

M Ű E G Y E T E M 1 7 8 2

Budapest University of Technology and Economics
Department of Measurement and Information Systems
Fault Tolerant Systems Research Group

Critical Architectures Laboratory
Spring Semester 2019/2020

Distributed Data Storage in NoSQL Databases – Cassandra

Syllabus
v1.5

Author: Gábor Szárnyas
(szarnyas@mit.bme.hu)

April 8, 2020

1 Overview of Apache Cassandra



Figure 1: The logo of Apache Cassandra

Cassandra is one of the most widely used NoSQL databases [2]. Originally developed by Facebook [9], Cassandra is now an Apache project. Cassandra is a column family database with advanced fault-tolerance mechanisms. It allows the application to balance between availability and consistency by allowing it to tune the consistency constraints. Cassandra is used mainly by Web 2.0 companies, including Digg, Netflix, Reddit, SoundCloud and Twitter. It is also used for research purposes at CERN and NASA [4].

1.1 Data Model

Cassandra uses the *column family data model*. A column family is similar to a table of a relational database: it consists of rows and columns. However, unlike in a relational database's table, the rows do not have to have the same fixed set of columns. Instead, each row can have a different set of columns. This makes the data structure more dynamic and avoids the problems associated with NULL values.

1.2 Serialization and Client Options

Cassandra uses the Apache Thrift framework [3] for serializing data. Like Cassandra, Thrift was originally developed by Facebook, but unlike Cassandra it is still used there.

The most straightforward option is to use the command-line *CQL Shell for Apache Cassandra* provided by the `bin/cqlsh` executable. While new Cassandra projects are encouraged to use CQL (the official Java driver is available at [7]), there are also Thrift-based ORM libraries available (e.g. [8]).

1.3 Sharding in Cassandra

1.3.1 Consistent Hashing

To distribute the data across the cluster, Cassandra uses a partitioner mechanism. The basic partitioners distribute the rows evenly based on their key's hash value using *consistent hashing*.

In a cluster with n nodes, the naïve method for determining the location for a row with key x is computing $h(x) \bmod n$, where $h(x)$ is the hash function. Currently, Cassandra provides partitioners based on the MD5 and the Murmur3 hash functions¹. However, this approach has a serious limitation: if we remove or add nodes to the cluster, we have to recompute the hash values and possibly relocate almost all rows in the cluster. To avoid this, Cassandra uses a special kind of hashing called *consistent hashing*.

¹From Cassandra 1.2, the default partitioner uses the Murmur3 hash function. Although the MD5 hash function has important cryptographic properties, these are irrelevant for distributing data. Hence, the MD5-based partitioner is no longer recommended.

The ring is divided into ranges equal to the number of nodes, with each node being responsible for one or more ranges of the data. Before a node can join the ring, it must be assigned a *token*. The token value determines the node's position in the ring and its range of data. The ring is walked clockwise until it locates the node with a token value greater than that of the row key. Each node is responsible for the region of the ring between itself (inclusive) and its predecessor (exclusive). With the nodes sorted in token order, the last node is considered the predecessor of the first node; hence the ring representation² [6].

For example, consider a simple four-node cluster, where all of the row keys managed by the cluster are in the range of 0 to 100. Each node is assigned a token that represents a point in this range. In this example, the token values are 0, 25, 50 and 75. The first node (with token 0), is responsible for the wrapping range (76+), the second node (with token 25) is responsible for the data range 1 – 25, and so on [6]. Figure 2 shows the token ring for this cluster.

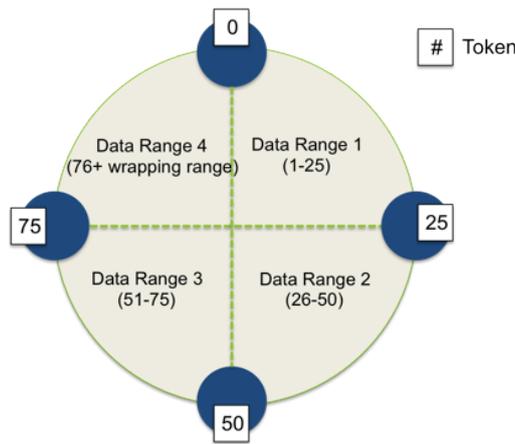


Figure 2: Cassandra's ring for data partitioning [6]

1.3.2 Virtual Nodes

In older versions of Cassandra (prior to version 1.2), you had to calculate and assign a single token to each node in a cluster. Each token determined the node's position in the ring and its portion of data according to its hash value. Although the design of consistent hashing used in older versions (compared to other distribution designs), allowed moving a single node's worth of data when adding or removing nodes from the cluster, it still required substantial effort to do so.

Starting in version 1.2, Cassandra changed this paradigm from one token and range per node to many tokens per node. This paradigm is called *virtual nodes* (vnodes). Vnodes allow each node to own a large number of small partition ranges distributed throughout the cluster. Vnodes also use consistent hashing to distribute data but using them doesn't require token generation and assignment [1].

1.4 Cassandra Query Language

Cassandra provides the Cassandra Query Language (CQL).³ The reference for CQL3 is available in [5]. CQL's goal is to offer a query language similar to the Structured Query Language (SQL) used in most relational databases.

²Note that *consistent* here is different from both the idea of consistency in *data consistency* and in the *ACID* (*atomicity, consistency, isolation, durability*) properties guaranteed by transactions. It refers to the fact that tries to map the same rows to the same machine, even if the number of machines (n) changes over time slightly.

³Not to be confused with *Cypher Query Language* of the Neo4j graph database.

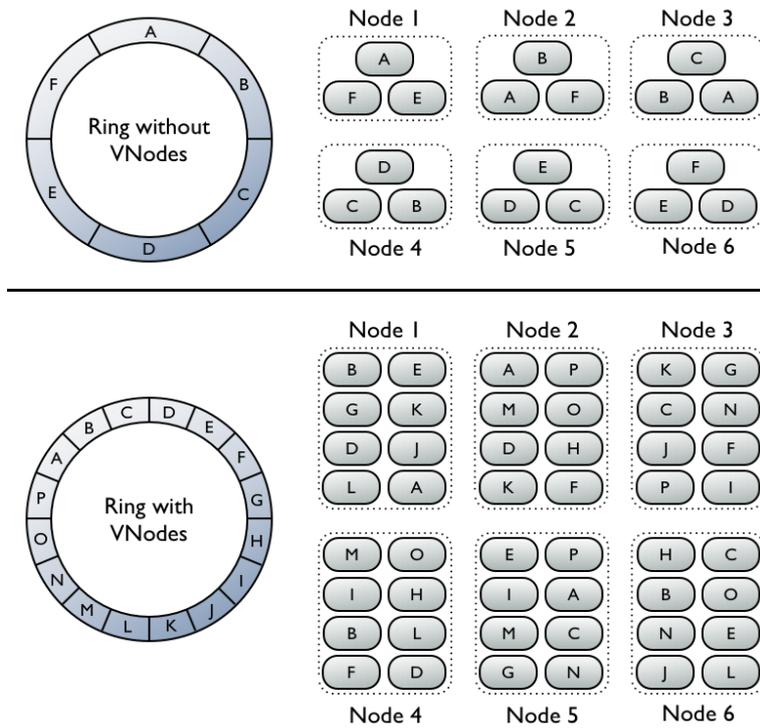


Figure 3: Virtual nodes in Cassandra

CQL allows the user to set the data consistency on a per operation basis. The desired level can be set with the **CONSISTENCY** [**ANY** | **ONE** | **TWO** | **THREE** | **QUORUM** | **ALL** | ...] command. The **CONSISTENCY** command displays the current consistency level.

2 Exercises

2.1 Setting up the Environment

The topology consists of three machines (Figure 4), based on the same VMware virtual machine image. The virtual image contains the following software components:

- Xubuntu 16.04 (64-bit)
- Apache Cassandra 3.10
- ClusterSSH
- tmux

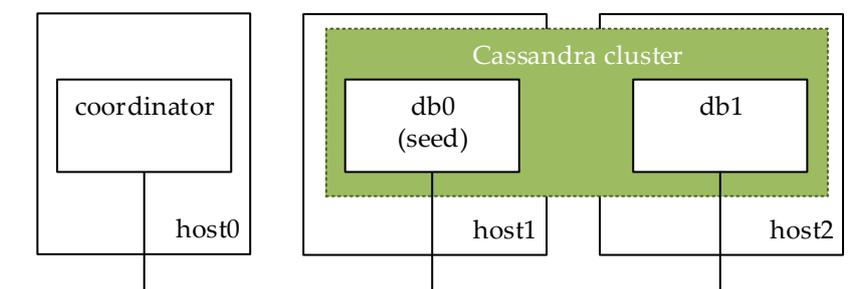


Figure 4: Topology

2.1.1 Hostnames

Set hostnames on the database nodes Because the machines are based on the same image, they have the same hostname (`cassandra-vm`). To distinguish between them, change the hostnames in the `/etc/hostname` file to different ones, e.g. `coordinator`, `db0`, `db1`. Also change the hostname associated with the `127.0.1.1` IP address in the `/etc/hosts` file as well⁴.

Set hostnames on the coordinator To allow easier access (e.g. when using `ssh`), add the hostnames to the *coordinator's* `/etc/hosts` file.

```
10.0.0.100 db0
10.0.0.101 db1
```

Listing 1: `hosts` file

This way, we connect to all database servers using Cluster SSH:

```
cssh db0 db1
```

⁴This is due to a quirk in Debian-based Linux distributions. For details, visit <http://qref.sourceforge.net/quick/ch-gateway.en.html#s-net-dns>.

2.1.2 Running Cassandra on a Single Node

In the provided virtual machines, Cassandra is already set up to run in a single node. This means that the appropriate directory for storing the log files (`logs/`) and the data (`data/`) are created and their permissions are set.

```
sudo service cassandra start
```

To see the Cassandra log, use

```
tail -f /var/log/cassandra/system.log
```

2.1.3 Restarting Cassandra

Restarting Cassandra can be achieved with the following command.

```
sudo service cassandra restart
```

To delete all previous data and restart the server, issue the following command:

```
sudo service cassandra stop && \  
  sudo rm -rf /var/lib/cassandra/data/* && \  
  sudo service cassandra start
```

2.1.4 Running Cassandra on a Cluster of Nodes

The next step is to configure the Cassandra cluster. The main configuration file is stored in the `conf/cassandra.yaml` file. Follow the documentation [5] and configure Cassandra to run in a two node cluster.

2.2 Data Modeling

Create a database that stores road cars. Cars have a *manufacturer*, a *type*. Each car has a *maximum performance* and a *maximum torque* value. Table 1 shows an example.

manufacturer	type	maximum performance	maximum torque
Ford	Focus	100	170
Mercedes	E class	184	270

Table 1: Column family for storing cars

2.3 Replication and Consistency

To test Cassandra's replication schema and consistency models, we define a few steps. The task is to follow these steps, observe the cluster's behaviour and try to explain it.

It is important to note that some of the steps may throw error messages as their expected behaviour. In these cases, try to determine the reason of the error message and proceed to the next step.

Between the steps, you should delete all previous data (from the `/var/lib/cassandra/` directory) and restart the server.

2.3.1 Network Partition without Replication

1. Create a keyspace with the *replication factor* set to 1.
2. Create a column family and fill it with data.
3. Query the data from both db0 and db1.
4. Disconnect db0 from the network. The easiest way to do this is to disable the network interface in VMware.
5. Query the data from db1.
6. Modify the data from db1, e.g. add new rows to the column family.
7. Reconnect db0 and query the data from db0.

2.3.2 Network Partition with Replication and Weak Consistency

1. Create a keyspace with the *replication factor* set to 2.
2. Create a column family and fill it with data.
3. Query the data from both db0 and db1.
4. Disconnect db0 from the network.
5. Query the data from db1.
6. Modify the data from db1, e.g. add new rows to the column family.
7. Disconnect db1 from the network.
8. Reconnect db0 from the network.
9. Also query and modify from db0. Modify the same rows as previously so it will *conflict* with the previous write operations.
10. Reconnect db1 to the network.
11. Query the data from both db0 and db1.

2.3.3 Network Partition with Replication and Quorum Consistency

1. Create a keyspace with the *replication factor* set to 2.
2. Create a column family, fill it with data, query it and disconnect db0.
3. On db1, set the consistency to **QUORUM**.
4. Query and modify the data.
5. Observe how Cassandra interprets the Quorum Consistency.

2.4 Advanced Data Modeling

Cars have different powertrains (e.g. hybrid systems [10]). Each type can be described with different parameters:

- Internal combustion engine: fuel type, displacement, maximum torque, maximum power
- Electric motor: maximum torque, maximum power
- Both: all of the above and the combined maximum torque and power values

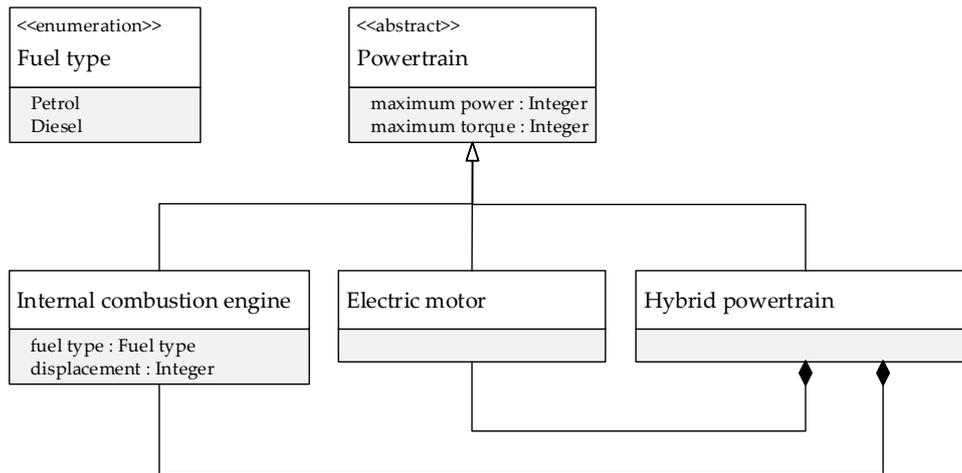


Figure 5: UML class diagram showing different powertrain types

The class hierarchy for different powertrain types are shown on Figure 5.

1. Extend the `cars` column family to store the powertrain of each car.
2. Write a query that collects the cars with an internal combustion engine.
3. Write a query that collects the cars with an internal combustion engine or an electric motor.

3 Additional Exercises

Cars have different transmissions systems. The main categories are the following:

- Automatic
 - fixed ratios: list of gear ratios, type (e.g. semi-automatic, dual-clutch, etc.)
 - continuously variable (CVT): gear ratio range
- Manual: list of gear ratios
- No gearbox: gear ratio

Figure 6 shows the class hierarchy for different transmission types.

Extend the `cars` column family to store the transmission for each car. To store the list of gear ratios, use CQL's `list` collection type.

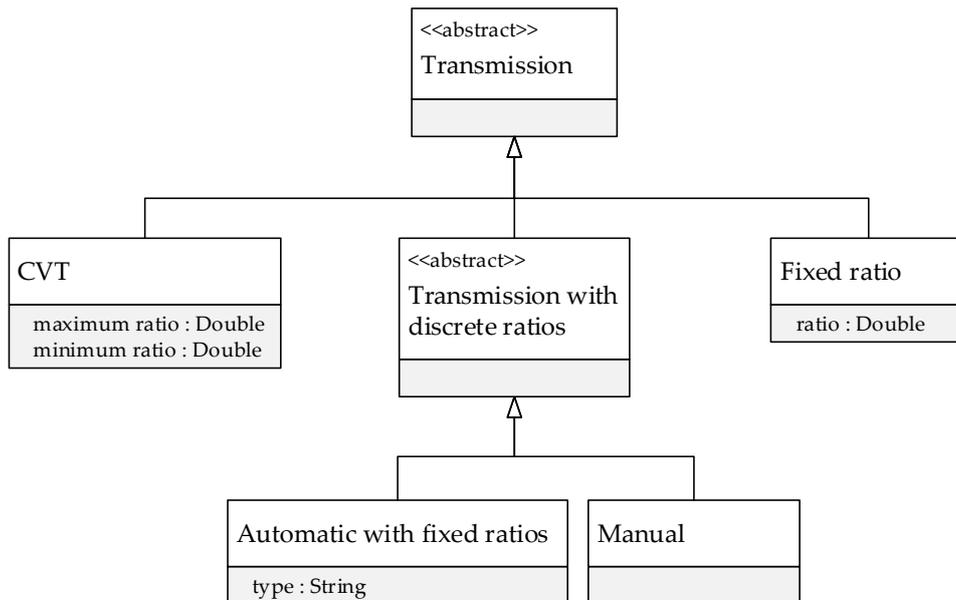


Figure 6: UML class diagram showing different transmission types

4 Tips

- Set the broadcast address to the IP address of the machine.
- The `ssh-keygen -t rsa` and `ssh-copy-id` commands come in handy.
- VM: Bridged. To get a new IP address (in case of collision), change the adapter to something else (e.g. NAT) and change it back to Bridged.
(Fortunately, Disconnect & Reconnect does not assign a new IP address.)
- `Ctrl` + `Alt` + `T` launches a terminal emulator application. You can open new tabs with the `Ctrl` + `↑` + `T` hotkey. `Ctrl` + `D` sends an end-of-transmission character, which closes the terminal.
- Use `ifconfig` to determine a machine's IP address.
- The most important hotkeys for `mcedit` are the following:
 - `F7` – search
 - `↑` + `F7` – find next
 - `Ctrl` + `0` – exit
- For `nano`, use `Ctrl` + `W` for search.
- If you use `vim`, you should already know the hotkeys ;-).
- Use the `watch` command to repeatedly run a tool, e.g.
`watch -n 1 bin/nodetool status`
- To check if Cassandra is running in the background, run:
`ps aux | grep cassandr[a]`
To repeat the check later, simply run `!ps`.
- Useful Bash hotkeys: Use `↩` and `↑` a lot. `Ctrl` + `A` jumps back to the start of the line, `Ctrl` + `U` deletes the current line.

References

- [1] Virtual nodes in Cassandra 1.2. <http://www.datastax.com/dev/blog/virtual-nodes-in-cassandra-1-2>, December 2012.
- [2] Apache Cassandra. <http://cassandra.apache.org/>, October 2013.
- [3] Apache Thrift. <http://thrift.apache.org/>, October 2013.
- [4] Planet Cassandra – Companies. <http://planetcassandra.org/Company/ViewCompany>, October 2013.
- [5] Apache Cassandra 3.9 documentation. <http://cassandra.apache.org/doc/3.9/>, March 2017.
- [6] Nick Bailey. Introduction to Cassandra architecture. <https://www.slideshare.net/nickmbailey/introduction-to-cassandra-architecture>, June 2016.
- [7] DataStax. Java Driver for Apache Cassandra. <https://github.com/datastax/java-driver>, October 2013.
- [8] impetus-opensource. Kundera. <https://github.com/impetus-opensource/Kundera>, October 2013.
- [9] Avinash Lakshman and Prashant Malik. Cassandra: a decentralized structured storage system. *SIGOPS Oper. Syst. Rev.*, 44(2):35–40, April 2010.
- [10] Wikipedia. Hybrid vehicle drivetrain. http://en.wikipedia.org/wiki/Hybrid_vehicle_drivetrain, October 2013.