

Month 3, Week 1: The Unabridged Learning Guide

Introduction: The Application's Memory

For the past month, we have constructed a powerful but flawed machine. Our Express.js application can handle requests, process data, and respond like a real backend server, but it suffers from amnesia. Every time it restarts, all the data it has created—the new users, the posts, the products—vanishes. Its memory is volatile.

This week, we give our application a permanent memory. We step into the world of **databases**.

A database provides **persistent storage**, a place on the disk where our application's data can live securely, surviving restarts, crashes, and deployments. This is the single most critical step in transforming our application from a learning exercise into a real-world service.

We will begin by understanding the two major philosophies of database design: the structured world of **SQL** and the flexible world of **NoSQL**. We will learn why, for the vast majority of applications with structured, relational data, SQL is the time-tested, reliable, and correct architectural choice. We will set up **PostgreSQL**, the world's most advanced open-source relational database, as our system of record.

Finally, we will master the two fundamental sublanguages of SQL. First, **Data Definition Language (DDL)**, the architectural language used to create the “blueprint” of our database—the tables, columns, and the relationships between them. Second, **Data Manipulation Language (DML)**, the language we use to bring that blueprint to life by creating, reading, updating, and deleting the data within it. By the end of this lesson, you will have the foundational skills to design and interact with the very heart of any backend application: its database.

Table of Contents

1. Module 1: The World of Databases
 - The Great Divide: SQL vs. NoSQL
 - Why We Choose PostgreSQL
 - Setting Up Your Database Environment
 2. Module 2: DDL - Architecting the Schema
 - Core DDL Commands
 - A Deeper Look at PostgreSQL Data Types
 - Constraints: Enforcing Data Integrity
 3. Module 3: DML - Manipulating the Data
 - INSERT: Creating Data
 - SELECT: Reading Data
 - UPDATE: Modifying Data
 - DELETE: Removing Data
 4. Take-Home Assessment: Designing a Multi-Table Schema
-

Module 1: The World of Databases

The Great Divide: SQL vs. NoSQL

- **SQL (Relational Databases):**
 - **Model:** Data is stored in tables, which are composed of rows and columns.
 - **Schema:** A strict, predefined structure. You must define your columns and their data types before you can store data.
 - **Relationships:** Relationships between tables are explicitly defined using primary and foreign keys. This is a core strength.
 - **Language:** The powerful, standardized Structured Query Language (SQL).
 - **Analogy:** A well-organized library where every book has a specific shelf and a card catalog entry (the schema) telling you exactly where to find it.
 - **Best for:** Applications with structured, related data where consistency and data integrity are paramount (e.g., e-commerce, financial systems, most traditional web applications).
- **NoSQL (Non-Relational Databases):**
 - **Model:** A broad category including Document stores (JSON-like objects), Key-Value stores, Column-Family stores, and Graph databases.
 - **Schema:** Flexible or “schema-on-read.” You can often store objects with different structures in the same collection.
 - **Relationships:** Relationships are less formally defined, often handled within the application code.
 - **Language:** Each NoSQL database has its own query language.
 - **Analogy:** A messy, creative workshop. You have bins for different types of materials (collections), but the items in each bin can vary wildly. It’s fast and flexible, but you are responsible for making sense of the organization.
 - **Best for:** Big Data applications, real-time systems, or applications with unstructured or rapidly changing data requirements.

Why We Choose PostgreSQL For our journey as backend architects, starting with SQL is essential. It teaches the fundamentals of data integrity, structure, and relational theory that are applicable everywhere. We choose PostgreSQL specifically because: * **It’s Open-Source and Free:** No licensing costs. * **It’s Extremely Powerful:** It supports advanced data types (like JSONB, for storing JSON documents within a relational structure), powerful indexing, and is known for its reliability. * **It’s SQL Compliant:** The SQL you learn with PostgreSQL is transferable to almost any other SQL database (MySQL, SQL Server, etc.).

Setting Up Your Database Environment

1. **Install PostgreSQL:** Use your system’s package manager.
 - **macOS:** `brew install postgresql`
 - **Ubuntu/Debian:** `sudo apt update && sudo apt install postgresql postgresql-contrib`
2. **Start the Service:** Your package manager will provide instructions to start the PostgreSQL server so it’s running in the background.
3. **Install a GUI Client:** While you can do everything through the command-line client `psql`, a GUI makes learning and exploration much easier.

- **DBeaver:** A free, open-source, and cross-platform client that can connect to almost any database. Highly recommended.
- **Postico:** A popular, user-friendly native macOS client.

Module 2: DDL - Architecting the Schema

DDL (Data Definition Language) commands do not touch the data itself; they build the “container” that will hold the data.

Core DDL Commands

- **CREATE TABLE:** Defines a new table.
- **ALTER TABLE:** Modifies an existing table (e.g., adds/removes a column).
- **DROP TABLE:** Permanently deletes a table and all its data.

A Deeper Look at PostgreSQL Data Types

Type	Description	Example
SERIAL	Auto-incrementing 4-byte integer.	1, 2, 3, ...
BIGSERIAL	Auto-incrementing 8-byte integer (for very large tables).	1, 2, 3, ...
INTEGER or INT	4-byte integer.	12345
BIGINT	8-byte integer.	9000000000
VARCHAR(n)	Variable-length string with a max length of n .	'Hello World'
TEXT	Variable-length string with no predefined limit.	A long blog post.
BOOLEAN	true or false .	true
NUMERIC(p, s)	Exact decimal number. p =total digits, s =digits after decimal.	123.45
TIMESTAMP WITH TIME ZONE (TIMESTAMPTZ)	A date and time, stored in UTC and converted to local time zone on retrieval. This is the recommended type for timestamps.	'2025-10-10 10:30:00-05'
DATE	Just the date (no time).	'2025-10-10'
UUID	A universally unique identifier.	'a0eebc99-9c0b-4ef8-bb6d-6bb9bd380a'
JSONB	Stores JSON data in an efficient binary format. You can query inside the JSON.	'{ "name": "Alex", "tags": ["dev", "js"] }'

Constraints: Enforcing Data Integrity

- **PRIMARY KEY:** The most important constraint. A column (or set of columns) that uniquely identifies a row. It implies both **UNIQUE** and **NOT NULL**.
 - **FOREIGN KEY:** Creates a link between a column in one table and the **PRIMARY KEY** of another, enforcing referential integrity.
 - **NOT NULL:** Ensures a column cannot be left empty.
 - **UNIQUE:** Ensures every value in the column is unique across all rows. A primary key is always unique, but you can have other unique columns (e.g., email).
 - **DEFAULT:** Provides a default value if one isn't specified during an **INSERT**.
 - **CHECK:** A powerful constraint that lets you define a custom rule. For example, to ensure a product price is always positive: `price NUMERIC CHECK (price > 0)`.
-

Module 3: DML - Manipulating the Data

DML (Data Manipulation Language) commands are the “CRUD” operations of SQL.

INSERT: Creating Data

```
INSERT INTO users (username, email, password_hash)
VALUES ('alex_architect', 'alex@example.com', 'some_bcrypt_hash_here');
```

It's a strong best practice to always specify the column names. This makes your code more readable and resilient to schema changes.

SELECT: Reading Data This is the most powerful DML command. * **SELECT * ...:** The asterisk is a wildcard for “all columns.” Convenient for exploration, but in application code, it's better to specify the exact columns you need. * **WHERE:** The filter clause. You can combine conditions with **AND** and **OR**.

```
SELECT id, username, email FROM users WHERE username = 'alex_architect';
```

UPDATE: Modifying Data The **WHERE** clause is critical. Without it, you will update every row in the table. There is no “undo” button.

```
UPDATE users
SET email = 'alex.arch@newdomain.com', username = 'alex_arch'
WHERE id = 1;
```

DELETE: Removing Data Even more dangerous than **UPDATE**. Without a **WHERE** clause, this command will empty your entire table.

```
DELETE FROM users WHERE id = 1;
```

DELETE vs. TRUNCATE * **DELETE FROM my_table;** removes all rows one by one. It is a DML operation. * **TRUNCATE my_table;** removes all rows instantly by deallocating the data pages. It is a DDL operation and is much faster for emptying large tables.

Take-Home Assessment: Designing a Multi-Table Schema

Objective: To demonstrate mastery of DDL and DML by designing and populating a small, relational database schema.

The Task: Write a single SQL script file named `schema.sql`. When executed, this script should create two tables, `authors` and `books`, and then populate them with some sample data.

1. The DDL Section: * **authors table:** * `id`: Auto-incrementing primary key. * `name`: Text, required. * `bio`: Text, optional. * **books table:** * `id`: Auto-incrementing primary key. * `title`: Text, required. * `publication_year`: Integer. * `author_id`: An integer that must reference the `id` in the `authors` table. If an author is deleted, their books should also be deleted.

2. The DML Section: * Write at least two `INSERT` statements for the `authors` table. * Write at least three `INSERT` statements for the `books` table, making sure to use the IDs of the authors you just created. * Write a `SELECT` query to retrieve all books written by a specific author. * Write an `UPDATE` query to change the `bio` of one of your authors.

Submission: Submit your `schema.sql` file via a Pull Request to your personal assignments repository. ‘