HBase Feature Summary

Use Cases

Amazon DynamoDB and Apache HBase are optimized to process massive amounts of data. Popular use cases for Amazon DynamoDB and Apache HBase include the following:

- **High volume special events**—Popular events, such as the Super Bowl, Olympics, and World Cup, or even seasonal events, such as national electoral campaigns, are of relatively short duration and have variable workloads with the potential to consume large amounts of resources. Amazon DynamoDB lets you dial up capacity when you need it and dial back down as needed to handle variable workloads. This quality renders Amazon DynamoDB a suitable choice for such high volume special events.
- **Social media applications**—Community-based applications, such as online-gaming, photo-sharing, location-aware applications, and so on, have unpredictable usage patterns with the potential to go viral anytime. The elasticity and flexibility of Amazon DynamoDB make it suitable for such high volume, variable workloads.
- **Batch-oriented processing**—For large datasets, such as log data, weather data, product catalogs, and so on, you might already have large amounts of historical data that you want to maintain for historical trend analysis, but need to ingest and batch process current data for predictive purposes. For these types of workloads, Apache HBase is a good choice because of its high read and write throughput and efficient storage of sparse data.
- **Reporting**—To process and report on high volume transactional data such as daily stock market trades, Apache HBase is a good choice. This is because Apache HBase supports high throughput writes and update rates, which make it suitable for storage of high frequency counters and complex aggregations.
- **Real-time analytics**—The payload or message size in event data, such as tweets, E-commerce, and so on, is relatively small when compared with application logs. If you want to ingest streaming event data in real-time for sentiment analysis, ad serving, trending analysis, and so on, Amazon DynamoDB lets you dial up throughput capacity when you need it, and dial it back down when you are done, with no downtime. Apache HBase can handle real-time ingestion of data, such as application logs, with ease due to its high write throughput and efficient storage of sparse data. Combining this with Hadoop's ability to handle sequential reads and scans in a highly optimized way renders Apache HBase a powerful tool for real-time data analytics.

Data Models

Amazon DynamoDB and Apache HBase key/value stores are designed with the goal to deliver significant performance benefits with low latency and high throughput. To achieve this goal, key/value stores are designed with simpler and less constrained data models than traditional relational databases. Although the fundamental data model building-

blocks are similar in both Amazon DynamoDB and Apache HBase, each database uses a distinct terminology to describe its specific data model.

At a high level, a database is a collection of tables, and each table is a collection of rows. A row can contain one or more columns. In most cases, NoSQL database tables typically do not require a formal schema except for a mandatory primary key that uniquely identifies each row. The following table illustrates the high-level concept of a NoSQL database.

Table					
Row	Primary Key	Column 1	Column 2	Column 3	Column n

Table 2: High Level NoSQL Database Table Representation

Columnar databases are devised to store each column separately so that aggregate operations for one column of the entire table are significantly quicker than the traditional row storage model.

From a comparative standpoint, a row in Amazon DynamoDB is referred to as an *item*, and each *item* can have any number of *attributes*. An attribute has a key and a value, and is commonly referred to as a name-value pair. An Amazon DynamoDB table can have unlimited items indexed by primary key, as shown in the following example.

Table					
Item 1	Primary Key	Attribute 1	Attribute 2	Attribute 3	Attribute ...n
Item 2	Primary Key	Attribute 1		Attribute 3	
Item n	Primary Key		Attribute 2	Attribute 3	

Table 3: High-level Representation of Amazon DynamoDB Table

Amazon DynamoDB defines two types of primary keys: a single *hash* primary key (Table 4) and a composite *hash-range* primary key with two attributes (Table 5).

Table					
Item	Hash Key	Attribute 1	Attribute 2	Attribute 3	Attribute ...n

Table 4: Amazon DynamoDB Hash Primary Key

Table						
Item	Hash Key	Range Key	Attribute 1	Attribute 2	Attribute 3	attribute ...n

Table 5: Amazon DynamoDB Hash-Range Primary Key

In Amazon DynamoDB, a single attribute hash primary key is useful for quick reads and writes of data. For example, *PersonID* serves as the hash primary key in the following *Person* table.

Person Table					
Item	PersonId	FirstName	LastName	Zipcode	Gender
Item 1	1001	Fname-1	Lname-1	00000	
Item 2	1002	Fname-2	Lname-2		M
Item 3	2002	Fname-3	Lname-3	10000	F

Table 6: Person Amazon DynamoDB Table

A composite hash-range key in Amazon DynamoDB is indexed as a hash key element and a range key element. This multi-part key maintains a hierarchy between the first and second element values. Holding the hash key element constant facilitates searches across the range key element to retrieve items quickly for a given hash key. In the following *GameScores* table, the composite hash-range key is a combination of *PersonId* (hash) and *GameId* (range).

GameScores Table						
	PersonId (HashKey)	GameId RangeKey	TopScore	TopScoreDate	Wins	Losses
item1	1001	Game01	67453	2013-12-09:17:24:31	73	21
item2	1001	Game02	98567	2013-12-11:14:14:37	98	27
Item3	1002	Game01	43876	2013-12-15:19:24:39	12	23
Item4	2002	Game02	65689	2013-10-01:17:14:41	23	54

Table 7: GamesScores DynamoDB Table

Although there is no explicit limit on the number of attributes associated with an individual item in an Amazon DynamoDB table, there are restrictions on the aggregate size of an item or payload, including all attribute names and values. A small payload can potentially improve performance and reduce costs because it requires fewer resources to process. For information about how to handle items that exceed the maximum item size, see [Guidelines for Working with Items](http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForItems.html)¹⁵ in the [Amazon DynamoDB Developer Guide](http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/).¹⁶

¹⁵ <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/GuidelinesForItems.html>

¹⁶ <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/>

In Apache HBase, the most basic unit is a column. One or more columns form a *row*. Each row is addressed uniquely by a primary key referred to as a *row key*. A row in Apache HBase can have millions of columns. Each column can have multiple versions with each distinct value contained in a separate *cell*.

One fundamental modeling concept in Apache HBase is that of a *column family*. A column family is a container for grouping sets of related data together within one table, as shown in the following example.

Table							
		Column Family 1		Column Family 2		Column Family 3	
row	row key	Column 1	Column 2	Column 3	Column 4	Column 5	Column 6

Table 8: Apache HBase Row Representation

Apache HBase groups columns with the same general access patterns and size characteristics into column families to form a basic unit of separation. For example, in the following *Person* table, you can group personal data into one column family called *personal_info* and the statistical data into a *demographic* column family. Any other columns in the table would be grouped accordingly as well, as shown in the following example.

Person Table					
		personal_info		demographic	
	row key	firstname	lastname	zipcode	gender
row 1	1001	Fname-1	Lname-1	00000	
row 2	1002	Fname-2	Lname-2		M
row 3	2002	Fname-3	Lname-3	10000	F

Table 9: Person Table in Apache HBase

Columns are addressed as a combination of the column family name and the column qualifier expressed as *family:qualifier*. All members of a column family have the same prefix. In the preceding example, the *firstname* and *lastname* column qualifiers can be referenced as *personal_info:firstname* and *personal_info:lastname*, respectively.

Column families allow you to fetch only those columns that are required by a query. All members of a column family are physically stored together on a disk. This means that optimization features, such as performance tunings, compression encodings, and so on, can be scoped at the column-family level.

The row key is a combination of user and game identifiers in the following Apache HBase *GameScores* table. A row key can consist of multiple parts concatenated to

provide an immutable way of referring to entities. From an Apache HBase modeling perspective, the resulting table is *tall-narrow*. This is because the table has few columns relative to the number of rows, as shown in the following example.

GameScores Table					
		top_scores		metrics	
	row key	score	date	wins	loses
row 1	1001-game01	67453	2013-12-09:17:24:31	73	21
row 2	1001-game02	98567	2013-12-11:14:14:37	98	27
row 3	1002-game01	43876	2013-12-15:19:24:39	12	23
row 4	2002-game02	65689	2013-10-01:17:14:41	23	54

Table 10: Tall-Narrow GameScores Apache HBase Table

Alternatively, you can model the game identifier as a column qualifier in Apache HBase. This approach facilitates precise column lookups and supports usage of filters to read data. The result is a *flat-wide* table with few rows relative to the number of columns. This concept of a flat-wide Apache HBase table is shown in the following table.

GameScores Table							
		top_scores			metrics		
	row key	gameId	score	top_score_date	gameId	wins	loses
row 1	1001	game01	98567	2013-12-11:14:14:37	game01	98	27
		game02	43876	2013-12-15:19:24:39	game02	12	23
row 2	1002	game01	67453	2013-12-09:17:24:31	game01	73	21
row 3	2002	game02	65689	2013-10-01:17:14:41	game02	23	54

Table 11: Flat-Wide GameScores Apache HBase Table

For performance reasons, it is important to keep the number of column families in your Apache HBase schema low. Anything above three-column families can potentially degrade performance. The recommended best practice is to maintain a one-column family in your schemas and introduce a second-column family and third-column family only if data access is limited to a one-column family at a time. Note that Apache HBase does not impose any restrictions on row size.

Data Types

Both Amazon DynamoDB and Apache HBase support unstructured datasets with a wide range of data types.

Amazon DynamoDB supports the data types shown in the following table.

	Type	Description	Example (JSON Format)
Scalar			
	String	Unicode with UTF8 binary encoding	{ "S": "Game01" }
	Number	Positive or negative exact-value decimals and integers	{ "N": "67453" }
	Binary	Encoded sequence of bytes	{ "B": "dGhpcyB0ZXh0IGlzlGJhc2U2NC1lbnNvZGVk" }
Multi-valued			
	String Set	Unique set of strings	{ "SS": ["Black", "Green"] }
	Number Set	Unique set of numbers	{ "NS": ["42.2", "-19.87"] }
	Binary Set	Unique set of binary values	{ "BS": ["U3Vubnk=", "UmFpbmk="] }

Table 12: Amazon DynamoDB Data Types

Each Amazon DynamoDB attribute is a name-value pair that is either a single-valued, scalar data type or a multi-valued set. Primary key attributes can be any scalar types, but not multi-valued types. Individual items in an Amazon DynamoDB table can have any number of attributes. Binary type attributes can store any binary data, for example, compressed data, encrypted data, or even images.

To recap, Apache HBase defines the following concepts:

- **Row**—An atomic byte array or key/value container.
- **Column**—A key within the key/value container inside a row.
- **Column Family**—Divides columns into related subsets of data that are stored together on disk.
- **Timestamp**—Apache HBase adds the concept of a fourth dimension column that is expressed as an explicit or implicit timestamp. A timestamp is usually represented as a *long* integer in *milliseconds*.
- **Value**—A time-versioned value in the key/value container. This means that a cell can contain multiple versions of a value that can change over time. Versions are stored in decreasing timestamp, with the most recent first.

Apache HBase defines its key/value pairs as arbitrary arrays of bytes. Because row keys and column qualifiers are also arbitrary arrays of bytes, almost anything can serve as a row key or column qualifier, from strings to binary representations of longs or even serialized data structures.

Column family names *must* be composed of printable characters in human-readable format. This is because column family names are used as part of the directory name in the file system. Furthermore, column families must be declared up front at schema definition time. Column qualifiers are not subjected to this restriction and can be composed of any arbitrary binary characters and created at runtime.

Indexing

In general, data is indexed by primary key for fast retrieval in both Amazon DynamoDB and Apache HBase. Secondary indexes extend the basic indexing functionality and provide an alternate query path in addition to queries against the primary key.

Amazon DynamoDB supports two kinds of secondary indexes on a table that already implements a hash-and-range key:

- **Local secondary index**—An index that has the same hash key as the table, but a different range key.
- **Global secondary index**—An index with a hash and range key that can be different from those on the table.

You can define one or more local secondary indexes and one or more global secondary indexes per table.

A local secondary index organizes data by the index range key, rather than by the range key for the source table. Every local secondary index automatically contains the hash and range attributes from its parent table. A local secondary index can copy or project some or all of the attributes from the source table and provide an alternate query path for efficiently accessing data.

In the example *GameScores* table introduced in the preceding section, you can define *PersonTopScoresIndex* as a local secondary index for that table. This index contains the same hash key, *PersonId*, as the source table and also defines *TopScoreDate* as its range key. The range key value from the source table (in this example, *GameId*) is automatically projected or copied into the index, but it is not a part of the index key, as shown in the following table.

PersonTopScoresIndex			
Index Key		Attribute1	Attribute2
PersonId (HashKey)	TopScoreDate (RangeKey)	GameId	TopScore
1001	2013-12-09:17:24:31	Game01	67453
1001	2013-12-11:14:14:37	Game02	98567

PersonTopScoresIndex			
1002	2013-12-15:19:24:39	Game01	43876
2002	2013-10-01:17:14:41	Game02	65689

Table 13: Local Secondary Index in Amazon DynamoDB

You can decide whether to copy or project additional non-key attributes from the source table into the index. For example, *PersonTopScoresIndex* includes a copy of the non-key attribute *TopScore* as part of its definition, whereas the following *PersonTopScoresIndex_1* excludes the same attribute from its definition.

PersonTopScoresIndex_1		
Index Key		Attribute1
PersonId (HashKey)	TopScoreDate (RangeKey)	GameId
1001	2013-12-09:17:24:31	Game01
1001	2013-12-11:14:14:37	Game02
1002	2013-12-15:19:24:39	Game01
2002	2013-10-01:17:14:41	Game02

Table 14: Local Secondary Index in Amazon DynamoDB without Non-Key Attribute Projection

The main difference between a global secondary index and a local secondary index is that a global secondary index defines a completely new hash/range index on a table. You can define any attribute as the hash key for the global secondary index as long as its data type is scalar rather than a multi-value set.

When you work with global secondary indexes, keep in mind the following points:

- Unlike associated tables, key values in a global secondary index do not need to be unique.
- Any table attribute can be a key, including attributes that are not present in all items, as long as the data types are scalar rather than multi-value sets.
- Every global secondary index must have a hash key and an optional range key.
- The source table’s primary key attributes are automatically projected into the index, as is the case with local secondary indexes.
- Unlike local secondary indexes that consume read and write capacity units from the table, every global secondary index has its own provisioned throughput settings for

read and write activity. For more information about provisioned throughput, see the [Throughput Model](#) section in this whitepaper.

The following *LeaderBoardIndex* example shows a global secondary index for the *GameScores* table, as discussed in preceding sections.

LeaderBoardIndex		
Index Key		Attribute1
GameId (HashKey)	TopScore (RangeKey)	PersonId
Game01	98567	1001
Game02	43876	1001
Game01	65689	1002
Game02	67453	2002

Table 15: Global Secondary Index for a GameScores Table in Amazon DynamoDB

The preceding *LeaderBoardIndex* global secondary index defines *GameId* as its hash key and *TopScore* as its range key. The index key does not need to have any of the key attributes from the source table. However, the table's primary key attributes are always present in the global secondary index. In this example, *PersonId* is automatically projected or copied into the index.

The *LeaderBoardIndex*, in this example, allows you to easily obtain a list of top scores for a specific game by querying the index. The results are ordered by *TopScore*, the range key. You can choose to project additional attributes from the source table into the index.

All rows in Apache HBase are always sorted lexicographically by row key. The sort is byte-ordered. This means that each row key is compared on a binary level, byte by byte, from left to right. Row keys are always unique and act as the primary index in Apache HBase.

Although Apache HBase does not have native support for built-in indexing models such as Amazon DynamoDB, you can implement custom secondary indexes to serve as alternate query paths by using these techniques:

- **Create an index in another table**—You can maintain a secondary table that is periodically updated. However, depending on the load strategy, the risk with this method is that the secondary index can potentially become out of sync with the main table. You can mitigate this risk if you build the secondary index while publishing data to the cluster, and perform concurrent writes into the index table.

- **Use the coprocessor framework**—You can leverage the coprocessor framework to implement custom secondary indexes. Coprocessors act like triggers that are similar to stored procedures in RDBMS.

In summary, both Amazon DynamoDB and Apache HBase define data models that allow efficient storage of data to optimize query performance. Amazon DynamoDB imposes a restriction on its item size to allow efficient processing and reduce costs. Apache HBase uses the concept of column families to provide data locality for more efficient read operations.

Amazon DynamoDB supports both scalar and multi-valued sets to accommodate a wide range of unstructured datasets. Similarly, Apache HBase stores its key/value pairs as arbitrary arrays of bytes, giving it the flexibility to store any data type.

Amazon DynamoDB supports built-in secondary indexes and automatically updates and synchronizes all indexes with their parent tables. With Apache HBase, you can implement and manage custom secondary indexes yourself.

From a data model perspective, you can choose Amazon DynamoDB if your item size is relatively small. Although Amazon DynamoDB provides a number of options to overcome row size restrictions, Apache HBase is better equipped to handle large complex payloads with minimal restrictions.

Data Processing

This section highlights foundational elements for processing and querying data within Amazon DynamoDB and Apache HBase.

Throughput Model

Amazon DynamoDB uses a provisioned throughput model to process data. With this model, you can specify your read and write capacity needs in terms of number of input operations per second that a table is expected to achieve. During table creation time, Amazon DynamoDB automatically partitions and reserves the appropriate amount of resources to meet your specified throughput requirements.

To decide on the required read and write throughput values for a table, consider the following factors:

- **Item size**—The read and write capacity units that you specify are based on a predefined data item size per read or per write operation. For more information about

provisioned throughput data item size restrictions, see [Provisioned Throughput in Amazon DynamoDB](#)¹⁷ in the [Amazon DynamoDB Developer Guide](#).¹⁶

- **Expected read and write request rates**—You must also determine the expected number of read and write operations your application will perform against the table, per second.
- **Consistency**—Whether your application requires strongly consistent or eventually consistent reads is a factor in determining how many read capacity units you need to provision for your table. For more information about consistency and Amazon DynamoDB, see the [Consistency Model](#) section in this whitepaper.
- **Local secondary indexes**—Queries against indexes consume provisioned read throughput. For more information, see [Provisioned Throughput Considerations for Local Secondary Indexes](#)¹⁸ in the [Amazon DynamoDB Developer Guide](#).¹⁶
- **Global secondary indexes**—The provisioned throughput settings of a global secondary index are separate from those of its parent table. Therefore, the expected workload on the global secondary index also needs to be taken into consideration when specifying the read and write capacity at index creation time.

Although read and write requirements are specified at table creation time, Amazon DynamoDB lets you increase or decrease the provisioned throughput to accommodate load with no downtime.

In Apache HBase, the number of nodes in a cluster can be driven by the required throughput for reads and/or writes. The available throughput on a given node can vary depending on the data, specifically:

- Key/value sizes
- Data access patterns
- Cache hit rates
- Node and system configuration

You should plan for peak load if load will likely be the primary factor that increases node count within an Apache HBase cluster..

Consistency Model

A database consistency model determines the manner and timing in which a successful

¹⁷ <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/ProvisionedThroughputIntro.html>

¹⁸ <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/LSI.html#LSI.ThroughputConsiderations>

write or update is reflected in a subsequent read operation of that same value.

Amazon DynamoDB lets you specify the desired consistency characteristics for each read request within an application. You can specify whether a read is eventually consistent or strongly consistent.

The eventual consistency option is the default in Amazon DynamoDB and maximizes the read throughput. However, an eventually consistent read might not always reflect the results of a recently completed write. Consistency across all copies of data is usually reached within a second.

A strongly consistent read in Amazon DynamoDB returns a result that reflects all writes that received a *successful* response prior to the read. To get a strongly consistent read result, you can specify optional parameters in a request. It takes more resources to process a strongly consistent read than an eventually consistent read. For more information about read consistency, see [Data Read and Consistency Considerations](#)¹⁹ in the [Amazon DynamoDB Developer Guide](#).¹⁶

Apache HBase reads and writes are strongly consistent. This means that all reads and writes to a single row in Apache HBase are atomic. Each concurrent reader and writer can make safe assumptions about the state of a row. Multi-versioning and time stamping in Apache HBase contribute to its strongly consistent model.

Transaction Model

Unlike RDBMS, NoSQL databases typically have no domain-specific language, such as SQL, to query data. Amazon DynamoDB and Apache HBase provide simple application programming interfaces (APIs) to perform the standard create, read, update, and delete (CRUD) operations.

Neither Amazon DynamoDB nor Apache HBase support multi-item/cross-row or cross-table transactions due to performance considerations. However, both databases provide batch operations for reading and writing multiple items/rows across multiple tables with no transaction guarantees.

Amazon DynamoDB provides atomic item and attribute operations for adding, updating, or deleting data. Further, item-level transactions can specify a condition that must be satisfied before that transaction is fulfilled. For example, you can choose to update an item only if it already has a certain value.

¹⁹ <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/APISummary.html#DataReadConsistency>

Conditional operations allow you to implement [optimistic concurrency control](#) systems²⁰ on Amazon DynamoDB. For conditional updates, Amazon DynamoDB allows atomic increment and decrement operations on existing scalar values without interfering with other write requests.

Apache HBase also supports atomic high update rates (the classic read-modify-write) within a single row key, enabling storage for high frequency counters. Unlike Amazon DynamoDB, Apache HBase uses multi-version concurrency control to implement updates. This means that an existing piece of data is not overwritten with a new one; instead, it becomes obsolete when a newer version is added.

Row data access in Apache HBase is atomic and includes any number of columns, but there are no further guarantees or transactional features spanning multiple rows. Similar to Amazon DynamoDB, Apache HBase supports only single-row transactions.

Table Operations

Amazon DynamoDB and Apache HBase provide scan operations to support large-scale analytical processing. A scan operation is similar to [cursors](#) in RDBMS.²¹ By taking advantage of the underlying sequential, sorted storage layout, a scan operation can easily iterate over wide ranges of records or entire tables. Applying filters to a scan operation can effectively narrow the result set and optimize performance.

Amazon DynamoDB uses parallel scanning to improve performance of a scan operation. A parallel scan logically sub-divides an Amazon DynamoDB table into multiple segments, and then processes each segment in parallel. Rather than using the default scan operation in Apache HBase, you can implement a custom parallel scan by means of the API to read rows in parallel.

Amazon DynamoDB provides a Query API for complex query processing in addition to its scan operation. The Query API is accessible only in tables that define a composite primary key. A query request retrieves items from an Amazon DynamoDB table or index using the hash-and-range primary key.

In summary, Amazon DynamoDB and Apache HBase have similar data processing models in that they both support only atomic single-row transactions. Both databases also provide batch operations for bulk data processing across multiple rows and tables.

²⁰ http://en.wikipedia.org/wiki/Optimistic_concurrency_control

²¹ http://en.wikipedia.org/wiki/Cursor_%28databases%29

One key difference between the two databases is the flexible provisioned throughput model of Amazon DynamoDB. The ability to dial up capacity when you need it and dial it back down when you are done is useful for processing variable workloads with unpredictable peaks.

For workloads that need high update rates to perform data aggregations or maintain counters, Apache HBase is a good choice. This is because Apache HBase supports a multi-version concurrency control mechanism, which contributes to its strongly consistent reads and writes. Amazon DynamoDB gives you the flexibility to specify whether you want your read request to be eventually consistent or strongly consistent depending on your specific workload.

Architecture

This section summarizes key architectural components of Amazon DynamoDB and Apache HBase.

Amazon DynamoDB Architecture Overview

The origins of Amazon DynamoDB architecture can be traced to the genesis of Dynamo, a NoSQL database described in a paper by DeCandia, *et al.*²² At a high level, Amazon DynamoDB is designed for high availability, durability, and consistently low latency (typically in the single digit milliseconds) performance.

Amazon DynamoDB runs on a fleet of AWS managed servers that leverage solid state drives (SSDs) to create an optimized, high-density storage platform. This platform decouples performance from table size, and eliminates the need for the working set of data to fit in memory while still returning consistent, low latency responses to queries. As a managed service, Amazon DynamoDB abstracts its underlying architectural details from the user.

Apache HBase Architecture Overview

Apache HBase is typically deployed on top of the [Hadoop Distributed File System \(HDFS\)](#), which provides a scalable, persistent, storage layer.²³ [Apache ZooKeeper](#) is a critical component for maintaining configuration information²⁴ and managing the entire Apache HBase cluster.

²² Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall and Werner Vogels. 2007. [Dynamo: Amazon's Highly Available Key-value Store](#).

²³ <http://wiki.apache.org/hadoop/HDFS>

²⁴ <http://zookeeper.apache.org/>

The three major Apache HBase components are the following:

- Client API
- Master server
- Region servers

Apache HBase stores data in indexed store files called HFiles on HDFS. The store files are sequences of blocks with a block index stored at the end for fast lookups. The store files provide an API to access specific values as well as to scan ranges of values, given a start and end key.

During a write operation, data is first written to a commit log called a write-ahead-log (WAL) and then moved into memory in a structure called Memstore. When the size of the Memstore exceeds a given maximum value, it is flushed as an HFile to disk. Each time data is flushed from Memstores to disk, new HFiles must be created. As the number of HFiles builds up, a *compaction* process merges the files into fewer, larger files.

A read operation essentially is a merge of data stored in the Memstores and in the HFiles. The WAL is never used in the read operation. It is meant only for recovery purposes if a server crashes before writing the in-memory data to disk.

A region in Apache HBase acts as a store per column family. Each region contains contiguous ranges of rows stored together. Regions can be merged to reduce the number of store files. A large store file that exceeds the configured maximum store file size can trigger a region split.

A region server can serve multiple regions. Each region is mapped to exactly one region server. Region servers handle reads and writes, as well as keep data in-memory until enough is collected to warrant a flush. Clients communicate directly with region servers to handle all data-related operations.

The master server is responsible for monitoring and assigning regions to region servers and uses Apache ZooKeeper to facilitate this task. Apache ZooKeeper also serves as a registry for region servers and a bootstrap location for region discovery.

The master server is also responsible for handling critical functions such as load balancing of regions across region servers, region server failover, and completing region splits, but it is not part of the actual data storage or retrieval path.

You can run Apache HBase in a multi-master environment. All masters compete to run the cluster in a multi-master mode. However, if the active master shuts down, then the remaining masters contend to take over the Master role.

Apache HBase on Amazon EMR Architecture Overview

Amazon EMR defines the concept of instance groups, which are collections of Amazon EC2 instances. The Amazon EC2 virtual servers perform roles analogous to the master and slave nodes of Hadoop. For best performance, Apache HBase clusters should run on at least two Amazon EC2 instances. There are three types of instance groups in an Amazon EMR cluster:

- **Master**—Contains one master node that manages the cluster. You can use the Secure Shell (SSH) protocol to access the master node if you want to view logs or administer the cluster yourself. The master node runs the Apache HBase master server and Apache Zookeeper.
- **Core**—Contains one or more core nodes that run HDFS and store data. The core nodes run the Apache HBase region servers.
- **Task**—(Optional). Contains any number of task nodes.

Partitioning

Amazon DynamoDB stores three geographically distributed replicas of each table to enable high availability and data durability within a region. Data is auto-partitioned primarily using the hash key. As throughput and data size increase, Amazon DynamoDB will automatically repartition and reallocate data across more nodes.

Partitions in Amazon DynamoDB are fully independent, resulting in a shared nothing cluster. However, provisioned throughput is divided evenly across the partitions.

A region is the basic unit of scalability and load balancing in Apache HBase. Region splitting and subsequent load-balancing follow this sequence of events:

1. Initially there is only one region for a table, and as more data is added to it, the system monitors the load to ensure that the configured maximum size is not exceeded.
2. If the region size exceeds the configured limit, the system dynamically splits the region into two at the row key in the middle of the region, creating two roughly equal halves.
3. The master then schedules the new regions to be moved off to other servers for load balancing, if required.

Behind the scenes, Apache Zookeeper tracks all activities that take place during a region split and maintains the state of the region in case of server failure. Apache HBase regions are equivalent to [range partitions](#) that are used in RDBMS sharding.²⁵ Regions can be spread across many physical servers that consequently distribute the load, resulting in scalability.

In summary, as a managed service the architectural details of Amazon DynamoDB are abstracted from you to let you focus on your application details. With the self-managed Apache HBase deployment model, it is crucial to understand the underlying architectural details to maximize scalability and performance. AWS gives you the option to offload Apache HBase administrative overhead, if you opt to launch your cluster on Amazon EMR.

Performance Optimizations

Amazon DynamoDB and Apache HBase are inherently optimized to process large volumes of data with high performance. NoSQL databases typically use an on-disk, column-oriented storage format for fast data access and reduced I/O when fulfilling queries. This performance characteristic is evident in both Amazon DynamoDB and Apache HBase.

Amazon DynamoDB stores items with the same hash key contiguously on disk to optimize fast data retrieval. Similarly, Apache HBase regions contain contiguous ranges of rows stored together to improve read operations. You can enhance performance even further if you apply techniques that maximize throughput at reduced costs, both at the infrastructure and application tiers.

Tip: A recommended best practice is to monitor Amazon DynamoDB and Apache HBase performance metrics to proactively detect and diagnose performance bottlenecks.

The following section focuses on several common performance optimizations that are specific to each database or deployment model.

Amazon DynamoDB Performance Considerations

Performance considerations for Amazon DynamoDB focus on how to define an appropriate read and write throughput and how to design a suitable schema for an application. These performance considerations span both infrastructure level and application tiers.

²⁵ http://en.wikipedia.org/wiki/Partition_%28database%29

Provisioned Throughput Considerations

Factors that must be taken into consideration when determining the appropriate throughput requirements for an application are item size, expected read and write rates, consistency, and secondary indexes, as discussed in the Throughput Model section of this whitepaper.

If an application performs more reads per second or writes per second than a table's provisioned throughput capacity allows, requests above the provisioned capacity will be throttled. For instance, if a table's write capacity is 1,000 units and an application performs 1,500 writes per second for the maximum data item size, Amazon DynamoDB will allow only 1,000 writes per second to go through, and the extra requests will be throttled.

Tip: A recommended best practice is to provision throughput capacity sufficiently far in advance to ensure that it is there when needed, and to monitor performance with Amazon CloudWatch configuring notification alarms for alerting purposes when a certain threshold of consumed capacity units is reached. In addition, Amazon DynamoDB lets you adjust throughput capacity up and down without any downtime.

Primary Key Design Considerations

Primary key design is critical to the performance of Amazon DynamoDB. When storing data, Amazon DynamoDB divides a table's items into multiple partitions, and distributes the data primarily based on the hash key element. The provisioned throughput associated with a table is also divided evenly among the partitions with no sharing of provisioned throughput across partitions.

Tip: To efficiently utilize the overall provisioned throughput, spread the workload across hash key values.

For example, if a table has a very small number of heavily accessed hash key elements, possibly even a single very heavily used hash key element, traffic can become concentrated on a single partition and create "hot spots" of read and write activity within a single item collection. Amazon DynamoDB will throttle unbalanced workloads in such situations.

To get the most out of Amazon DynamoDB throughput, you can build tables where the hash key element has a large number of distinct values. Ensure that values are requested fairly uniformly and as randomly as possible. The same guidance applies to global secondary indexes. Choose hash and range keys that provide uniform workloads to achieve the overall provisioned throughput.

Local Secondary Index Considerations

When querying a local secondary index, the number of read capacity units consumed depends on how the data is accessed. For example, when you create a local secondary

index and project non-key attributes into the index from the parent table, Amazon DynamoDB can retrieve these projected attributes efficiently.

In addition, when you query a local secondary index, the query can also retrieve attributes that are *not* projected into the index. Avoid these types of index queries that read attributes that are not projected into the local secondary index. Fetching attributes from the parent table that are not specified in the local secondary index causes additional latency in query responses and incurs a higher provisioned throughput cost.

Tip: Project frequently accessed non-key attributes into a local secondary index to avoid fetches and improve query performance.

Maintain multiple local secondary indexes in tables that are updated infrequently but are queried using many different criteria to improve query performance. This guidance does not apply to tables that experience heavy write activity.

If very high write activity to the table is expected, one option to consider is to minimize interference from reads by not reading from the table at all. Instead, create a global secondary index with a structure that is identical to that of the table, and then direct all queries to the index rather than to the table.

Global Secondary Index Considerations

If a query exceeds the provisioned read capacity of a global secondary index, that request will be throttled. Similarly, if a request performs heavy write activity on the table, but a global secondary index on that table has insufficient write capacity, then the write activity on the table will be throttled.

Tip: For a table write to succeed, the provisioned throughput settings for the table and global secondary indexes must have enough write capacity to accommodate the write; otherwise, the write will be throttled.

Global secondary indexes support eventually consistent reads, each of which consume one half of a read capacity unit. The number of read capacity units is the sum of all projected attribute sizes across all of the items returned in the index query results. With write activities, the total provisioned throughput cost for a write consists of the sum of write capacity units consumed by writing to the table and those consumed by updating the global secondary indexes.

Apache HBase Performance Considerations

Apache HBase performance tuning spans hardware, network, Apache HBase configurations, Hadoop configurations, and the Java Virtual Machine Garbage Collection settings. It also includes applying best practices when using the client API. To optimize performance, it is worthwhile to monitor Apache HBase workloads with tools such as

Ganglia to identify performance problems early and apply recommended best practices based on observed performance metrics.

Memory Considerations

Memory is the most restrictive element in Apache HBase. Performance-tuning techniques are focused on optimizing memory consumption.

From a schema design perspective, it is important to bear in mind that every cell stores its value as fully qualified with its full row key, column family, column name, and timestamp on disk. If row and column names are long, the cell value coordinates might become very large and take up more of Apache HBase allotted memory. This can cause severe performance implications, especially if the dataset is large.

Tip: Keep the number of column families small to improve performance and reduce the costs associated with maintaining HFiles on disk.

Apache HBase Configurations

Apache HBase supports built-in mechanisms to handle region splits and compactions. Split/compaction storms can occur when multiple regions grow at roughly the same rate, and eventually split at about the same time. This can cause a large spike in disk I/O because of the compactions needed to rewrite the split regions.

Tip: Rather than relying on Apache HBase to automatically split and compact the growing regions, you can perform these tasks manually.

If you handle the splits and compactions manually, you can perform them in a time-controlled manner and stagger them across all regions to spread the I/O load as much as possible to avoid potential split/compaction storms. With the manual option, you can further alleviate any problematic split/compaction storms and gain additional performance.

Schema Design

A region can run hot when dealing with a write pattern that does not distribute load across all servers evenly. This is a common scenario when dealing with streams processing events with *time series* data. The gradually increasing nature of time series data can cause all incoming data to be written to the same region.

This concentrated write activity on a single server can slow down the overall performance of the cluster. This is because inserting data is now bound to the performance of a single machine. This problem is easily overcome by employing key design strategies such as the following:

- Applying salting prefixes to keys; in other words, prepending a random number to a row.

- Randomizing the key with a hash function.
- Promoting another field to prefix the row key.

These techniques can achieve a more evenly distributed load across all servers.

Client API Considerations

There are a number of optimizations to take into consideration when reading or writing data from a client using the Apache HBase API. For example, when performing a large number of PUT operations, you can disable the auto-flush feature. Otherwise, the PUT operations will be sent one at a time to the region server.

Whenever you use a scan operation to process large numbers of rows, use filters to limit the scan scope. Using filters can potentially improve performance. This is because column over-selection can incur a nontrivial performance penalty, especially over large data sets.

Tip: As a recommended best practice, set the scanner-caching to a value greater than the default of 1, especially if Apache HBase serves as an input source for a [MapReduce](#) job.²⁶

Setting the scanner-caching value to 500, for example, will transfer 500 rows at a time to the client to be processed, but this might potentially cost more in memory consumption.

Compression Techniques

Data compression is an important consideration in Apache HBase production workloads. Apache HBase natively supports a number of compression algorithms that you can enable at the column family level.

Tip: Enabling compression yields better performance.

In general, compute resources for performing compression and decompression tasks are typically less than the overhead for reading more data from disk.

Apache HBase on Amazon EMR

Apache HBase on Amazon EMR is optimized to run on AWS with minimal administration overhead. You still can access the underlying infrastructure and manually configure Apache HBase settings, if desired.

Cluster Considerations

You can resize an Amazon EMR cluster using core and task nodes. You can add more core nodes, if desired. Task nodes are useful for managing the Amazon EC2 instance

²⁶ <http://en.wikipedia.org/wiki/MapReduce>

capacity of a cluster. You can increase capacity to handle peak loads and decrease it later during demand lulls.

Tip: As a recommended best practice, in production workloads you can launch Apache HBase on one cluster and any analysis tools, such as Apache Hive, on a separate cluster to improve performance. Managing two separate clusters ensures that Apache HBase has ready access to the infrastructure resources it requires.

Amazon EMR provides a feature to backup Apache HBase data to Amazon S3. You can perform either manual or automated backups with options to perform full or incremental backups as needed.

Tip: As a best practice, every production cluster should always take advantage of the backup feature available on Amazon EMR.

Hadoop and Apache HBase Configurations

You can use a [bootstrap action](#) to install additional software²⁷ or change Apache HBase or Apache Hadoop configuration settings on Amazon EMR. Bootstrap actions are scripts that are run on the cluster nodes when Amazon EMR launches the cluster. The scripts run before Hadoop starts and before the node begins processing data.

You can write custom bootstrap actions or use predefined bootstrap actions provided by Amazon EMR. For example, you can install Ganglia to monitor Apache HBase performance metrics using a predefined bootstrap action on Amazon EMR.

In summary, whether you are running a managed NoSQL database such as Amazon DynamoDB or Apache HBase on Amazon EMR, or managing your Apache HBase cluster yourself on Amazon EC2 or on-premises, you should take performance optimizations into consideration if you want to maximize performance at reduced costs.

The key difference between a hosted NoSQL solution and managing it yourself is that a managed solution such as Amazon DynamoDB or Apache HBase on Amazon EMR lets you offload the bulk of the administration overhead so that you can focus on optimizing your application.

If you are a developer who is getting started with NoSQL, Amazon DynamoDB or the hosted Apache HBase on the Amazon EMR solution are suitable options, depending on your use case. For developers with in-depth Apache Hadoop/Apache HBase knowledge who need full control of their Apache HBase clusters, the self-managed Apache HBase deployment model offers the most flexibility from a cluster management standpoint.

²⁷ <http://docs.aws.amazon.com/ElasticMapReduce/latest/DeveloperGuide/emr-plan-bootstrap.html>

Conclusion

Amazon DynamoDB lets you offload operating and scaling a highly available distributed database cluster, making it a suitable choice for today's real-time, web-based applications. As a managed service, Apache HBase on EMR is optimized to run on AWS with minimal administration overhead. For advanced users who want to retain full control of their Apache HBase clusters, the self-managed Apache HBase deployment model is a good fit.

Amazon DynamoDB and Apache HBase exhibit inherent characteristics that are critical for successfully processing massive amounts of data. With use cases ranging from batch-oriented processing to real-time data-serving, Amazon DynamoDB and Apache HBase are both optimized to handle large datasets. However, knowing your dataset and access patterns are key to choosing the right NoSQL database for your workload.

Further Reading

For additional information, consult the following sources:

- [Amazon DynamoDB](#)²⁸
- [Apache HBase](#)²⁹
- [Apache Hadoop](#)³⁰
- [Amazon Elastic Compute Cloud \(Amazon EC2\)](#)³¹
- [Amazon Elastic MapReduce \(Amazon EMR\)](#)³²
- [Amazon Kinesis](#)³³
- [Amazon Redshift](#)³⁴
- [AWS Data Pipeline](#)³⁵
- [Amazon CloudWatch](#)³⁶

²⁸ <http://aws.amazon.com/dynamodb/>

²⁹ <http://hbase.apache.org/>

³⁰ <http://hadoop.apache.org/>

³¹ <http://aws.amazon.com/ec2/>

³² <http://aws.amazon.com/elasticmapreduce/>

³³ <http://aws.amazon.com/kinesis/>

³⁴ <http://aws.amazon.com/redshift/>

³⁵ <http://aws.amazon.com/datapipeline/>

³⁶ <http://aws.amazon.com/cloudwatch/>

- [Amazon Identity and Access Management \(Amazon IAM\)](#)³⁷
- [Amazon Simple Storage Service \(Amazon S3\)](#)³⁸
- [AWS Command Line Interface \(Amazon CLI\)](#)³⁹
- [Ganglia](#)⁴⁰
- [Dynamo: Amazon's Highly Available Key-value Store](#)⁴¹
- [Amazon DynamoDB Developer Guide](#)⁴²
- [Amazon EC2 User Guide](#)⁴³
- [Amazon Elastic MapReduce Developer Guide](#)⁴⁴
- [Amazon S3 Developer Guide](#)⁴⁵
- [HBase: The Definitive Guide, by Lars George](#)⁴⁶
- [The Apache HBase™ Reference Guide](#)⁴⁷

³⁷ <http://aws.amazon.com/iam/>

³⁸ <http://aws.amazon.com/s3/>

³⁹ <http://docs.aws.amazon.com/general/latest/gr/GetTheTools.html>

⁴⁰ <http://ganglia.sourceforge.net/>

⁴¹ <http://www.read.seas.harvard.edu/~kohler/class/cs239-w08/decandia07dynamo.pdf>

⁴² <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Introduction.html>

⁴³ <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html>

⁴⁴ <http://docs.aws.amazon.com/ElasticMapReduce/latest/DeveloperGuide/emr-what-is-emr.html>

⁴⁵ <http://docs.aws.amazon.com/AmazonS3/latest/dev/Welcome.html>

⁴⁶ <http://www.hbasebook.com/>

⁴⁷ <http://hbase.apache.org/book.html>