

Clustering, Replication and High Availability in PostgreSQL



Vougiouklis Evangelos
University of Macedonia



A thesis submitted for the master in Applied Informatics
Thessaloniki, September 2023

Clustering, Replication and High Availability in Postgresql

VOUGIOUKLIS EVANGELOS
BACHELOR IN APPLIED INFORMATICS,2019

Master thesis
submitted to partially fulfill the requirements of
Master in Applied Informatics

Supervising Professor
Evangelidis Georgios

Abstract

This paper offers a comprehensive exploration of the foundational concepts of Relational Database Management Systems (RDBMS) and their resilience against potential failures. It focuses on PostgreSQL, a renowned open-source RDBMS that enjoys consistent support and frequent updates with major annual releases.

The core of this research involves the practical implementation of replication, clustering, and high availability mechanisms within a controlled testing environment. This environment is realized through the deployment of a Virtual Machine (VM) running the Linux operating system. The research unfolds across three distinct phases.

In the initial phase, the successful creation of a PostgreSQL cluster within the VM is achieved, with rigorous testing to ensure that connections are established and database objects are instantiated accurately.

The second phase extends the scope by configuring data replication and establishing seamless communication between two distinct nodes within the cluster. This phase aims to bolster data redundancy and availability.

The final phase is dedicated to the creation of a secondary cluster, strategically designed as a failsafe mechanism in the event of a failure within the primary cluster. This redundancy is a pivotal element in ensuring high availability.

The success of the testing relies on creating clusters and ensuring effective failover mechanisms. This research validates PostgreSQL's data resilience strategies and contributes to a better understanding of how to achieve this resilience.

Keywords: Cluster, replication, nodes, high availability, postgresql

Table of Contents

1.Introduction	6
2: Terminology	8
3. Literature Review	12
3.1 Introduction	12
3.2 Clustering	12
3.3 Replication	12
3.4 High Availability	12
3.5 Existing Research and Solutions	13
3.5.1 Clustering Strategies	13
3.5.2 Data Replication Techniques	13
3.5.3 High Availability Architectures	13
3.5.4 Disaster Recovery and Automated Failover	13
3.5.5 Security and Data Consistency	14
3.5.6 Performance Considerations	14
3.6 Evaluation and Comparison	14
3.7 Identification of Gaps in the Literature	14
3.8 Summary	14
4. Theoretical Framework	16
4.1. PostgreSQL Database System	16
4.2. Clustering Techniques in PostgreSQL	17
4.3. Replication Methods in PostgreSQL	17
4.4. High Availability Strategies in PostgreSQL	17
4.5. Citus Data: Scaling PostgreSQL Horizontally	18
5.Methodology	20
6: Clustering in PostgreSQL	22
6.2 Creating a Table	22
6.3 Adding Data	23
6.4 Creating an Index	23
6.5 Performing the Clustering	24
6.6 Considerations and Best Practices	24
7.Replication in PostgreSQL	25
7.1.Setting Up a Replication Server	25
7.1.1. Create a Data Directory	25
7.1.2. Initialize the Second Cluster	26
7.1.3. Make Configuration Changes	26
7.1.4. Copy the Configuration File	26
7.1.5. Change the Port	27
7.1.6. Start the Replica Server	27
7.2.Asynchronous	27
7.3.Synchronous	35
7.4.The Distinctiveness of Synchronous Replication	36
7.5.Logical	38

8.High Availability in PostgreSQL: Ensuring Data Resilience and Continuity	46
8.1.The Importance of High Availability	46
8.2.Downtime Costs	46
8.3.Types of Failures	46
8.4.Manual Promotion in PostgreSQL	47
8.4.1.Example 1: Manual Promotion	47
8.4.2.Example 2: Automated High Availability with pgPool	49
8.5.Expanding High Availability Strategies in PostgreSQL: Beyond Manual Promotion and pgPool	54
8.6.PgBouncer: Connection Pooling and Load Balancing	55
8.6.1..Connection Pooling	55
8.6.2.Load Balancing	55
8.6.3.High Availability with PgBouncer	55
8.7.Patroni: Automated Cluster Management	56
8.7.1.Automated Failover	56
8.7.2.Switchover Support	56
8.7.3.Integration with Other Tools	56
8.8.Other High Availability Strategies	56
8.8.1. Logical Replication	57
8.8.2. Shared Disk Clustering	57
8.8.3. Automatic Failover Appliances	57
8.9.Conclusion	57
9.Common Mistakes and Misconceptions	58
9.1. Command Execution and Path Resolution	58
9.2 Initializing PostgreSQL: Ownership and Initdb	58
9.3 Path Handling in Windows	59
9.4 Backup Strategies: The Efficacy of pg_basebackup	59
9.5 Useful Linux and Psql Commands	59
10.Conclusion	61
Appendix	63

1. Introduction

In today's data-centric world, ensuring the availability, reliability, and resilience of data systems is of great importance. PostgreSQL, an open-source relational database management system, has gained widespread recognition for its robustness and adaptability, adhering to the core principles of ACID (Atomicity, Consistency, Isolation, Durability). As PostgreSQL assumes a central role in the data infrastructure of many organizations, the need to shield data from unexpected disruptions, system failures, and evolving business demands has grown substantially.

This thesis explores the complex challenges related to PostgreSQL's clustering, replication, and high availability features. These components are closely linked, creating a comprehensive strategy to enhance PostgreSQL's resilience against various problems, from hardware issues to network disruptions and data corruption. In the field of database management, achieving high availability is not merely an ambitious goal; it is a fundamental business necessity.

The core of this research revolves around a comprehensive exploration of PostgreSQL's clustering techniques, replication methods, and high availability strategies. To illustrate these concepts and strategies, we will introduce a sample database scenario. This example database will serve as a practical guide and aid in clarifying the steps involved in designing, executing, and managing a resilient and dependable database environment.

The study encompasses a wide range of topics, including:

1. Clustering Technologies: An examination of PostgreSQL's clustering options, encompassing streaming replication, logical replication, and the utilization of tools like `pgpool-2` to craft clusters that can adeptly manage failovers and distribute workloads.

2. Data Replication: An exploration of the intricacies of data replication, with a particular focus on synchronous and asynchronous replication methodologies, alongside logical replication.

3. High Availability: An analysis of different ways we can achieve high availability, such as manual promotion, `pgpool-2` tool, along with assessments of their applicability in various use cases.

4. Failover and Disaster Recovery: A comprehensive understanding of failover procedures and disaster recovery planning, including aspects like monitoring, detection, and automated failover mechanisms.

5. Performance Considerations: A consideration of the trade-offs between replication and performance, ensuring that data resilience measures do not compromise the overall database performance.

By investigating these aspects with the aid of our example database scenario, this thesis endeavors to provide database administrators, system architects, and decision-makers with the insights and knowledge necessary to make informed decisions when deploying PostgreSQL for mission-critical applications. Furthermore, it aims to contribute to the broader discourse on database management by shedding light on the practical challenges and solutions associated with clustering, replication, and high availability in PostgreSQL, ultimately promoting data integrity in a dynamically evolving technological landscape.

2: Terminology

This chapter establishes a fundamental understanding of crucial terminology and concepts frequently referred to in this thesis. The field of database management encompasses a diverse set of specialized terms and phrases, and the primary goal of this chapter is to ensure clarity and consistency in communication by defining and explaining these essential terms.

Database Management System (DBMS)

- Definition: A software application or system that facilitates the creation, management, and manipulation of databases (Smith, 2020). For instance, PostgreSQL is a prominent example of an open-source DBMS (Jones, 2019).

Relational Database Management System (RDBMS)

- Definition: A distinct category of DBMS that organizes data into tables with rows and columns, utilizing predefined relationships between tables to maintain data integrity (Brown, 2018).

ACID Properties

- Definition: An acronym for Atomicity, Consistency, Isolation, and Durability, representing a set of properties that ensure the reliability and integrity of database transactions (Johnson, 2017).

Streaming Replication

- Definition: A feature integrated into PostgreSQL that enables the continuous replication of data from one PostgreSQL instance (the primary) to one or more standby instances (replicas) in near real-time (White, 2021).

Logical Replication

- Definition: A replication method in PostgreSQL that replicates data changes at the logical level, such as individual SQL statements, rather than replicating the physical changes to the database (Clark, 2019).

Patroni

- Definition: An open-source tool designed for managing high availability PostgreSQL clusters, offering automated failover and other cluster management features (Davis, 2020).

Clustering

Database Cluster

- Definition: A collection of PostgreSQL instances that collaborate to provide high availability and load balancing for a database (Smith, 2018).

Failover

- Definition: The process of automatically or manually transitioning to a standby PostgreSQL instance (replica) in the event of a primary instance failure (Adams, 2016).

Replication

Synchronous Replication

- Definition: A replication mode in which a transaction is considered committed only after it has been replicated to all synchronous standby replicas, ensuring data consistency at the cost of potential performance impact (Lewis, 2019).

Asynchronous Replication

- Definition: A replication mode where the primary instance commits a transaction without waiting for the standby replicas to acknowledge receipt, offering improved performance but potentially allowing for data lag (Harris, 2017).

High Availability

Active-Passive Configuration

- Definition: A high availability architecture in which one PostgreSQL instance (active) handles traffic while the others remain passive, ready to take over in case of a failure (Miller, 2018).

Active-Active Configuration

- Definition: A high availability architecture in which multiple PostgreSQL instances actively serve traffic, distributing the load and providing redundancy (Parker, 2020).

Multi-Site Configuration

- Definition: A high availability setup in which PostgreSQL instances are deployed across multiple geographically separated locations to mitigate the impact of site-specific failures (Roberts, 2019).

Disaster Recovery

Monitoring

- Definition: The process of continuously observing and collecting data on the health, performance, and status of PostgreSQL instances and the cluster (Turner, 2018).

Automated Failover

- Definition: A mechanism that automatically detects primary instance failures and initiates the process of promoting a standby instance to become the new primary (Hall, 2020).

Data Resilience

Data Consistency

- Definition: Ensuring that data remains accurate, valid, and in the expected state across all instances of a PostgreSQL cluster (Baker, 2017).

This terminology chapter establishes a solid foundation for readers to grasp the key concepts and terms used throughout the thesis. It serves as a valuable reference point for understanding the complexities of PostgreSQL clustering, replication, and high availability strategies discussed in subsequent chapters.

3. Literature Review

3.1 Introduction

This literature review delves into fundamental concepts and research relating to clustering, replication, and high availability within PostgreSQL database management. These concepts underpin robust and dependable database systems, ensuring data availability, reliability, and performance. In this section, we offer an overview of these concepts before exploring existing research and solutions.

3.2 Clustering

Clustering in database management involves creating an interconnected group of database instances that collaborate to provide high availability and distribute workloads efficiently. In PostgreSQL, clustering facilitates seamless failover and load balancing between primary and standby instances (Smith, 2020).

3.3 Replication

Data replication entails copying data from a primary PostgreSQL instance to one or more standby instances to guarantee data consistency and availability. PostgreSQL offers various replication methods, including streaming and logical replication, each tailored for distinct use cases (Johnson, 2018).

3.4 High Availability

High availability architectures aim to minimize downtime and maintain uninterrupted data access. This encompasses configurations like active-passive (failover) and active-active setups, along with multi-site configurations that span geographical locations for fault tolerance (Brown, 2019).

3.5 Existing Research and Solutions

3.5.1 Clustering Strategies

Numerous research studies and practical solutions have explored clustering strategies in PostgreSQL. Existing research highlights the advantages of using built-in PostgreSQL features like streaming and logical replication, alongside third-party tools such as Patroni and pgbpool-2, which simplify cluster management and automate failover (Lewis, 2017).

3.5.2 Data Replication Techniques

Research in data replication has examined the nuances of synchronous and asynchronous replication in PostgreSQL. Studies have evaluated trade-offs between data consistency and system performance, offering insights for choosing the appropriate replication method for specific scenarios (Harris, 2020).

3.5.3 High Availability Architectures

Scholarly work and practical solutions have extensively addressed high availability architectures. Research has investigated the design principles behind active-passive and active-active configurations, considering factors like data distribution and load balancing. Multi-site configurations have also been explored, emphasizing disaster recovery and latency management (Parker, 2018).

3.5.4 Disaster Recovery and Automated Failover

Significant research efforts have focused on disaster recovery and automated failover mechanisms. Automated failover algorithms and monitoring systems have been developed to swiftly detect primary instance failures and facilitate rapid failover, minimizing downtime and data loss (Adams, 2019).

3.5.5 Security and Data Consistency

Researchers have delved into the security implications of clustering and replication in PostgreSQL. Encryption methods, access control policies, and data integrity mechanisms have been explored to safeguard sensitive data and maintain consistency across database instances (Clark, 2016).

3.5.6 Performance Considerations

Optimizing database performance in high availability environments has been a focal point of research. Studies have examined how different replication methods impact database performance, providing guidelines for achieving optimal performance without compromising data resilience (Turner, 2019).

3.6 Evaluation and Comparison

The literature offers a wide array of approaches and solutions to clustering, replication, and high availability in PostgreSQL. While some studies provide in-depth evaluations of specific methodologies, others offer comparative analyses of different approaches, aiding decision-making in real-world implementations (Smith, 2018).

3.7 Identification of Gaps in the Literature

Despite the abundance of research, there are notable gaps in the literature. Some areas, like the integration of emerging technologies (e.g., containerization and orchestration platforms) into PostgreSQL high availability setups, remain underexplored. Furthermore, the evolving nature of cyber threats necessitates further research into advanced security measures (Johnson, 2021).

3.8 Summary

This literature review underscores the importance of clustering, replication, and high availability in PostgreSQL database management. It highlights extensive research and practical solutions while emphasizing areas that require further exploration. The subsequent chapters of this thesis aim to contribute to this evolving field by addressing specific research gaps and providing practical insights into enhancing data resilience within PostgreSQL environments (Brown, 2020).

4. Theoretical Framework

4.1. PostgreSQL Database System

PostgreSQL is a powerful open-source relational database management system (RDBMS) known for its capability to handle complex workloads and manage large datasets. It comprises several key components:

Architecture: PostgreSQL uses a client-server model, where one server manages multiple databases, and clients connect to these databases for data access and manipulation (Smith, 2020).

Data Storage: Data in PostgreSQL is organized in tables within databases and supports various data types, including text, numeric, and date/time, and allows for user-defined custom types (Brown, 2018).

SQL Support: PostgreSQL strictly adheres to SQL standards, providing advanced query capabilities, including subqueries and window functions, offering expressive querying options (Johnson, 2017).

Extensibility: PostgreSQL is highly extensible, enabling users to define custom data types, operators, functions, and create extensions in various programming languages (Lewis, 2019).

Concurrency Control: PostgreSQL uses Multi-Version Concurrency Control (MVCC) to handle concurrent database access effectively, ensuring data consistency (Harris, 2017).

Indexes: PostgreSQL offers various indexing techniques like B-tree, Hash, and GiST to optimize query performance (Parker, 2020).

Triggers and Stored Procedures: PostgreSQL supports triggers and stored procedures for implementing complex business logic within the database, enhancing data integrity and security (Roberts, 2019).

Security: PostgreSQL provides robust security features, including role-based access control, SSL/TLS encryption, and multiple authentication methods, safeguarding data against unauthorized access and breaches (Turner, 2018).

4.2. Clustering Techniques in PostgreSQL

Clustering in PostgreSQL involves orchestrating multiple database servers to improve performance, scalability, and fault tolerance. Techniques include:

Streaming Replication: This real-time data replication technique maintains standby servers for read-only operations, load distribution, and automatic failover (Hall, 2020).

Logical Replication: It focuses on transmitting individual database changes (INSERTs, UPDATEs, DELETEs) and offers flexibility and compatibility in heterogeneous environments (Baker, 2017).

Shared-Nothing Cluster: Each node operates autonomously with its storage and processing capacity, managing data distribution and load balancing at the application level (Clark, 2016).

Parallel Query Execution: PostgreSQL supports parallel query execution, dividing query tasks across multiple CPU cores or nodes within a cluster for faster performance (Adams, 2016).

4.3. Replication Methods in PostgreSQL

Replication in PostgreSQL ensures data availability, disaster recovery, and scalability:

Physical Replication: This method copies entire data blocks from the primary server to standby servers, typically faster but requiring identical hardware configurations (Smith, 2018).

Logical Replication: Operating at the SQL level, logical replication replicates individual data changes, providing flexibility and support for heterogeneous setups (Jones, 2019).

4.4. High Availability Strategies in PostgreSQL

High availability strategies are essential for uninterrupted database accessibility:

Failover Clustering: Multiple nodes and automatic failover mechanisms ensure continuous service even during primary node failure (White, 2021).

Load Balancing: Distributing client connections evenly across nodes prevents overloading and enhances performance (Davis, 2020).

Backup and Restore: Regular backups and restoration processes are crucial for data recovery (Miller, 2018).

Monitoring and Alerting: Vigilant monitoring systems detect issues for swift corrective actions (Johnson, 2021).

Replication: Techniques like streaming or logical replication ensure standby nodes are ready in case of a primary node failure, minimizing downtime (Smith, 2020).

Data Center Redundancy: Deploying servers across geographically dispersed data centers offers added HA by seamlessly shifting services in case of data center failures (Adams, 2019).

This comprehensive overview of the theoretical framework provides a deeper understanding of PostgreSQL, clustering techniques, replication methods, and high availability strategies, empowering organizations to ensure data integrity, scalability, and uninterrupted database access (Brown, 2020).

4.5. Citus Data: Scaling PostgreSQL Horizontally

Citus Data is a prominent player in the world of PostgreSQL database management, specializing in horizontal scaling. Citus extends the capabilities of PostgreSQL, allowing it to efficiently handle large volumes of data and demanding workloads by distributing data across multiple nodes.

Architecture: Citus deploys a distributed architecture, enabling PostgreSQL to scale horizontally. It partitions tables into smaller "shards" distributed across multiple servers, providing parallel query execution and high performance (Citus Data, 2021).

Scalability: Citus offers an automatic sharding feature, facilitating the addition of new nodes as data volumes grow. This approach helps PostgreSQL databases easily adapt to the evolving needs of organizations (Smith, 2022).

Complex Queries: Citus ensures that complex queries can be executed in parallel across shards, improving the overall query performance and response time. It provides a seamless experience for users dealing with intricate analytical and transactional workloads (Jones, 2020).

Data Distribution: Citus offers various data distribution strategies, allowing users to choose between distributing data based on specific columns, such as date or region, to optimize query performance (Brown, 2021).

Citus Data's contributions to PostgreSQL horizontal scaling make it a valuable solution for organizations that require both the robustness of PostgreSQL and the ability to scale their database horizontally to meet the demands of large-scale applications.

5. Methodology

Installing Postgresql in Linux

In this chapter, we will cover the steps to install PostgreSQL and access it for the first time.

Installation

1. Use the package manager to install the desired PostgreSQL version:

```
sudo sh -c 'echo "deb http://apt.postgresql.org/pub/repos/apt $(lsb_release
-cs)-pgdg main" > /etc/apt/sources.list.d/pgdg.list'
wget --quiet -O - https://www.postgresql.org/media/keys/ACCC4CF8.asc | sudo
apt-key add -
sudo apt-get update
sudo apt-get -y install postgresql
```

This installation process follows PostgreSQL's official repository and package management guidelines (PostgreSQL Documentation, 2021).

Starting the Server

2. Verify if the PostgreSQL server has started after installation:

```
sudo systemctl start postgresql
sudo systemctl enable postgresql
```

Starting the PostgreSQL server and enabling it as a system service ensures that it will run upon system startup (PostgreSQL Wiki, 2022).

Accessing PostgreSQL

3. Access PostgreSQL for the first time using the following commands:

```
sudo -i -u postgres
```

```
psql
```

After executing these steps, you will be logged into the PostgreSQL database with the **'postgres'** user. You are now ready to interact with the PostgreSQL database system.

These revised instructions provide a clear sequence of steps for installing PostgreSQL, starting the server, and accessing it for the first time. Please ensure that you have the necessary permissions and dependencies for these actions to be successful on your system.

6: Clustering in PostgreSQL

In this chapter, we delve into the concept of clustering in PostgreSQL, a technique that can significantly enhance query performance by optimizing the physical storage of data (Schönig, 2020). Clustering aims to align the physical order of rows in a table with the logical order defined by an index, thereby minimizing disk I/O and improving query response times.

6.1 Initializing the PostgreSQL Database

Before creating a table and performing clustering, we need to initialize a PostgreSQL database. This is done using the `initdb` command, which prepares the directory structure and configuration files required for a new database cluster (Schönig, 2020).

```
initdb -D /path/to/data/directory
```

Replace `/path/to/data/directory` with the actual path where you want to store the database files. This step is essential for setting up the PostgreSQL environment.

6.2 Creating a Table

Once the PostgreSQL database is initialized, we can proceed to create a table that will serve as the foundation for our data clustering experiment. In PostgreSQL, this is accomplished using the `CREATE TABLE` command. Here's an example of creating a table named `table1`:

```
CREATE TABLE table1 (  
  
    id serial PRIMARY KEY,  
  
    name varchar(255),  
  
    age int,  
  
    employment boolean  
  
);
```

This command creates a table with four columns: ``id``, ``name``, ``age``, and ``employment``, along with their respective data types.

6.3 Adding Data

Once the table is created, we can proceed to populate it with data using the `INSERT` command. Here's an example of inserting data into `table1`:

```
INSERT INTO table1 (name, age, employment)
```

```
VALUES
```

```
(`John`, 20, false),
```

```
(`Vagelis`, 30, true),
```

```
(`George`, 40, true);
```

This command inserts three rows into the table, each with values for the ``name``, ``age``, and ``employment`` columns.

6.4 Creating an Index

Before we can cluster the table, we need to create an index on the column by which we want to cluster it. The index defines the logical order of the data and will be used as a reference during the clustering process. In this example, we create an index named `table1_index` on the ``name`` column:

```
CREATE INDEX table1_index ON table1(name);
```

6.5 Performing the Clustering

With the index in place, we can proceed to perform the clustering using the `CLUSTER` command. This command rearranges the physical order of rows in the table to match the order specified by the index. Here's how to execute the clustering process:

```
CLUSTER table1 USING table1_index;
```

It's important to note that the clustering process can take a noticeable amount of time, especially for large tables. Additionally, if the table experiences frequent insertions, deletions, or updates, it may impact performance. Therefore, it's advisable to periodically rerun the `CLUSTER` command to maintain optimal data organization.

6.6 Considerations and Best Practices

It's essential to use clustering judiciously, primarily on high-traffic tables where the performance gains outweigh the cost of the clustering process. Since clustering utilizes memory resources, it's recommended to apply it thoughtfully, considering the specific requirements of your PostgreSQL database (Schönig, 2020).

In the following chapters, we'll explore more advanced clustering strategies and optimizations to further enhance the performance of PostgreSQL databases, using additional insights from the PostgreSQL 13 Cookbook (Naga, 2021) and the PostgreSQL Documentation (PostgreSQL Documentation, 2021).

7.Replication in PostgreSQL

7.1.Setting Up a Replication Server

In our replication experiment, we aim to configure asynchronous, synchronous, and logical replication. To achieve this, we need to create a second PostgreSQL server to act as the replication target. Here are the steps to set up the replication server:

7.1.1. Create a Data Directory

Start by creating a dedicated data directory for the replication server. This directory should be owned by the PostgreSQL user, which will be used to replicate data into it.

```
sudo mkdir /var/lib/postgresql/replica  
sudo chown postgres:postgres /var/lib/postgresql/replica
```

7.1.2. Initialize the Second Cluster

Next, initialize the second PostgreSQL cluster using the `initdb` command. This command prepares the necessary directory structure and configuration files for the new cluster.

```
sudo -u postgres /usr/lib/postgresql/Postgresql14/bin/initdb -D  
/var/lib/postgresql/replica
```

7.1.3. Make Configuration Changes

To ensure a smooth replication process, make configuration changes to the new instance while the cluster is offline. First, stop the new cluster using `pg_ctl`:

```
pg_ctl stop -D /var/lib/postgresql/14/replica
```

7.1.4. Copy the Configuration File

Each PostgreSQL instance has its own configuration file. To replicate the configuration from the main cluster, copy the main configuration file to the replica's data folder.

```
sudo cp /etc/postgresql/14/main/postgresql.conf  
/etc/postgresql/Postgresql14/second-server/postgresql.conf
```

Now, open the replica's configuration file for editing:

```
sudo nano /etc/postgresql/Postgresql14/second-server/postgresql.conf
```

7.1.5. Change the Port

In the replica's configuration file (**postgresql.conf**), locate the **Port** option and change its value to a port number that is not already in use. Note that the default port for the main cluster is usually 5432. Changing the port prevents conflicts when running both instances simultaneously.

7.1.6. Start the Replica Server

Finally, restart the replication server service using the following command:

```
pg_ctl start -D /var/lib/postgresql/14/replica
```

With these steps completed, we have successfully set up a second PostgreSQL server to serve as the replication target. We can now proceed with configuring and testing asynchronous, synchronous, and logical replication as per your experiment requirements.

Important: not to use postgres or admin user to replicate for security reasons and separate them.

7.2.Asynchronous

Setting up Asynchronous Replication in PostgreSQL

To establish asynchronous replication in PostgreSQL, follow these steps:

1. Stop the PostgreSQL service:

```
pg_ctl stop -D /postgres/master
```

2. Edit the `postgresql.conf` file on the primary server using the `nano` command:

```
nano /postgresmaster/postgresql.conf
```

Inside the file, navigate using `Ctrl+W` and update the following configurations:

```
# Enable Write-Ahead Logging (WAL) archiving
```

```
wal_level = replica
```

```
hot_standby = on
```

```
archive_mode = on
```

```
archive_command = 'command to archive WAL segments'
```

```
listen_addresses = '*'
```

```
# Configure replication settings (optional but recommended for performance)
```

```
max_wal_senders = max_number_of_standbys
```

```
wal_keep_segments = number_of_WAL_segments_to_keep
```

Note: The last two settings are optional but can enhance performance.

3. Edit the `pg_hba.conf` file:

```
nano postgres/master/pg_hba.conf
```

Add this line to the IPv4 local connections section:

```
local replication all replication_user localhost md5
```

If working locally and trusting the connection, you can use `'trust'` instead of `'md5'`.

4. Restart the PostgreSQL service:

```
pg_ctl start -D /postgres/master
```

5. Connect to `'psql'` and create the `'replication_user'` with replication privileges:

```
CREATE USER replication_user REPLICATION LOGIN ENCRYPTED  
PASSWORD 'strong';
```

a. Create a new table named `'employees'`:

```
CREATE TABLE employees (id int PRIMARY KEY, name varchar);
```

b. Populate the `'employees'` table using the `'generate_series'` function:

```
INSERT INTO employees  
SELECT generate_series(1, 100), 'somename' || generate_series(101, 200);
```

6. Full Backup of the Master:

a. Create a folder for the slave server or delete existing data:

```
mkdir /postgres/slave
```

b. Run `pg_basebackup`:

```
pg_basebackup -h <master_ip/localhost> -U replication_user -p <port> -D  
<slave_folder> -Fp -Xs -R
```

Example:

```
pg_basebackup -h localhost -U replication_user -p 1111 -D /postgres/slave -Fp  
-Xs -R
```

- `-h` = *host*

- `-U` = *user*

- `-p` = *port*

- `-D` = *directory*

- `-Fp` = *format plain (use `-Ft` for tar format, requiring unzipping)*

- `-Xs` = *WAL method stream*

- `-R` = *creates the configuration for it to be a replica*

7. Change the port of the slave:

Edit `postgresql.conf`:

```
nano /postgres/slave/postgresql.conf
```

Set the port to your desired value, e.g., `port = 2111`.

8. Save the file and start the cluster:

```
pg_ctl start -D /postgres/slave
```

9. Test Replication:

To verify successful replication, insert or delete a row from the `employees` table on the master:

```
DELETE FROM employees WHERE id = 2;
```

Then, on the slave, run:

```
SELECT * FROM employees WHERE id = 2;
```

The latter query should yield no rows.

```
# WRITE-AHEAD LOG
#-----
# - Settings -
wal_level = replica          # minimal, replica, or logical
                             # (change requires restart)
#fsync = on                  # flush data to disk for crash safety
```

Set wal_level to replica.

```
#primary_slot_name =         # replication slot on sending server
                             # (change requires restart)
#promote_trigger_file = ''  # file name whose presence ends recovery
hot_standby = on            # "off" disallows queries during recovery
                             # (change requires restart)
```

Set hot_standby to on.

```
listen_addresses = '*'      # what IP address(es) to listen on;
                             # comma-separated list of addresses;
                             # defaults to 'localhost'; use '*' for all
```

Set listen_address to * for every ip or x.x.x.x for specific ip.

```
# - Archiving -
archive_mode = on           # enables archiving; off, on, or always
                             # (change requires restart)
archive_command = 'test ! -f /postgress/temp/%f && cp %p /postgress/temp/%f'
                             # placeholders: %p = path of file to archive
                             #             %f = file name only
                             # e.g. 'test ! -f /mnt/server/archivedir/%f && cp %p /mnt/server/
archive_timeout = 30        # force a logfile segment switch after this
```

Set archive mode to on and set the command to create the archive.

```
GNU nano 4.8 postgresql.conf
# If external_pid_file is not explicitly set, no extra PID file is written.
#external_pid_file = '          # write an extra PID file
                                # (change requires restart)

#-----
# CONNECTIONS AND AUTHENTICATION
#-----

# - Connection Settings -

#listen_addresses = 'localhost'          # what IP address(es) to listen on;
                                          # comma-separated list of addresses;
                                          # defaults to 'localhost'; use '*' for all
                                          # (change requires restart)
port = 1111                              # (change requires restart)
max_connections = 100                    # (change requires restart)
#superuser_reserved_connections = 3      # (change requires restart)
#unix_socket_directories = '/var/run/postgresql' # comma-separated list of directories
                                          # (change requires restart)
#unix_socket_group = ''                  # (change requires restart)
#unix_socket_permissions = 0777         # begin with 0 to use octal notation
                                          # (change requires restart)
#bonjour = off                           # advertise server via Bonjour
                                          # (change requires restart)
#bonjour_name = ''                       # defaults to the computer name
                                          # (change requires restart)

# - TCP settings -
# see "man 7 tcp" for details
```

Set the port to an not allocated port.

```
# Allow replication connections from localhost, by a user with the
# replication privilege.
local  replication  all                                     trust
host   replication  all 127.0.0.1/32                          trust
host   replication  all ::1/128                               trust
host   replication  rep_user localhost                          md5
```

[Wrote 98 lines]

Append the last line to pg_hba.conf.

```
postgres@vagelis-VirtualBox:/postgres/master$ nano postgresql.conf
postgres@vagelis-VirtualBox:/postgres/master$ /usr/lib/postgresql/12/bin/pg_ctl restart -D /postgres/m
aster
```

Restart the server.

```
postgres=# \conninfo
You are connected to database "postgres" as user "postgres" via socket in "/var/run/postgresql" at port
"1111".
postgres=# CREATE USER rep_user REPLICATION LOGIN ENCRYPTED PASSWORD 'strong';
```

Create the user for replication.

```
postgres=# create table employees(id int PRIMARY KEY,name varchar);
CREATE TABLE
postgres=# INSERT INTO employees select generate_series(1,100),'somename' || generate_series(101,200);
INSERT 0 100
postgres=# select * from employees;
postgres=# select * from employees limit 10;
 id | name
-----+-----
  1 | somename101
  2 | somename102
  3 | somename103
  4 | somename104
  5 | somename105
  6 | somename106
  7 | somename107
  8 | somename108
  9 | somename109
 10 | somename110
(10 rows)
```

Create and populate a table.

```
postgres@vagelis-VirtualBox:/postgres/slave$ pg_basebackup -h localhost -U rep_user -p 1111 -D /postgres/slave -Fp -Xs -R
```

Run the pg_basebackup command.

```
port = 2111
max_connections = 100
# defaults to 'localhost'; use '*' for all
# (change requires restart)
# (change requires restart)
# (change requires restart)
```

Change port on slave.

```
postgres@vagelis-VirtualBox:/postgres/master$ /usr/lib/postgresql/12/bin/pg_ctl start -D /postgres/slave
waiting for server to start...2023-09-17 21:15:00.532 EEST [11563] LOG: starting PostgreSQL 12.16 (Ubuntu 12.16-0ubuntu0.20.04.1) on x86_64-pc-linux-gnu, compiled by gcc (Ubuntu 9.4.0-1ubuntu1~20.04.1) 9.4.0, 64-bit
2023-09-17 21:15:00.532 EEST [11563] LOG: listening on IPv4 address "0.0.0.0", port 2111
2023-09-17 21:15:00.532 EEST [11563] LOG: listening on IPv6 address ":::", port 2111
2023-09-17 21:15:00.548 EEST [11563] LOG: listening on Unix socket "/var/run/postgresql/.s.PGSQL.2111"
2023-09-17 21:15:00.594 EEST [11564] LOG: database system was interrupted; last known up at 2023-09-17 21:13:08 EEST
2023-09-17 21:15:01.543 EEST [11564] LOG: entering standby mode
2023-09-17 21:15:01.553 EEST [11564] LOG: redo starts at 0/C000028
2023-09-17 21:15:01.558 EEST [11564] LOG: consistent recovery state reached at 0/C000138
2023-09-17 21:15:01.559 EEST [11563] LOG: database system is ready to accept read only connections
2023-09-17 21:15:01.570 EEST [11568] LOG: started streaming WAL from primary at 0/D000000 on timeline 1
done
server started
```

Start the slave server.


```
postgres@vagelis-VirtualBox:/postgres/master$ psql -p 1111
psql (12.16 (Ubuntu 12.16-0ubuntu0.20.04.1))
Type "help" for help.
```

```
postgres=# delete from employees where id = 2;
DELETE 1
postgres=#
```

```
postgres=# select * from employees limit 10;
 id |      name
----+-----
  1 | somename101
  2 | somename102
  3 | somename103
  4 | somename104
  5 | somename105
  6 | somename106
  7 | somename107
  8 | somename108
  9 | somename109
 10 | somename110
(10 rows)
```

```
postgres=# select * from employees limit 10;
 id |      name
----+-----
  1 | somename101
  3 | somename103
  4 | somename104
  5 | somename105
  6 | somename106
  7 | somename107
  8 | somename108
  9 | somename109
 10 | somename110
 11 | somename111
(10 rows)
```

This the transaction that happens to master and the aforementioned result on the slave

7.3.Synchronous

In the pursuit of enhancing the reliability and data consistency of a PostgreSQL database, the transition from asynchronous to synchronous replication is a pivotal step. In this section, we delve into the intricacies of configuring synchronous replication, highlighting the critical steps involved in this transformation.

Configuration Steps

To transition into synchronous replication, one must first access the PostgreSQL command-line client, commonly referred to as the 'psql' client, on the master server. Alternatively, configuration can also be achieved by modifying the 'postgresql.conf' file.

Option 1: Using the psql Client

Execute the following command within the psql client on the master server:

```
alter system set synchronous_standby_names to '*';
```

Option 2: Modifying 'postgresql.conf'

Alternatively, open the 'postgresql.conf' file and modify the parameter 'synchronous_standby_names' as follows:

```
synchronous_standby_names = '*'
```

Once the chosen setting is applied, it is imperative to restart the PostgreSQL service to enact the changes effectively. This can be accomplished using the following command:

```
pg_ctl restart -D /postgress/master
```

```
postgres=# alter system set synchronous_standby_names to '*';
ALTER SYSTEM
postgres=# █
```

```
postgres@vagelis-VirtualBox:/postgress/slave1$ /usr/lib/postgresql/12/bin/pg_ctl restart -D /postgress/master █
```

7.4. The Distinctiveness of Synchronous Replication

One of the most striking distinctions between synchronous and asynchronous replication is how they handle replication in the event of standby server unavailability. In synchronous replication, when the standby server experiences downtime or interruption, the master server halts any further transaction processing. This pause in operations persists until the slave or standby server resurfaces and can resume replication.

This behavior manifests explicitly after the standby server has been initiated or 'started.' It emphasizes the synchronous nature of the replication process, where the master insists on real-time verification from the slave for each transaction before it proceeds.

Conversely, in asynchronous replication, the master does not impose such stringent constraints. Even if the standby server encounters a disruption, the master continues to function and accumulate Write-Ahead Log (WAL) segments. Asynchronous replication, in essence, follows a fire-and-forget paradigm, allowing the master to forge ahead while the standby may lag behind.

To mitigate potential data loss inherent to asynchronous replication, several best practices must be observed. These include the implementation of comprehensive monitoring systems to detect and address issues with standbys, ensuring sufficient WAL retention on the master, and having well-defined backup and recovery procedures in place to synchronize a standby with the master when necessary.

By grasping the distinctions between synchronous and asynchronous replication, database administrators can make informed decisions regarding which replication mode best suits their data integrity and availability requirements. The choice between these modes reflects a delicate balance between immediate consistency and potentially increased latency, illustrating the nuanced nature of database replication in PostgreSQL.

This chapter provides a comprehensive overview of the steps involved in transitioning to synchronous replication and elucidates the fundamental disparities between synchronous and asynchronous replication modes. Understanding these intricacies is essential for designing robust and resilient database systems that align with specific business needs and data integrity objectives.

```
postgres@vagelis-VirtualBox:~/postgres$ /usr/lib/postgresql/12/bin/pg_ctl stop -D /postgres/slave
waiting for server to shut down...done
server stopped
postgres@vagelis-VirtualBox:~$

postgres@vagelis-VirtualBox:~/postgres/slave1$ psql -p 1111
psql (12.16 (Ubuntu 12.16-0ubuntu0.20.04.1))
Type "help" for help.

postgres=# select * from employees limit 10;
 id | name
----+-----
  1 | somename101
  3 | somename103
  4 | somename104
  5 | somename105
  6 | somename106
  7 | somename107
  8 | somename108
  9 | somename109
 10 | somename110
 11 | somename111
(10 rows)

postgres=# insert into employees values(112,'test');
```

In this image we can see what happens when the slave is unavailable and the master is trying to insert data.

```
postgres=# insert into employees values(112,'test');
2023-09-18 02:40:30.706 EEST [16573] LOG: standby "walreceiver" is now a synchronous standby with priority 1
2023-09-18 02:40:30.706 EEST [16573] STATEMENT: START_REPLICATION 0/1200000 0 TIMELINE 1
INSERT 0 1
postgres=#
```

This will occur after the slave has been started.

```
postgres=# insert into employees values(114,'test');
INSERT 0 1
postgres=#
```

And this is what happens when the slave is unavailable and the master is using asynchronous replication.

7.5.Logical

In this chapter, we will outline the steps required to configure logical replication between a publisher and a subscription server.

Publisher Configuration

1. Stop the PostgreSQL service on the publisher server:

```
pg_ctl stop -D /var/lib/postgresql/14/main
```

2. Open the `postgresql.conf` file of the main server using the `nano` command:

```
nano /var/lib/postgresql/14/main/postgresql.conf
```

3. Change the following configuration value to enable logical replication:

```
wal_level = logical
```

4. Save the changes and exit the editor.
5. Restart the PostgreSQL service on the publisher server:

```
pg_ctl start -D /var/lib/postgresql/14/main
```

Subscription Server Configuration

1. Stop the PostgreSQL service on the subscription server:

```
pg_ctl stop -D /var/lib/postgresql/14/replica
```

2. Open the `postgresql.conf` file of the subscription server using the `nano` command:

```
nano /var/lib/postgresql/14/main/postgresql.conf
```

3. Choose an available port (e.g., 5434) for the subscription server:

```
port = 5434
```

4. Save the changes and exit the editor.

5. Restart the PostgreSQL service on the subscription server:

```
pg_ctl start -D /var/lib/postgresql/14/replica
```

Test Logical Replication

To verify that logical replication is working correctly, perform the following steps:

1. Connect to the publisher server using `psql` with the specified port:

```
psql -p 5433 postgres
```

2. Create a new database named `testdb`:

```
CREATE DATABASE testdb;
```

```
\c testdb;
```

3. Create a new table named `employees`:

```
CREATE TABLE employees (id int PRIMARY KEY, name varchar);
```

4. Populate the `employees` table using the `generate_series` function:

```
INSERT INTO employees  
SELECT generate_series(1, 100), 'name' || generate_series(101, 200);
```

5. Check if data has been successfully added to the `employees` table:

```
SELECT * FROM employees;
```

Data Transfer to Subscription Server

To transfer data from the publisher server to the subscription server, use the following `pg_dump` and `psql` commands:

```
pg_dump -t employees -s testdb -p 5432 | psql -p 5434 test_sub_db
```

If `test_sub_db` does not exist on the subscription server, create it before running the `pg_dump` command:

```
CREATE DATABASE test_sub_db;
```

You can then verify the data transfer by connecting to the subscription server:

```
psql -p 5434 postgres  
\c test_sub_db; -- Switch to the test_sub_db database  
\d; -- List tables  
SELECT * FROM employees; -- Query the data
```

Finalizing the Connection

On the publication server, run the following command to publish all tables of the `testdb` database:

```
CREATE PUBLICATION publication_1 FOR ALL TABLES;
```

If you only want to publish the `employees` table, use this command instead:

```
CREATE PUBLICATION publication_1 FOR TABLE employees;
```

On the subscription server, create a subscription for the publication:

```
CREATE SUBSCRIPTION sub_1 CONNECTION 'dbname=testdb  
host=localhost user=postgres port=5432' PUBLICATION publication_1;
```

This completes the setup of logical replication between the publisher and subscription servers.

Please ensure that you replace placeholders like database names, port numbers, and server paths with the actual values applicable to your setup. Additionally, always exercise caution when making configuration changes to your PostgreSQL servers.

```
vagelis@vagelis-VirtualBox:~$ sudo mkdir /postg  
vagelis@vagelis-VirtualBox:~$ sudo mkdir /postg/pub  
vagelis@vagelis-VirtualBox:~$ sudo mkdir /postg/sub  
vagelis@vagelis-VirtualBox:~$ sudo chon postgres:postgres /postg/pub  
sudo: chon: command not found  
vagelis@vagelis-VirtualBox:~$ sudo chown postgres:postgres /postg/pub  
vagelis@vagelis-VirtualBox:~$ sudo chown postgres:postgres /postg/sub
```

Create the folders that will have the clusters and give ownership to postgres user.


```
vagelis@vagelis-VirtualBox:~$ locate initdb
/etc/alternatives/initdb.1.gz
/usr/lib/postgresql/12/bin/initdb
/usr/lib/postgresql/15/bin/initdb
/usr/share/locale/de/LC_MESSAGES/initdb-15.mo
/usr/share/locale/el/LC_MESSAGES/initdb-15.mo
/usr/share/locale/es/LC_MESSAGES/initdb-15.mo
/usr/share/locale/fr/LC_MESSAGES/initdb-15.mo
/usr/share/locale/it/LC_MESSAGES/initdb-15.mo
/usr/share/locale/ja/LC_MESSAGES/initdb-15.mo
/usr/share/locale/ka/LC_MESSAGES/initdb-15.mo
/usr/share/locale/ko/LC_MESSAGES/initdb-15.mo
/usr/share/locale/pt_BR/LC_MESSAGES/initdb-15.mo
/usr/share/locale/ru/LC_MESSAGES/initdb-15.mo
/usr/share/locale/sv/LC_MESSAGES/initdb-15.mo
/usr/share/locale/uk/LC_MESSAGES/initdb-15.mo
/usr/share/locale/zh_CN/LC_MESSAGES/initdb-15.mo
/usr/share/man/man1/initdb.1.gz
/usr/share/postgresql/12/man/man1/initdb.1.gz
/usr/share/postgresql/15/man/man1/initdb.1.gz
```

Search for the initdb command.

```
vagelis@vagelis-VirtualBox:~$ sudo -u postgres /usr/lib/postgresql/15/bin/initdb
-D /postg/pub
The files belonging to this database system will be owned by user "postgres".
This user must also own the server process.

The database cluster will be initialized with this locale configuration:
provider:      libc
LC_COLLATE:   en_US.UTF-8
LC_CTYPE:     en_US.UTF-8
LC_MESSAGES:  en_US.UTF-8
LC_MONETARY:  eL_GR.UTF-8
LC_NUMERIC:   eL_GR.UTF-8
LC_TIME:      eL_GR.UTF-8
The default database encoding has accordingly been set to "UTF8".
The default text search configuration will be set to "english".

Data page checksums are disabled.

fixing permissions on existing directory /postg/pub ... ok
creating subdirectories ... ok
selecting dynamic shared memory implementation ... posix
selecting default max_connections ... 100
selecting default shared_buffers ... 128MB
selecting default time zone ... Europe/Athens
creating configuration files ... ok
running bootstrap script ... ok
performing post-bootstrap initialization ... ok
syncing data to disk ... ok
```

Creation of pub cluster.

```
initdb: warning: enabling "trust" authentication for local connections
initdb: hint: You can change this by editing pg_hba.conf or using the option -A
, or --auth-local and --auth-host, the next time you run initdb.
```

Success. You can now start the database server using:

```
/usr/lib/postgresql/15/bin/pg_ctl -D /postg/pub -l logfile start
```

```
vagelis@vagelis-VirtualBox:~$ sudo -u postgres /usr/lib/postgresql/15/bin/initdb
-D /postg/sub
```

The files belonging to this database system will be owned by user "postgres".
This user must also own the server process.

The database cluster will be initialized with this locale configuration:

```
provider:      libc
LC_COLLATE:    en_US.UTF-8
LC_CTYPE:      en_US.UTF-8
LC_MESSAGES:   en_US.UTF-8
LC_MONETARY:   eL_GR.UTF-8
LC_NUMERIC:    eL_GR.UTF-8
LC_TIME:       eL_GR.UTF-8
```

The default database encoding has accordingly been set to "UTF8".
The default text search configuration will be set to "english".

Data page checksums are disabled.

Creation of sub cluster.

```
vagelis@vagelis-VirtualBox:~$ locate pg_ctl
/etc/alternatives/pg_ctl.1.gz
/etc/postgresql/12/main/pg_ctl.conf
/etc/postgresql/15/main/pg_ctl.conf
/usr/bin/pg_ctlcluster
/usr/lib/postgresql/12/bin/pg_ctl
/usr/lib/postgresql/15/bin/pg_ctl
/usr/share/locale/cs/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/de/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/el/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/es/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/fr/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/it/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/ja/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/ka/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/ko/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/pt_BR/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/ru/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/sv/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/tr/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/uk/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/zh_CN/LC_MESSAGES/pg_ctl-15.mo
/usr/share/man/man1/pg_ctl.1.gz
/usr/share/man/man1/pg_ctlcluster.1.gz
/usr/share/postgresql/12/man/man1/pg_ctl.1.gz
/usr/share/postgresql/15/man/man1/pg_ctl.1.gz
```

Locating the pg_ctl command in order to start the cluster.

```
vagelis@vagelis-VirtualBox:~$ sudo su postgres
```

Change to postgres user.

```
postgres@vagelis-VirtualBox:/home/vagelis$ cd /postg/pub
postgres@vagelis-VirtualBox:/postg/pub$ ls
base          pg_ident.conf  pg_serial      pg_tblspc      postgresql.auto.conf
global        pg_logical     pg_snapshots   pg_twophase     postgresql.conf
pg_commit_ts  pg_multixact   pg_stat         PG_VERSION
pg_dynshmem   pg_notify      pg_stat_tmp     pg_wal
pg_hba.conf   pg_replslot    pg_subtrans     pg_xact
postgres@vagelis-VirtualBox:/postg/pub$ nano postgresql.conf
postgres@vagelis-VirtualBox:/postg/pub$ cd ..
postgres@vagelis-VirtualBox:/postg$ cd /sub
bash: cd: /sub: No such file or directory
postgres@vagelis-VirtualBox:/postg$ cd sub
postgres@vagelis-VirtualBox:/postg/sub$ nano postgresql.conf
postgres@vagelis-VirtualBox:/postg/sub$ locate pg_ctl
/etc/alternatives/pg_ctl.1.gz
/etc/postgresql/12/main/pg_ctl.conf
/etc/postgresql/15/main/pg_ctl.conf
/usr/bin/pg_ctlcluster
/usr/lib/postgresql/12/bin/pg_ctl
/usr/lib/postgresql/15/bin/pg_ctl
/usr/share/locale/cs/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/de/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/el/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/es/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/fr/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/it/LC_MESSAGES/pg_ctl-15.mo
/usr/share/locale/ja/LC_MESSAGES/pg_ctl-15.mo
```

Locating the postg folder to change the postgresql.conf files.

```
# defaults to localhost, use
# (change requires restart)
port = 5555 # (change requires restart)
max_connections = 100 # (change requires restart)
#superuser_reserved_connections = 3 # (change requires restart)

# - Settings -
wal_level = logical # minimal, replica, or logical
# (change requires restart)
```

Port 5555 was used for Publisher cluster wal_level was set to logical.

```
# (change requires restart)
port = 6666 # (change requires restart)
max_connections = 100 # (change requires restart)
#superuser_reserved_connections = 3 # (change requires restart)
#unix_socket_directory = '/var/run/postgresql' # comma-separated list of
```

Port 6666 was used for Subscription cluster.

```

postgres@vagelis-VirtualBox:/postg/sub$ /usr/lib/postgresql/15/bin/pg_ctl -D /p
ostg/pub start
waiting for server to start....2023-08-31 19:39:04.690 EEST [18476] LOG: start
ing PostgreSQL 15.4 (Ubuntu 15.4-1.pgdg20.04+1) on x86_64-pc-linux-gnu, compile
d by gcc (Ubuntu 9.4.0-1ubuntu1~20.04.1) 9.4.0, 64-bit
2023-08-31 19:39:04.705 EEST [18476] LOG: listening on IPv4 address "127.0.0.1
", port 5555
2023-08-31 19:39:04.741 EEST [18476] LOG: listening on Unix socket "/var/run/p
ostgresql/.s.PGSQL.5555"
2023-08-31 19:39:04.761 EEST [18479] LOG: database system was shut down at 202
3-08-31 19:33:49 EEST
2023-08-31 19:39:04.798 EEST [18476] LOG: database system is ready to accept c
onnections
done

```

Starting Publisher cluster

```

postgres@vagelis-VirtualBox:/postg/sub$ /usr/lib/postgresql/15/bin/pg_ctl -D /p
ostg/sub start
waiting for server to start....2023-08-31 19:39:09.625 EEST [18486] LOG: start
ing PostgreSQL 15.4 (Ubuntu 15.4-1.pgdg20.04+1) on x86_64-pc-linux-gnu, compile
d by gcc (Ubuntu 9.4.0-1ubuntu1~20.04.1) 9.4.0, 64-bit
2023-08-31 19:39:09.626 EEST [18486] LOG: listening on IPv4 address "127.0.0.1
", port 6666
2023-08-31 19:39:09.661 EEST [18486] LOG: listening on Unix socket "/var/run/p
ostgresql/.s.PGSQL.6666"
2023-08-31 19:39:09.681 EEST [18489] LOG: database system was shut down at 202
3-08-31 19:33:58 EEST
2023-08-31 19:39:09.711 EEST [18486] LOG: database system is ready to accept c
onnections
done
server started

```

Starting Subscription cluster

```

postgres@vagelis-VirtualBox:/postg/sub$ psql -p 5555
psql (15.4 (Ubuntu 15.4-1.pgdg20.04+1))
Type "help" for help.

postgres=# create database testdb;
CREATE DATABASE

```

Login to psql and create the new testdb for Publisher cluster.

```

testdb=# create publication pub1 for table employees;
CREATE PUBLICATION

```

Creation of the publication for table employees.

```

test_sub_db=# create subscription sub1 connection 'dbname=testdb host=localhost
user=postgres port=5555' publication pub1;
NOTICE: created replication slot "sub1" on publisher
CREATE SUBSCRIPTION
test_sub_db=# █

```

Creation of subscription on table publication pub1.

8.High Availability in PostgreSQL: Ensuring Data Resilience and Continuity

In today's data-driven world, databases form the backbone of countless applications and services, making high availability (HA) a paramount concern for organizations that rely on PostgreSQL as their relational database management system (RDBMS). PostgreSQL is celebrated for its robust features and extensibility, but even the most well-architected systems can encounter failures. This is where the concept of high availability comes into play: ensuring that your PostgreSQL database remains accessible, responsive, and reliable, even in the face of hardware failures, software crashes, or other unforeseen issues.

8.1.The Importance of High Availability

High availability, in the context of databases, refers to the ability of a system to provide uninterrupted access to data and services despite various potential disruptions. It is a crucial component of business continuity and disaster recovery planning, as downtimes can lead to substantial financial losses, damage to reputation, and, in some industries, even safety risks.

8.2.Downtime Costs

Downtime is costly. A brief outage can inconvenience users, disrupt transactions, and tarnish a company's reputation. In more severe cases, it can lead to data loss, financial penalties, and legal repercussions. Therefore, ensuring that your PostgreSQL database remains operational 24/7 is a strategic imperative.

8.3.Types of Failures

Various factors can lead to database outages, including hardware failures (e.g., disk crashes, server hardware issues), software issues (e.g., database corruption, software bugs), network problems (e.g., loss of connectivity), and even human error (e.g., accidental data deletion). An effective high availability strategy must account for these potential failures.

8.4.Manual Promotion in PostgreSQL

One of the methods for achieving high availability in PostgreSQL is through manual promotion of a standby server to a primary server. This approach is based on PostgreSQL's built-in support for streaming replication and cascading replication slots.

8.4.1.Example 1: Manual Promotion

In this practical example, we illustrate the process of manual promotion and the subsequent resynchronization of a demoted master in a PostgreSQL database cluster. This operation is critical for maintaining high availability and data integrity within the system.

Manual Promotion:

1. Initiating the process, we gracefully shut down the current master using the following commands:

```
pg_ctl -D $PGDATA stop
```

2. To designate a new master, we promote the standby server, ensuring a smooth transition of roles:

```
pg_ctl -D /postgress/slave1 promote
```

3. In cases involving multiple standby servers, it is imperative to update the `postgresql.conf` file on each standby server with the new primary connection information, enhancing fault tolerance:

```
primary_conninfo = 'user=rep_user password=secret host=localhost  
port=4444'
```

4. Subsequently, restart the second standby server to apply the configuration changes, thus completing the manual promotion process.

Switchover Validation:

Before initiating a manual switchover, it is crucial to verify that the old master is indeed unreachable and unavailable for read and write operations. Failure to do so may result in split-brain scenarios. To mitigate this risk, we employ the following command to shut down the PostgreSQL instance on the old master that requires demotion:

```
pg_ctl -D /postgress/master stop
```

Upon successful execution of step 2, we proceed to the switchover process.

Promote Standby Server:

1. With confidence in the unavailability of the old master, we promote the selected standby server to assume the role of the new master:

```
pg_ctl -D $PGDATA promote
```

2. Subsequently, reconfigure the remaining servers to recognize the new master as the leader by updating their respective `postgresql.conf` files, as exemplified in step 3.

Resynchronization of Demoted Master:

To efficiently resynchronize the demoted master as a standby server, eliminating the need for a complete recreation, the following steps must be followed:

- Ensure that the `wal_log_hints` or `data_checksums` feature is enabled. If not, enable it on both the new master and all its standby servers. Achieve this by connecting to the PostgreSQL server via `psql` and executing the command:

```
ALTER SYSTEM SET wal_log_hints TO 'ON';
```

- Note: The availability of WAL segments since the failover event is crucial for the successful execution of the subsequent steps.

1. Commence the resynchronization process by shutting down the old master gracefully:

```
pg_ctl -D $PGDATA stop
```

2. Utilize the `pg_rewind` utility to resynchronize the old master with the new master. The command should resemble the following:

```
pg_rewind -D /postgres/master --source-server="host=localhost port=4444 user=postgres"
```

3. Apply the updated primary connection information settings in the `postgresql.conf` file of the old master, as demonstrated in step 3.

4. Create a `standby.signal` file within the data directory of the old master to indicate its role as a standby server:

```
touch $PGDATA/standby.signal
```

5. Conclude the process by starting the old master using the command:

```
pg_ctl -D $PGDATA start
```

This meticulous procedure ensures the seamless resynchronization of the demoted master as a standby server, minimizing disruption and maintaining data consistency within the PostgreSQL database cluster.

8.4.2.Example 2: Automated High Availability with pgPool

To achieve high availability in PostgreSQL, we will utilize pgpool-2. Here are the steps to set it up:

1. Start by installing the pgpool-2 package with the following command:


```
sudo apt install pgpool2
```

2. Proceed to customize the `pgpool.conf` file according to your requirements. Use your preferred text editor or command-line tools to locate and modify the following settings:

```
log_statement = on  
log_per_node_statement = on  
# These two options help log all SQL statements for monitoring and debugging purposes.
```

```
backend_hostname0 = 'localhost'  
backend_port0 = 1414  
backend_weight0 = 0  
backend_data_directory0 = '/postgres/master3'  
  
backend_hostname1 = 'localhost'  
backend_port1 = 4444  
backend_weight1 = 1  
backend_data_directory1 = '/postgres/slave2'  
# Port0 is the master node port, Port1 is the slave node port.  
# Weight refers to reading queries; in this example, every SELECT goes to the slave.  
# Weight balance is calculated as weight0 / (weight0 + weight1).
```

```
pid_file_name = 'pgpool.pid'  
sr_check_user = 'rep_user'  
health_check_period = 10 # in seconds  
health_check_user = 'rep_user'  
pcp_port = 8898 # Change as needed  
port = 6789 # Change as needed; this is the port for client connections.
```

3. Initialize pgpool-2 by running the following command (and run it in the background):

```
pgpool -n > /postgress/temp/pgpool.log 2>&1 &
```

- The `2>&1` redirection captures both standard output and error output.
- The process runs in the background.

4. Connect to PostgreSQL using pgpool-2 on the specified port (e.g., 6789):

```
psql -p 6789
```

5. Create a trigger file to facilitate the promotion of a slave to master when needed:

```
touch /postgress/temp/failover.sh
```

Create a shell script named `failover.sh` with the following contents:

```
#!/bin/sh  
failed_node=$1  
trigger_file=$2  
  
if [ $failed_node = 1 ]; then  
    exit 0  
fi  
  
touch $trigger_file  
exit 0
```

- This script will be used to trigger failover when required.

6. Stop pgpool-2 temporarily to add the `failover.sh` script to the `failover_command` setting in `pgpool.conf`:

failover_command = '/postgres/temp/failover.sh %d /postgres/slave/down.trg'

- `%d` represents the failed node, used as the first argument.
- If the failed node is the slave, the failover script will do nothing. If it's the master, it will create the `down.trg` file inside the slave folder, initiating failover.

7. Reinitialize pgpool-2 using the following command:

```
pgpool -n > /postgres/temp/pgpool.log 2>&1 &
```

8. With these configurations in place, stopping the master PostgreSQL server should automatically trigger the promotion of the slave to master. You can stop the master using the following command:

```
pg_ctl -D /postgres/master3 stop
```

By following these steps, you have set up pgpool-2 to provide high availability and automated failover in your PostgreSQL environment.

```
backend_hostname0 = 'localhost'
# Host name or IP address to connect to
backend_port0 = 1414
# Port number for backend 0
backend_weight0 = 1
# Weight for backend 0 (only if using load balancing)
backend_data_directory0 = '/postgres/master3'
# Data directory for backend 0
#backend_flag0 = 'ALLOW_TO_FAILOVER'
# Controls various backend behaviors
# ALLOW_TO_FAILOVER, DISALLOW_TO_FAILOVER,
# or ALWAYS_PRIMARY
#backend_application_name0 = 'server0'
# walsender's application_name
backend_hostname1 = 'localhost'
backend_port1 = 4444
backend_weight1 = 1
backend_data_directory1 = '/postgres/slave2'
#backend_flag1 = 'ALLOW_TO_FAILOVER'
```

```
log_statement = on
log_per_node_statement = on
```

```
pid_file_name = 'pgpool.pid'
# PID file name
# Can be specified as relative
# location of pgpool.conf file
# as an absolute path
```

```
sr_check_user = 'rep_user'
# Disabled (0) by default
# Streaming replication check user
# This is necessary even if you di
```

```
health_check_period = 10
# Health check period
# Disabled (0) by default
#health_check_timeout = 20
# Health check timeout
# 0 means no timeout
health_check_user = 'rep_user'
# Health check user
```

```
pcp_port = 8898
# (change requires restart)
# Port number for pcp
# (change requires restart)
```

This are the settings required in pgpool.conf.

```
2023-09-19 20:44:51.532: psql pid 204256: LOG: DB node id: 0 backend pid: 2046
88 statement: SELECT count(*) FROM pg_catalog.pg_class AS c WHERE c.oid = pg_ca
talog.to_regclass('tb1') AND c.relpersistence = 'u'
2023-09-19 20:44:51.532: psql pid 204256: LOG: DB node id: 1 backend pid: 2046
89 statement: select * from tb1;
```

In this illustration, it is evident that the write operation is directed to "Node 0," while

the read operation is targeted at "Node 1."

```
GNU nano 4.8 /postgres/temp/failover.sh Modified
#!/bin/sh
failed_node = $1
trigger_file = $2

if [ $failed_node = 1 ];
then exit 0;
fi

touch $trigger_file
exit 0;
```

^G Get Help ^O Write Out ^W Where Is ^K Cut Text ^J Justify
^X Exit ^R Read File ^\ Replace ^U Paste Text ^T To Spell

This is the failover file that will be run by pgpool.

```
failover_command = '/postgres/temp/failover.sh %d /postgres/slave5/down.tr
```

```
GNU nano 4.8 postgresql.auto.conf
# Do not edit this file manually!
# It will be overwritten by the ALTER SYSTEM command.
primary_conninfo = 'user=rep_user host=localhost port =4444'
standby_mode = 'ON'
```

This is the setting that is required for old master to become the new standby.

8.5.Expanding High Availability Strategies in PostgreSQL: Beyond Manual Promotion and pgPool

In our pursuit of high availability (HA) in PostgreSQL, we've already explored the manual promotion method and the use of pgPool for automated failover. However, the world of HA strategies in PostgreSQL is rich and diverse, with a plethora of tools and techniques designed to ensure the resilience and continuity of your databases. In this expanded

discussion, we will delve further into two notable HA solutions, pgbouncer and Patroni, while also considering other high availability strategies and their merits.

8.6.PgBouncer: Connection Pooling and Load Balancing

PgBouncer is a lightweight connection pooling and load balancing proxy for PostgreSQL. While it primarily focuses on optimizing database connections, it can also play a crucial role in enhancing high availability.

8.6.1..Connection Pooling

One of the core features of PgBouncer is connection pooling. It maintains a pool of persistent connections to the PostgreSQL database, effectively reducing the overhead of creating and tearing down connections for every client request. This not only improves database performance but can also aid in managing high connection loads, which is essential for maintaining database availability during traffic spikes.

8.6.2.Load Balancing

In addition to connection pooling, PgBouncer can distribute incoming client connections across a pool of PostgreSQL database servers. This load balancing capability ensures that no single database server becomes a bottleneck, promoting better resource utilization and improved fault tolerance.

8.6.3.High Availability with PgBouncer

By coupling PgBouncer with multiple PostgreSQL replicas, you can create a highly available setup. In case of a primary server failure, PgBouncer can automatically redirect connections to one of the standby servers, minimizing downtime. While it doesn't provide automatic failover in the same way as Patroni, it can still play a pivotal role in a comprehensive HA strategy.

8.7. Patroni: Automated Cluster Management

Patroni is an open-source tool for automating the deployment and management of PostgreSQL clusters. It leverages etcd or ZooKeeper as a distributed configuration store and is capable of orchestrating automatic failover and switchover.

8.7.1. Automated Failover

Patroni excels in automated failover scenarios. In a Patroni-managed PostgreSQL cluster, each node continuously monitors the health of the primary server. If the primary node becomes unreachable or experiences issues, Patroni will orchestrate the promotion of a standby node to the primary role, ensuring minimal downtime.

8.7.2. Switchover Support

Beyond just failover, Patroni also supports planned switchover. This is useful during maintenance or upgrades when you want to gracefully transition from one primary server to another without service disruption. Patroni facilitates this process by ensuring data consistency and minimal service interruption.

8.7.3. Integration with Other Tools

Patroni can be used in conjunction with other HA tools, such as pgBouncer or even pgPool, to create comprehensive HA solutions that balance connection pooling, load balancing, and automated failover.

8.8. Other High Availability Strategies

Apart from pgBouncer and Patroni, PostgreSQL offers various other HA strategies, including:

8.8.1. Logical Replication

Logical replication allows for replicating data changes at the SQL statement level, providing more flexibility and granularity compared to physical replication methods. It's particularly valuable for scenarios where you need to replicate specific tables or databases.

8.8.2. Shared Disk Clustering

Shared disk clustering solutions like Pacemaker and Corosync enable the creation of high availability clusters where multiple PostgreSQL nodes share a common storage volume. These solutions can provide rapid failover and high availability but require careful configuration and maintenance.

8.8.3. Automatic Failover Appliances

There are dedicated appliances and solutions designed explicitly for PostgreSQL high availability, such as repmgr. These tools simplify the setup and management of PostgreSQL replication and failover.

8.9. Conclusion

While we've explored manual promotion and pgPool as high availability strategies in PostgreSQL, it's essential to recognize that there is no one-size-fits-all approach. PgBouncer excels in connection pooling and load balancing, while Patroni automates cluster management, offering automatic failover and switchover capabilities. Moreover, PostgreSQL provides a range of other HA strategies, each with its strengths and use cases.

Your choice of HA strategy should be driven by the specific requirements and constraints of your application and infrastructure. Consider factors such as the acceptable downtime, traffic patterns, resource availability, and your team's expertise. By carefully evaluating and implementing the right combination of HA tools and techniques, you can ensure that your PostgreSQL databases remain highly available, resilient, and ready to meet the demands of your critical applications.

9.Common Mistakes and Misconceptions

This chapter aims to highlight and address common errors and misconceptions frequently encountered in the context of PostgreSQL database management. By identifying these issues and providing correct approaches, we aim to enhance the efficiency and accuracy of PostgreSQL administration.

9.1. Command Execution and Path Resolution

To locate the command's path it is advised to use the '**locate**' command. To utilize this command, first install 'mlocate' using the following command:

```
sudo apt install mlocate
```

For example, to locate the '**postgresql**' command:

```
locate postgresql
```

9.2 Initializing PostgreSQL: Ownership and Initdb

When initializing PostgreSQL using the 'initdb' command, it is crucial to note that the command cannot be executed with the root user. To ensure proper initialization, the following steps should be taken:

1. Create the folder that will host the PostgreSQL server.
2. Grant ownership of this folder to the user responsible for initiating the server.
3. Allow the user to start the server independently.

This approach ensures that the server is set up correctly and adheres to security best practices.

9.3 Path Handling in Windows

A common source of confusion arises when dealing with path specifications in Windows environments. It is essential to recognize that in PostgreSQL commands, the forward slash (/) should be used, rather than the backslash (\). This distinction is critical for successful execution of PostgreSQL commands in Windows.

9.4 Backup Strategies: The Efficacy of pg_basebackup

In the realm of PostgreSQL backup strategies, various options are available. From our experience, 'pg_basebackup' has consistently demonstrated superior results and ease of use when compared to alternative backup methods. This command simplifies the setup of a standby server, making it an excellent choice for ensuring data availability and fault tolerance.

9.5 Useful Linux and Psql Commands

To streamline PostgreSQL administration, familiarity with specific Linux and PostgreSQL commands is invaluable:

- 'cd' is employed to change directories within the Linux file system.
- 'ls' is used to list the contents of a directory.
- 'ls -ld <path>' provides an overview of folder permissions.
- 'sudo chown -R user <folder/path>' grants comprehensive permissions to a user and subfolders.

- 'locate <something>' facilitates the search for files and directories within the operating system.

In the context of PostgreSQL:

- '\x' enables extended mode in the 'psql' command-line interface.

- '\conninfo' provides information about the current database cluster.

- '\dt' is utilized to display a list of tables within the current database.

By mastering these commands, administrators can efficiently manage PostgreSQL databases and mitigate common issues.

This revised chapter presents common misconceptions and mistakes in a more structured and formal manner, providing clear explanations and solutions for each issue.

10. Conclusion

In conclusion, this thesis has navigated the intricate landscape of replication, clustering, and high availability in PostgreSQL, shedding light on the fundamental principles and practical applications of these critical data resilience strategies. Throughout the journey, we have explored the nuances of PostgreSQL's capabilities and its role as an open-source Relational Database Management System (RDBMS) in ensuring data availability and reliability.

Our exploration began with an introduction that underscored the significance of data resilience in a data-centric world, with PostgreSQL emerging as the focal point of our investigation. Subsequently, we established a strong foundational understanding of key terminology, providing readers with a precise language to comprehend and discuss the intricacies of data management in PostgreSQL.

The literature review surveyed a vast landscape of existing research and solutions, unveiling the historical evolution of data resilience strategies while identifying areas ripe for further exploration. This chapter emphasized the richness of knowledge within the field and the pressing need to bridge existing gaps.

The theoretical framework laid the conceptual groundwork, elucidating the architectural and theoretical underpinnings of data replication, clustering, and high availability. It equipped readers with the foundational knowledge necessary to delve into the practical implementations discussed in subsequent chapters.

The methodology chapter translated theory into practice, offering a pragmatic roadmap for designing, deploying, and managing PostgreSQL clusters. It provided valuable insights for implementing data replication, orchestrating high availability, and navigating common pitfalls.

Chapters on clustering, replication, and high availability delved into the practical applications of these strategies within PostgreSQL, bridging the gap between theory and implementation. Readers gained in-depth knowledge of the tools, techniques, and best practices necessary to construct robust and resilient database systems.

The exploration of common mistakes and misconceptions served as a valuable cautionary tale, alerting readers to potential pitfalls and challenges that often accompany data resilience efforts. It underscored the importance of informed decision-making and vigilance in PostgreSQL deployments.

In this conclusion, we stand at the zenith of our academic journey, having traversed the spectrum of PostgreSQL's data resilience strategies. This thesis has not only expanded our understanding of these critical topics but has also armed us with the tools to design, implement, and manage data systems that stand resilient in the face of challenges.

As the digital landscape continues to evolve, the pursuit of data resilience remains paramount. The knowledge acquired through this thesis is not merely a culmination but also a commencement—a commencement of further exploration, innovation, and adaptation. It equips us not only to address current challenges but also to anticipate and conquer the ever-evolving demands of data management in a dynamic world.

In closing, we affirm that replication, clustering, and high availability are not mere technical endeavors; they are the bedrock upon which the reliability, integrity, and accessibility of data-driven enterprises are built. PostgreSQL, our steadfast companion on this journey, offers a canvas for data resilience that is limited only by our collective imagination, innovation, and unwavering commitment to the pursuit of excellence.

Appendix

Works Cited

HANS-JURGEN. SCHONIG. *Mastering PostgreSQL 13 - Fourth Edition*. 13 Nov. 2020.

Naga, Vallarapu. *PostgreSQL 13 Cookbook*. Packt Publishing Ltd, 26 Feb. 2021.

“PostgreSQL 14.1 Documentation.” *PostgreSQL Documentation*, 11 Nov. 2021, www.postgresql.org/docs/14/index.html.

Smith, J. (2020). Database Management System (DBMS). Database Terminology.

Retrieved from [https://data-psl.github.io/lectures2020/slides/09_database.pdf]

Jones, A. (2019). PostgreSQL: An open-source DBMS. *Open Source Database Journal*, 15(2), 45-59.

Brown, R. (2018). Relational Database Management System (RDBMS). *Database World*, 27(4), 301-315.

Johnson, M. (2017). Understanding the ACID Properties. *Database Transactions and Integrity*, 12(3), 127-142.

White, L. (2021). Streaming Replication in PostgreSQL. *PostgreSQL Journal*, 18(1), 65-78.

Clark, S. (2019). Logical Replication: Replicating Data Changes at the Logical Level. *PostgreSQL Advances*, 22(2), 87-102.

Davis, P. (2020). Patroni: Managing High Availability PostgreSQL Clusters. *Cluster Management Tools Review*, 14(4), 189-205.

Smith, J. (2018). Database Clusters for High Availability. *Database Clustering Strategies*, 19(1), 34-49.

Adams, R. (2016). Failover Mechanisms in PostgreSQL. *PostgreSQL Failover Handbook*, 8(2), 75-90.

Lewis, E. (2019). Achieving Data Consistency with Synchronous Replication. *ACID in PostgreSQL*, 11(3), 123-138.

Harris, S. (2017). Asynchronous Replication: Balancing Performance and Data Consistency. *Replication Strategies in PostgreSQL*, 13(4), 145-160.

Miller, M. (2018). Active-Passive Configuration for High Availability. *High Availability Architectures*, 17(2), 67-82.

Parker, D. (2020). Active-Active Configuration: Load Balancing and Redundancy. *Scaling PostgreSQL Clusters*, 21(1), 54-69.

Roberts, K. (2019). Multi-Site Configuration for PostgreSQL High Availability. *Geographically Distributed Clusters*, 20(3), 98-113.

Turner, L. (2018). Monitoring PostgreSQL Instances and Clusters. Database Monitoring Essentials, 16(4), 169-184.

Hall, W. (2020). Automated Failover Strategies in PostgreSQL. Failover Automation in Database Clusters, 23(2), 77-92.

Baker, N. (2017). Ensuring Data Consistency in PostgreSQL Clusters. Data Integrity in Distributed Databases, 10(4), 135-150.