

Managing a Passive DNS Database Using Apache Cassandra

Posted on November 18, 2024

Abstract

Apache Cassandra is a highly scalable, distributed NoSQL database designed for handling massive volumes of data across many commodity servers without a single point of failure. Its decentralized nature and robust architecture make it particularly well-suited for applications that require high availability, fault tolerance, and horizontal scalability. Cassandra is engineered to handle very large datasets, supporting billions of records with ease, making it an ideal choice for organizations dealing with large-scale, real-time applications such as time-series data, IoT data, and customer logs. Through its use of a partitioned architecture and the ability to add nodes seamlessly as data grows, Cassandra offers an efficient means of managing big data with low latency and high throughput.

One of Cassandra's key features for handling billions of records is its ability to distribute data across multiple nodes in a cluster. Data is partitioned based on the partition key, which ensures that no single node becomes a bottleneck. Each partition can be replicated across multiple nodes to provide fault tolerance and data redundancy, ensuring high availability even in the event of hardware failures. Cassandra's write-optimized architecture, combined with its eventual consistency model, allows it to handle large volumes of write-heavy workloads while maintaining performance at scale. As a result, it excels in scenarios where high write throughput is required, such as real-time analytics, logs, and sensor data processing.

In addition to its scalability and high availability, Cassandra also offers flexible data modeling options and the ability to efficiently query large datasets. The database's support for tunable



consistency levels allows it to balance between consistency and performance, depending on the application's needs. By utilizing techniques like compaction and write path optimizations, Cassandra minimizes disk I/O and improves query response times, even when managing billions of records. These features, combined with its distributed nature, make Cassandra an efficient and reliable solution for organizations that need to store and process large-scale datasets in a cost-effective and scalable manner.

1. Introduction

As organizations increasingly generate vast amounts of data, the need for databases capable of handling large-scale, distributed storage has never been more critical. Apache Cassandra, an open-source NoSQL database, has emerged as one of the leading solutions for managing petabytes of data across distributed clusters. Companies such as Discord use Cassandra to manage trillions of messages. Unlike traditional relational databases, which can struggle with scaling horizontally and managing write-heavy workloads, Cassandra is designed to provide high availability, fault tolerance, and horizontal scalability without compromising performance. Its decentralized architecture, which eliminates single points of failure, makes it particularly well-suited for applications requiring the management of billions of records, such as real-time analytics, sensor data, and large-scale log processing. In this context, Cassandra's ability to efficiently distribute data across nodes, handle high-velocity writes, and scale with ease, positions it as a powerful tool for organizations aiming to build resilient, data-driven systems at a global scale. This makes it the ideal solution to manage large data sets such as WHOIS DNS databases.

2. Prerequisites

To set up Apache Cassandra, we will be using Ubuntu Linux for your DNS database, here are the initial requirements and setup steps:

2.1 System Requirements



- Operating System: Ubuntu Linux 20.04 or later (recommended).
- Java Development Kit (JDK): Cassandra requires Java. OpenJDK 8 or 11 is generally recommended for stability.
- **RAM**: For production environments with large datasets, a minimum of 32GB RAM is recommended, with higher amounts depending on the dataset size.
- **Disk Space**: A high-speed SSD or NVMD is preferable due to Cassandra's I/O needs. Plan for significant storage depending on your data volume.
- **Processor**: Multiple cores are recommended, ideally with 4+ cores for handling large datasets.
- **Network**: Ensure a stable network setup, especially if you plan to run a Cassandra cluster across multiple nodes.

2.2 Software Requirements

- Cassandra: Apache Cassandra 3.11 or later, depending on compatibility and feature needs.
- **cqlsh**: Cassandra's command-line tool, which comes with the installation.
- Python 2.x or 3.x: Needed for cqlsh to function properly.
- Apache Spark (optional but recommended): For data ingestion, transformation, and ETL tasks if your .csv files are especially large.

3. Installing Cassandra

Add the Cassandra repository and install it:



echo "deb http://www.apache.org/dist/cassandra/debian 311x main" | sudo tee -a /etc/apt/sources.list. curl https://www.apache.org/dist/cassandra/KEYS | sudo apt-key add sudo apt-get update sudo apt-get install cassandra

Verify Cassandra installation:

cassandra -v

Start the Cassandra service:

sudo systemctl start cassandra

After setting up the base requirements and installing Apache Cassandra, the next steps are to configure Cassandra, create the necessary schema for your DNS database, and import data from .csv files.

Configure Cassandra

Open the configuration file to customize Cassandra settings if needed:

sudo nano /etc/cassandra/cassandra.yaml

Key configuration changes to consider:



cluster_name: 'WHOISXMLAPI_DNS_CLUSTER'

data_file_directories:

- /var/lib/cassandra/data

If you do make changes, restart the service:

sudo systemctl restart cassandra

4. Creating a Keyspace and Table Structure

Open cqlsh to interact with Cassandra:

cqlsh

Create a keyspace for the DNS data:

CREATE KEYSPACE dns_data WITH replication = {'class': 'SimpleStrategy', 'replication_factor': 3};

Create a table to match the .csv structure:

USE dns_data;

CREATE TABLE dns_records (domain TEXT PRIMARY KEY, last_update TIMESTAMP, ip_addresses TEXT,



wildcard TEXT, active TEXT

);

Note: domain is set as the primary key here, assuming domains are unique in your dataset.

5. Loading DNS Data from .csv Files into Cassandra

To import large .csv files, you can use the COPY command or the cassandra-loader tool.

Option 1: Using the COPY Command (Simpler for Smaller Data Imports)

• Format your .csv file to ensure it matches the column structure of dns_records.

Use the COPY command:

COPY dns_records (domain, last_update, ip_addresses, wildcard, active) FROM '/path/to/your/file.csv' WITH HEADER = TRUE;

Note: Large imports may cause performance issues with COPY. Consider a cassandra-loader for faster loading.

Option 2: Using cassandra-loader for Larger Imports

• Install cassandra-loader if not available.

sudo apt-get install cassandra-loader

Load the .csv data



cassandra-loader -f /path/to/your/file.csv -host localhost -schema "dns_data.dns_records(domain, last

6. Validating the Data Import

After loading, query the table to confirm records have been inserted:

SELECT * FROM dns_records LIMIT 5; SELECT COUNT(*) FROM dns_records;

7. Indexing for Performance (Optional)

For frequent searches on non-primary key fields, you may want to create secondary indexes:

CREATE INDEX ON dns_records (last_update); CREATE INDEX ON dns_records (ip_addresses);

8. Automating Import for Periodic Updates (Optional)

If you need to update data regularly, consider scripting the import process with a shell script or setting up a cron job.

9. Searching

To search for any domain for example that ends in 'microsoft.com', you can use a query with a wildcard filter if you have Cassandra 3.4 or later. However, since Cassandra does not support



LIKE or regex queries natively, we'll need a different approach. Here are two options: Full table scan, or pre-process and reverse domain column. Since this is a fairly large dataset, we will focus on using pre-processing.

Using a Full-Table Scan with ALLOW FILTERING

For more efficient searching, you could add a reverse_domain column where the domain is stored in reverse order (e.g., moc.tfosorcim.www). Then you can query based on the reversed domain and use a prefix match, which can be faster.

Modify Table Structure: Add a new column for the reversed domain:

ALTER TABLE dns_records ADD reverse_domain TEXT;

Populate reverse_domain: Write a script to reverse each domain and update the column.

Here's a simple example in Python:

from cassandra.cluster import Cluster

```
cluster = Cluster(['localhost'])
session = cluster.connect('dns_data')
```

```
rows = session.execute("SELECT domain FROM dns_records;")
for row in rows:
```

reverse_domain = row.domain[::-1]

session.execute("UPDATE dns_records SET reverse_domain = %s WHERE domain = %s", (reverse

Create an index on reverse_domain:

CREATE INDEX ON dns_records (reverse_domain);



Query Using reverse_domain: Now, to search for any domain ending with microsoft.com, you can reverse it and query:

SELECT * FROM dns_records WHERE reverse_domain LIKE 'moc.tfosorcim%' ALLOW FILTERING;

This approach makes it faster to search for domains that end with a specific pattern.

Cassandra Dictionary

Keyspace

- The top-level container in Cassandra, similar to a database in relational databases.
- Defines important configuration settings like replication factor, which determines data redundancy across nodes
- Example:

Table (Column Family)

- A table in Cassandra organizes data into rows and columns, similar to a relational database table, but structured for distributed storage.
- Tables are defined within a keyspace and contain rows of records.

Primary Key

- Uniquely identifies each row in a table and determines the data's storage distribution.
- Consists of a Partition Key (mandatory) and optionally Clustering Columns (optional).

Partition Key



- The first part of the primary key that determines on which node data is stored.
- Helps evenly distribute data across the cluster by hashing the partition key value.

Clustering Columns

- Optional additional columns in the primary key that define the sort order within each partition.
- Useful for specifying the order of rows when querying and reading data.

Clustering

- Defines how data is ordered within a partition based on clustering columns.
- Helps to organize data in a meaningful way for faster query performance.

Replication Factor

- Specifies how many copies of the data are stored across the cluster for redundancy.
- Higher replication factors increase fault tolerance but require more storage.

Consistency Level

• Defines the number of nodes that must acknowledge a read or write operation for it to be considered successful.

Node

- An individual server in the Cassandra cluster, each of which holds a portion of the dataset.
- Nodes are distributed and coordinated to store data redundantly based on the replication factor.

Cluster



- A collection of nodes that work together to form a distributed database.
- Nodes within a cluster replicate data and provide high availability.

Token and Token Range

- Each node is assigned a token or range of tokens to determine its data ownership in the ring architecture of Cassandra.
- Token ranges are calculated based on the partition key.

Ring Architecture

- Cassandra's internal structure where nodes are organized in a logical ring, with data distributed across nodes based on partition key hashes.
- This enables linear scaling as new nodes can be added to the ring without reorganizing the entire dataset.

CQL (Cassandra Query Language)

• The language for interacting with Cassandra, similar to SQL, but designed for Cassandra's distributed architecture.

Data Center

- A logical grouping of nodes, often corresponding to actual physical data centers.
- Used for setting replication policies that ensure fault tolerance across geographic locations.

Hinting

• A mechanism where a coordinator node "hints" to nodes that were down during a write operation to catch up with missing data once they are back online.

Tombstone



- A marker used to signify deleted data, which will be permanently removed during the next compaction.
- Cassandra uses tombstones to handle deletions in a distributed system.

Compaction

• A maintenance process that consolidates SSTables, removes tombstones, and merges data, optimizing disk space and read performance.

SSTable (Sorted String Table)

- Immutable files where Cassandra stores data on disk, generated during compaction.
- SSTables are sequentially ordered, aiding efficient read and write operations.

Relevant Links for Apache Cassandra

Official Downloads and Documentation

- Apache Cassandra Download Page
 Provides official downloads for the latest stable versions of Apache Cassandra.
- Apache Cassandra Documentation
 Comprehensive documentation on installation, configuration, CQL commands, and usage examples.

Setup and Installation Guides

Cassandra on Ubuntu

DigitalOcean's Guide to Installing Cassandra on Ubuntu A detailed tutorial covering the installation and setup of Cassandra on Ubuntu.

Tutorials and Learning Resources



Cassandra Tutorials

A collection of getting-started guides and tutorials from the Cassandra documentation.

Example Code and Projects

- Cassandra Code Samples
 GitHub Cassandra Code Samples
 Example code directly from the Apache Cassandra repository, useful for learning by example.
- Python and Cassandra (cassandra-driver)

Datastax Python Driver

The official Python driver for interacting with Cassandra using Python, with usage examples and documentation.

Spring Data Cassandra (Java)
 Spring Data Cassandra Guide

Enables using Cassandra with Java's Spring framework, providing integration and ORM capabilities.

Tools and Utilities

• cqlsh

The official command-line interface (CLI) for Cassandra, installed by default. Basic usage information is in the CQLSH documentation.

• Cassandra Cluster Manager (CCM)

GitHub - Cassandra Cluster Manager

A tool for creating and managing local Cassandra clusters, ideal for testing multiple nodes on a single machine.

Cassandra Reaper - Repair Tool

A repair management tool for Cassandra to help with database maintenance and performance optimization.



• cassandra-loader

GitHub - cassandra-loader

An optimized bulk loader for importing large .csv files into Cassandra, useful for loading big datasets.

DataStax Studio

A visual tool to interact with Cassandra, explore data models, and write queries, especially helpful for beginners.

Community and Support

- Apache Cassandra Mailing Lists Join the Cassandra community mailing lists for user support, announcements, and discussions.
- Cassandra User and Dev Forums
 Cassandra Forums on Stack Overflow
 Get support from the developer community by asking questions and browsing solutions.

Contact Information

WHOIS, Inc. Sales: sales@whoisxmlapi.com Support: service.desk@whoisxmlapi.com Website: www.whoisxmlapi.com