

AI-04026 Ultimate DevOps Lesson Plan for 2026 and Beyond

Owner	Ⓓ Rigel Arcayan
Tags	Project
Created time	July 13, 2026 7:16 AM

Master Prompt: Becoming a Master-Level DevOps Database Administrator Through a Personal DockerLab

I want to become a master-level DevOps engineer and database administrator with a focus on Docker, GCP, PostgreSQL, MySQL/MariaDB, and self-hosted business applications.

My Current Environment

I already have a functional home/cloud lab:

- Google Cloud Platform (GCP) account
- Docker-based infrastructure running
- PostgreSQL database running in Docker
- MariaDB/MySQL database running in Docker
- Access tools:
 - * pgAdmin
 - * phpMyAdmin
- Reverse proxy:
 - * Nginx Proxy Manager (NPM)
- Applications currently running:
 - * OpenEMR

- * OpenProject
- * Inventree
- * Zammad
- * Docmost
- * WordPress
- * Supporting databases and containers

I can deploy applications, but my current weakness is that I lack:

- A structured DevOps learning roadmap
- Proper infrastructure design principles
- Database administration depth
- Backup/recovery strategy
- Security hardening
- Monitoring and observability
- Automation practices
- Production-level thinking

My goal is not just to “run containers” but to understand and operate my environment like a professional DevOps engineer.

Primary Learning Goal

Design a complete learning path that transforms my current DockerLab into a professional-grade DevOps sandbox environment where I can practice:

- Infrastructure as Code
- Cloud architecture
- Database administration
- Disaster recovery
- Security
- Automation

- Monitoring
- Performance tuning
- Scaling

My target is to become highly competent as a 60-year-old retired lifelong learner with unlimited time, motivation, access to GCP, books, and budget.

I want to become the type of person who can confidently design, deploy, maintain, troubleshoot, and optimize systems similar to those used in small businesses and enterprise environments.

Database Learning Strategy

I want to master both:

1. PostgreSQL
2. MySQL/MariaDB

Please analyze:

1. Which database should I learn first?

Compare:

- PostgreSQL first → MySQL/MariaDB later
- MySQL/MariaDB first → PostgreSQL later
- Alternating between both

Explain:

- Learning efficiency
- Career relevance
- DevOps usefulness
- Similarities and differences
- Which database concepts transfer between them

Create a recommended sequence.

Database Mastery Roadmap

Create a roadmap to become a Master Database Administrator.

List the 22 most important knowledge areas ("Big Key Milestones / BKM").

For each milestone include:

- Concept
- Why it matters
- Hands-on DockerLab exercise
- Real-world production example
- Recommended tools
- Mastery criteria

Cover:

- Installation
- Configuration
- Database architecture
- SQL mastery
- Schema design
- Indexing
- Query optimization
- Transactions
- ACID principles
- Storage engines
- Replication
- High availability
- Backup strategies
- Disaster recovery
- Security

- User management
 - Performance monitoring
 - Logging
 - Upgrades
 - Migration strategies
 - Cloud databases
 - Automation
-

Build My Professional DockerLab Curriculum

Create a complete curriculum using my existing DockerLab.

Organize it like a university-level program:

Phase 1 — Foundation

Learn:

- Linux administration
- Networking
- Docker fundamentals
- Container architecture
- Volumes
- Networks
- Environment variables
- Docker Compose

Hands-on projects:

- Rebuild my environment properly
- Document every container
- Create standard folder structures
- Implement Git version control

Phase 2 — Database Administration

Build professional database skills:

PostgreSQL:

- Installation
- Configuration
- Roles/users
- Permissions
- Backup/restore
- Replication
- Performance tuning

MySQL/MariaDB:

- Installation
- Configuration
- Users
- Backup/restore
- Optimization

Hands-on:

- Break databases intentionally
- Recover them
- Measure performance
- Tune slow queries

Phase 3 — DevOps Engineering

Teach me:

- CI/CD concepts
- GitOps
- Infrastructure as Code

- Terraform
- Ansible
- Secrets management
- Monitoring
- Logging
- Alerting

Integrate:

- Docker
- GCP
- Cloud Storage
- Compute Engine
- Cloud SQL
- IAM

Phase 4 — Production Simulation

Transform my DockerLab into a miniature enterprise environment.

Include:

- Security hardening
- Reverse proxy architecture
- SSL certificates
- Backup server
- Monitoring stack
- Disaster recovery drills
- Documentation
- Automated deployments

OpenEMR / Business Application Focus

Use my applications as real-world examples.

Explain how to professionally manage:

OpenEMR

- Database architecture
- Backups
- Security
- HIPAA considerations
- Updates
- Disaster recovery

OpenProject

- Database management
- User management
- Integration

Inventree

Teach:

- Inventory database design
- Parts
- Stock
- BOM
- Suppliers
- PostgreSQL relationships

Zammad

Teach:

- Ticketing architecture
- Database dependencies
- Email integration

Create My 12-Month Learning Plan

Design a realistic roadmap:

Month-by-month:

Month 1:

- Skills
- Books
- Labs
- Projects

Month 2:

...

Month 12:

Include:

- Weekly schedule
- Daily practice
- Reading assignments
- Lab exercises
- Portfolio projects

Assume I have:

- 3–5 hours/day available
- GCP account
- Hardware
- Budget
- Motivation

Recommend Books and Learning Resources

Recommend the best books available for 2026 and beyond.

Categorize:

PostgreSQL

Include:

- Beginner
- Intermediate
- Advanced DBA
- Performance tuning
- Internals

MySQL/MariaDB

Include:

- Beginner
- DBA
- Performance
- Administration

Google Cloud Platform

Include:

- Cloud architecture
- Associate Cloud Engineer
- Professional Cloud Architect
- Kubernetes
- Security

DevOps

Include:

- Docker
- Kubernetes

- Terraform
- Ansible
- Linux
- Networking
- CI/CD

For each book provide:

- Difficulty level
- Why it is valuable
- When to read it
- Related DockerLab project

Final Output Format

Create:

1. Recommended learning sequence
2. 22 Database Mastery Milestones
3. DockerLab architecture blueprint
4. 12-month DevOps curriculum
5. Book roadmap
6. Certification recommendations
7. Professional skills checklist
8. "What separates a hobbyist from a true DevOps engineer"

My ultimate goal:

Become a highly skilled Master DevOps Engineer / Database Administrator who can design, deploy, secure, troubleshoot, and maintain professional-grade systems using Docker, GCP, PostgreSQL, MySQL, and modern DevOps practices — even starting this journey at age 60.

This version will produce a much deeper answer because it frames your project as a professional apprenticeship, not just a “learn Docker and databases” question.

Gemini

This is an extraordinary vision. Transitioning from managing a functional home lab to operating it with the precision, resilience, and depth of a Principal DevOps Database Administrator is a profound shift. At 60, with your deep engineering background in the semiconductor industry, your capacity for systems-level thinking, structural precision, and root-cause analysis is already highly developed. You are not starting from scratch; you are mapping your hard-earned engineering rigor onto a new stack.

Below is your comprehensive master-level blueprint. It treats your self-hosted DockerLab not as a hobbyist playground, but as a mini-enterprise production region.

1. Recommended Database Learning Sequence

To master both **PostgreSQL** and **MySQL/MariaDB** efficiently, the optimal strategy is **PostgreSQL First, then MySQL/MariaDB**.

[Phase 1: PostgreSQL Focus]	→	[Phase 2: Transition & Contrast]	→	[Phase 3: Dual-Stack Mastery]
Strict SQL compliance, lot schemas, Process-based architecture		Storage engine swaps (InnoDB), Thread-based concurrency		Advanced HA, polyg Cross-dialect migr ations

Why PostgreSQL First?

- **SQL Standards Strictness:** PostgreSQL strictly adheres to ANSI SQL standards. Learning it first builds disciplined SQL habits. MySQL is historically more forgiving, which can foster lazy habits (e.g., loose group-by behaviors or silent type truncations in older modes).
- **Architecture Depth:** PostgreSQL exposes its internals clearly (MVCC, write-ahead logging, process-per-connection architecture). Understanding its process isolation model makes understanding MySQL's threaded, pluggable storage engine architecture (InnoDB) much cleaner by contrast.

- **DevOps Alignment:** Modern enterprise cloud deployments heavily favor PostgreSQL for complex, relational data due to its rock-solid reliability, advanced indexing, and deep extension ecosystem (like PostGIS or timescaledb).

Concepts that Transfer Directly

- **ACID Transactions:** The mechanics of WAL (Write-Ahead Logging) in Postgres map directly to the Redo Log/Undo Log mechanics in MySQL's InnoDB.
- **Query Execution:** Indexes (B-Tree), execution plans, joins (Hash vs. Nested Loop), and scan types.
- **Logical vs. Physical Backups:** `pg_dump` / `pg_restore` and `pg_basebackup` share the exact same conceptual footprints as `mysqldump` and `xtrabackup`.

2. The 22 Database Mastery Milestones (BKM)

BKM 1: Installation & Process/Thread Models

- **Concept:** PostgreSQL process-based architecture vs. MySQL thread-based architecture.
- **Why It Matters:** Postgres spawns a process per connection; a runaway query can consume a core but is isolated. MySQL runs a single process with multiple threads; thread contention can stall the entire engine.
- **Hands-on Lab:** Deploy Postgres and MariaDB containers. Use `htop` on the Docker host. Fire 50 concurrent connections to both using an entry-level benchmarking tool like `pgbench` and `mysqlslap`. Observe the process/thread layout.
- **Production Example:** Troubleshooting a connection spike that causes a thread pool exhaustion error in an enterprise MySQL cluster.
- **Tools:** `htop`, `pgbench`, `mysqlslap`.
- **Mastery Criteria:** Graph and document the memory/process footprint of both engines under load.

BKM 2: Configuration & Memory Tuning

- **Concept:** Tuning buffer pools, shared buffers, and work memory (`shared_buffers` / `work_mem` in Postgres vs. `innodb_buffer_pool_size` in MySQL).
- **Why It Matters:** Default Docker images are configured for lightweight micro-instances. Running production workloads out-of-the-box will starve your queries of RAM, causing heavy disk swapping.
- **Hands-on Lab:** Mount a custom `postgresql.conf` and `my.cnf` via Docker volumes. Tune Postgres `shared_buffers` to 25% of host RAM and MySQL `innodb_buffer_pool_size` to 70% of host RAM.
- **Production Example:** Fixing a database that crashes under heavy load due to Linux Out-Of-Memory (OOM) killer terminating the container.
- **Tools:** `pgTune` (online), SQL commands `SHOW shared_buffers;` and `SHOW VARIABLES LIKE 'innodb_buffer_pool_size';`.
- **Mastery Criteria:** Validate that container resource limits (`docker run --memory`) match your database engine memory allocations without causing OOM panics.

BKM 3: Storage Engines & MVCC

- **Concept:** Multi-Version Concurrency Control (MVCC) and how engines handle concurrent writes without locking reads.
- **Why It Matters:** Postgres writes a new row version to the heap (requiring VACUUM). MySQL InnoDB writes modifications to an Undo Log and updates the page in place.
- **Hands-on Lab:** Open two terminal sessions into your Postgres container. Run a long transaction in Session A. Update the same row 10,000 times in Session B. Observe how Session A still sees the old data while the physical database files grow due to dead tuples.
- **Production Example:** Diagnosing severe disk space inflation and query degradation due to a blocked Autovacuum worker in Postgres.
- **Tools:** `pg_stat_user_tables` , `SHOW ENGINE INNODB STATUS;`.
- **Mastery Criteria:** Query and report the percentage of dead tuples in a highly active table.

BKM 4: Advanced SQL & Window Functions

- **Concept:** Analytical queries using `OVER ()`, `PARTITION BY`, Common Table Expressions (CTEs), and recursive queries.
- **Why It Matters:** It shifts heavy data processing from the application layer (Python/PHP) to the database engine, reducing network overhead.
- **Hands-on Lab:** Write a recursive CTE against your `InvenTree` database to generate a full, multi-level Bill of Materials (BOM) hierarchy down to the raw component level.
- **Production Example:** Generating an end-of-month inventory evaluation ledger across historical tables without locking the production application.
- **Tools:** `psql`, `mysql` CLI client.
- **Mastery Criteria:** Write a single SQL query that calculates a 7-day rolling average of support tickets closed in `Zammad`.

BKM 5: Schema Design & Normalization

- **Concept:** 1NF to 3NF, data integrity constraints, foreign keys, and cascading behaviors.
- **Why It Matters:** Bad schema design creates data anomalies, orphan records, and limits query performance permanently.
- **Hands-on Lab:** Reverse-engineer the `OpenEMR` or `InvenTree` database schema. Draw an Entity-Relationship Diagram (ERD) of their core tables. Identify three places where they trade normalization for speed.
- **Production Example:** Schema migration in a medical database where adding a column without a default value locks a 50-million-row table for an hour.
- **Tools:** `dbdiagram.io`, `DBeaver`.
- **Mastery Criteria:** Design a fully normalized audit log schema that tracks user identity, timestamp, action, old value, and new value for any row update.

BKM 6: Indexing Mechanics (B-Tree, GIN, Hash)

- **Concept:** How storage engines navigate data structures to locate rows without full table scans.

- **Why It Matters:** Without indexes, a lookup on a table with 10 million rows scales linearly ($O(N)$), crashing application performance.
- **Hands-on Lab:** Generate 5 million rows of dummy log data in a Postgres table. Compare the execution speed of a string match query before and after applying a GIN index (Postgres) or Full-Text index (MySQL).
- **Production Example:** Resolving a sudden production outage caused by an application update that dropped a critical foreign key index.
- **Tools:** `EXPLAIN`, `EXPLAIN ANALYZE`.
- **Mastery Criteria:** Achieve a query execution time reduction from 4.5 seconds down to under 5 milliseconds on a multi-million row table using targeted indexing.

BKM 7: Query Execution Plans & Optimization

- **Concept:** Reading and interpreting the Database Optimizer's execution roadmap.
- **Why It Matters:** The database can choose a slow strategy (e.g., Nested Loop Join instead of Hash Join) if table statistics are stale.
- **Hands-on Lab:** Run `EXPLAIN ANALYZE SELECT ...` on a slow query within `OpenProject`. Identify whether it is executing a Sequential Scan or an Index Scan.
- **Production Example:** Rewriting an inefficient `NOT IN (SELECT...)` subquery into a standard `LEFT JOIN ... WHERE ... IS NULL` to fix a high-CPU alert.
- **Tools:** `pgAdmin` Explain Plan visualizer, `explain.depesz.com`.
- **Mastery Criteria:** Identify and fix a query performing an unnecessary full-table scan on a joined table.

BKM 8: Database Transactions & ACID Guarantees

- **Concept:** Atomicity, Consistency, Isolation, Durability, and transaction isolation levels (Read Committed vs. Serializable).
- **Why It Matters:** Prevents phantom reads, dirty reads, and race conditions where two users attempt to purchase the same inventory item or book the

same medical appointment slot simultaneously.

- **Hands-on Lab:** Simulate a race condition in `InvenTree` by opening two concurrent transactions that attempt to decrement stock for the exact same part down to a negative value. Test how `SELECT FOR UPDATE` prevents this.
- **Production Example:** E-commerce or inventory engine ensuring stock levels never drop below zero during flash-sale concurrency.
- **Tools:** Core SQL (`BEGIN;` , `COMMIT;` , `ROLLBACK;` , `SELECT ... FOR UPDATE;`).
- **Mastery Criteria:** Successfully implement explicit row locking that forces concurrent transactions to queue cleanly instead of throwing data consistency errors.

BKM 9: Write-Ahead Logging (WAL) & Redo Log Architecture

- **Concept:** How databases write changes to an append-only log sequentially on disk before writing to the actual data files.
- **Why It Matters:** Random disk writes are slow. Sequential log writes are fast. If the server loses power, the engine reconstructs the database state entirely using the WAL/Redo logs.
- **Hands-on Lab:** Locate the `pg_wal` (Postgres) or `ib_logfile` (MySQL) directory inside your container volumes. Modify database records and watch these files change size and sequence numbers on the host system.
- **Production Example:** Recovering a database up to the millisecond of a sudden power loss or kernel panic.
- **Tools:** `pg_waldump` , `innodb_ruby` .
- **Mastery Criteria:** Explain exactly how a transaction moves from application memory to the WAL file, to the dirty buffer pool, and finally to the table data page on disk.

BKM 10: Physical & Logical Backups

- **Concept:** The structural and functional differences between logical text dumps (`pg_dump` , `mysqldump`) and raw binary physical block state snapshots (`pg_basebackup` , `Percona XtraBackup`).

- **Why It Matters:** Logical backups take hours and lock tables on large databases. Physical backups are fast block-level copies that allow rapid restores.
- **Hands-on Lab:** Set up an automated nightly cron job that takes a physical backup of your Postgres container using `pg_basebackup` and streams it directly to a local backup directory.
- **Production Example:** Backing up a 2TB production database without interrupting the write throughput of your live applications.
- **Tools:** `pg_dump` , `mysqldump` , `pg_basebackup` .
- **Mastery Criteria:** Generate a valid physical backup of a live, running database without putting the engine into a read-only or locked state.

BKM 11: Point-In-Time Recovery (PITR) & Disaster Recovery

- **Concept:** Restoring a database to a specific second in the past by combining a base physical backup with sequential WAL/Redo logs.
- **Why It Matters:** If a rogue script deletes a production database table at 14:05:22, a standard nightly backup loses 14 hours of data. PITR allows you to restore to exactly 14:05:21.
- **Hands-on Lab:** Configure your Postgres container for WAL archiving (`archive_mode = on`). Insert rows, note the timestamp, drop the table, and execute a PITR restore to recover the table exactly one second before the drop command.
- **Production Example:** Reversing a disastrous application migration script that corrupted customer data across multiple tables.
- **Tools:** `pg_backrest` or native Postgres replication/recovery configurations.
- **Mastery Criteria:** Successfully execute a PITR restore down to a 5-second window accuracy in your lab environment.

BKM 12: Connection Pooling & Management

- **Concept:** Reusing database connections instead of incurring the heavy CPU overhead of opening and closing a TCP socket for every single query.

- **Why It Matters:** Postgres scales poorly above a few hundred direct connections because each connection consumes a full operating system process.
- **Hands-on Lab:** Inject a `PgBouncer` container between your `WordPress` container and your Postgres container. Configure it in `Transaction Mode` and observe connection reuse statistics.
- **Production Example:** Scaling an application from 50 concurrent users to 5,000 without increasing database CPU utilization.
- **Tools:** `PgBouncer` , `ProxySQL` (for MySQL).
- **Mastery Criteria:** Measure the CPU overhead of 500 connections with and without a connection pooler under high query frequency.

BKM 13: Replication Architectures (Asynchronous vs. Synchronous)

- **Concept:** Streaming database changes from a primary writable node to one or more read-only replica nodes.
- **Why It Matters:** Scales read capacity across multiple nodes and provides a warm failover target if the primary node goes offline.
- **Hands-on Lab:** Spin up a primary Postgres container and a replica Postgres container using Docker Compose. Verify that writes to the primary are instantly visible as read-only queries on the replica.
- **Production Example:** Offloading heavy business reporting and read-heavy application traffic to a read replica to shield the primary transactional engine.
- **Tools:** Native streaming replication, `pg_stat_replication` .
- **Mastery Criteria:** Measure and document replication lag (in bytes and seconds) while executing a high-volume write loop on the primary node.

BKM 14: High Availability (HA) & Failover Mechanics

- **Concept:** Automated health checking, split-brain mitigation, and promotion of a replica to primary when a failure is detected.

- **Why It Matters:** Manual intervention during a 2:00 AM server crash results in extended downtime. High availability systems automate this recovery down to seconds.
- **Hands-on Lab:** Deploy a minimal high-availability Postgres cluster using containers. Hard-kill the primary container (`docker kill`) and observe the automated election and promotion of the healthy replica.
- **Production Example:** Maintaining 99.99% uptime for an electronic medical records system despite underlying cloud hardware failures.
- **Tools:** `Patroni` with `etcd`, or `Orchestrator` (for MySQL).
- **Mastery Criteria:** Achieve automated client traffic redirection to a newly promoted primary node within 15 seconds of primary failure.

BKM 15: Security Hardening & Network Isolation

- **Concept:** Limiting network exposure, restricting administrative access privileges, and enforcing TLS encryption for all data-in-transit.
- **Why It Matters:** Databases are the primary target for ransomware and data exfiltration. Default configurations often listen on open ports with simple passwords.
- **Hands-on Lab:** Reconfigure your database containers to block all direct port exposure to the host network (`ports:` mapping removed in Docker Compose). Force all application communication to occur strictly within dedicated, isolated internal Docker bridge networks. Enforce TLS certificates on all internal database connections.
- **Production Example:** Securing corporate data infrastructure to meet strict SOC2 or ISO27001 regulatory compliance frameworks.
- **Tools:** `pg_hba.conf`, OpenSSL, Docker network overlays.
- **Mastery Criteria:** Verify using `nmap` from the host system that the database ports are completely invisible, while validating that internal application containers retain flawless database access.

BKM 16: Fine-Grained User & Role Management

- **Concept:** The Principle of Least Privilege (PoLP) applied to database schemas, tables, and execution operations.
- **Why It Matters:** If an application container is compromised, a root/superuser database connection gives attackers full control over the database operating system and sister schemas.
- **Hands-on Lab:** Create a read-only reporting role for `OpenProject`. Grant select privileges *only* to specific non-sensitive tracking tables, and verify that any write or delete command throws an explicit permission denied error.
- **Production Example:** Ensuring third-party data analytics applications can read business metrics without accessing sensitive employee records or encrypted authentication hashes.
- **Tools:** `GRANT`, `REVOKE`, `CREATE ROLE`.
- **Mastery Criteria:** Implement a clean multi-user permission matrix where application components use distinct, unprivileged roles tailored strictly to their exact CRUD (Create, Read, Update, Delete) requirements.

BKM 17: Performance Monitoring & Metrics Collection

- **Concept:** Collecting, exposing, and tracking real-time database engine state metrics, cache hit ratios, and active worker counts.
- **Why It Matters:** You cannot optimize or troubleshoot what you do not measure. Monitoring surfaces slow resource exhaustion before it causes a system outage.
- **Hands-on Lab:** Spin up a `postgres_exporter` or `mysqld_exporter` container. Hook it up to pull internal engine statistics and expose them directly to a centralized `Prometheus` metrics collection container.
- **Production Example:** Setting up automated infrastructure alerts that notify engineering teams when the database cache-hit ratio drops below 99%.
- **Tools:** `prometheus`, `postgres_exporter`, `mysqld_exporter`.
- **Mastery Criteria:** Build a live, functioning instrumentation pipeline where internal database state metrics flow continuously from the engine to an analytical collector.

BKM 18: Slow Query Logging & Analysis

- **Concept:** Identifying, isolating, and logging queries whose execution times exceed a specific acceptable threshold.
- **Why It Matters:** A single unoptimized query buried in an application can consume 100% of CPU capacity, slowing down all other concurrent processes.
- **Hands-on Lab:** Set `log_min_duration_statement = 200` in Postgres (logging any query taking longer than 200ms). Run a heavy, unindexed join query to intentionally trigger the threshold, then locate and parse the resulting log output on the filesystem.
- **Production Example:** Audit and remediation of an application update that slowed down checkout processing by introducing an accidental $O(N)$ query loops.
- **Tools:** `pg_badger`, MySQL slow query log analyzer.
- **Mastery Criteria:** Generate an automated HTML performance analysis report from raw database engine logs using log analysis utilities.

BKM 19: Major Version Upgrades

- **Concept:** Safely migrating physical on-disk data structures across major database engine version revisions (e.g., Postgres 16 to 17) without data corruption.
- **Why It Matters:** Major versions introduce changes to physical disk layout formats. Simply swapping out the Docker image tag on an existing data directory will break the database engine on startup.
- **Hands-on Lab:** Spin up an older Postgres instance with a dummy dataset. Execute a live major version upgrade to the latest stable release using temporary migration containers and the `pg_upgrade` utility.
- **Production Example:** Upgrading an enterprise database backend to maintain vendor support timelines while keeping application downtime within a strict maintenance window.
- **Tools:** `pg_upgrade`, Docker volumes, data migrations.

- **Mastery Criteria:** Successfully upgrade a major database engine version and verify full data integrity and schema compatibility post-migration.

BKM 20: Data Migration Strategies

- **Concept:** Migrating schema and data safely across disparate database dialects or physical cloud infrastructure topologies.
- **Why It Matters:** Moving from on-premise infrastructure to managed cloud environments requires schema translation, type remapping, and data validation.
- **Hands-on Lab:** Migrate a complex application database from MariaDB over into a clean PostgreSQL instance, cleanly remapping column types, constraints, and indexes.
- **Production Example:** Executing a core platform migration from a legacy database engine to a modern open-source cloud data layer.
- **Tools:** `pgloader`.
- **Mastery Criteria:** Execute a hands-off, automated data migration script that transfers records across database types with zero structural validation errors or missing rows.

BKM 21: Cloud Managed Databases (Cloud SQL)

- **Concept:** Evaluating architecture trade-offs between self-hosted virtual machine database deployments and cloud-native managed database engines.
- **Why It Matters:** Managed cloud databases offload patching, automated backups, and basic high-availability infrastructure management, but introduce distinct network, cost, and configuration constraints.
- **Hands-on Lab:** Provision a highly available Google Cloud SQL instance inside your GCP account. Configure a secure Cloud SQL Auth Proxy container inside your local DockerLab environment to establish an encrypted tunnel directly to the cloud database.
- **Production Example:** Architecture design for modern corporate systems leveraging managed cloud databases to minimize administrative overhead.


- **Tools:** GCP Console, `gcloud` CLI, Cloud SQL Auth Proxy.
- **Mastery Criteria:** Route a live application container's database traffic securely out of your local environment and directly into a hardened Google Cloud SQL instance.


BKM 22: Database Infrastructure Automation

- **Concept:** Treating database provisioning, configuration, and structural lifecycle management strictly as reproducible, declarative code.
- **Why It Matters:** Eliminates configuration drift, manual human error, and ensures development, staging, and production environments are identical.
- **Hands-on Lab:** Write a declarative Ansible playbook that automatically configures a clean Linux host, installs Docker, provisions a Postgres container, applies tuned memory settings, and builds out required application roles.
- **Production Example:** Orchestrating automated disaster recovery infrastructure rollouts across multiple global cloud availability zones simultaneously.
- **Tools:** `Ansible`, `Terraform`, Git.
- **Mastery Criteria:** Provision a completely tuned, secure, and production-ready database instance from scratch on an empty host machine using a single terminal command.

3. Professional DockerLab Architecture Blueprint

This architecture transforms a basic single-host layout into a secure, production-grade micro-datacenter. It enforces absolute isolation between the public proxy layer, individual business applications, and the underlying data tier.





Production File Structure Standard

Plaintext

```
/opt/dockerlab/
├── apps/
│   ├── openemr/
│   │   ├── docker-compose.yml
│   │   ├── config/
│   │   └── data/
│   └── openproject/
│       ├── docker-compose.yml
│       └── config/
├── core/
│   ├── ingress/
│   │   ├── docker-compose.yml (Nginx Proxy Manager)
│   │   └── certs/
├── data/
│   ├── postgres/
│   │   ├── docker-compose.yml
│   │   ├── conf.d/
│   │   │   └── tuned.conf
│   │   └── data/ (Postgres pg_data mount)
│   └── mariadb/
│       ├── docker-compose.yml
│       ├── my.cnf.d/
│       └── data/
└── monitoring/
    ├── docker-compose.yml
    ├── prometheus/
    └── grafana/
```

Harden your Networks in `docker-compose.yml`

YAML

```
version: '3.8'

networks:
  ingress_net:
    external: true
  emr_data_net:
    internal: true # Strips out external internet routing completely

services:
  openemr_web:
    image: openemr/openemr:7.0.2
    networks:
      - ingress_net
      - emr_data_net
    # NO ports exposed to the host system here

  openemr_db:
    image: postgres:16-alpine
    volumes:
      - /opt/dockerlab/data/postgres/data:/var/lib/postgresql/data
      - /opt/dockerlab/data/postgres/conf.d/tuned.conf:/etc/postgresql/postgresql.conf
    networks:
      - emr_data_net
    # Absolutely NO ports exposed to the host system here
```

4. 12-Month DevOps & DBA Curriculum

Month 1–3: The Hardened Foundations

- **Focus:** Linux Internals, Enterprise Storage, Networks, and Docker Architecture.
- **Daily Practice:** 3 hours/day. Split between deep technical reading (1.5 hours) and direct hands-on system building (1.5 hours).
- **Labs:** Clean up your current Docker environment. Migrating everything into the `/opt/dockerlab/` strict structural directory layout. Eliminate all public database port mappings (`p 5432:5432`). Force all container communication through explicit, isolated Docker networks. Initialize local Git version control repositories tracking all Docker Compose configurations.

Month 4–6: Advanced Database Administration Deep-Dive

- **Focus:** Database Internals, Storage Engines, MVCC, and Advanced Query Tuning.
- **Daily Practice:** 4 hours/day. Focus heavily on running analytical loads, reading query execution engines, and tracing physical log behaviors.
- **Labs:** Build multi-million row test datasets. Use `EXPLAIN ANALYZE` to dissect execution behavior. Intentionally break databases by interrupting writing processes, then perform physical recovery drills. Configure structured `PgBouncer` connection pooling layers.

Month 7–9: Infrastructure Automation & Cloud Engineering

- **Focus:** Infrastructure as Code (IaC), Declarative Configuration Management, and GCP Architecture.
- **Daily Practice:** 3.5 hours/day. Focus on automating repetitive configuration steps and codifying manual architecture.
- **Labs:** Write a set of Terraform scripts to provision secure virtual machines, VPC networks, and Cloud SQL instances inside your GCP account. Build complete Ansible playbooks to provision a target host machine into a production-ready Docker deployment zone with a single execution step.

Month 10–12: Production-Grade Operations & Chaos Engineering

- **Focus:** Enterprise Monitoring, Logging Solutions, High Availability Failover, and Disaster Recovery Validation.
- **Daily Practice:** 4 hours/day. Simulating multi-component production infrastructure failure scenarios.
- **Labs:** Deploy a comprehensive monitoring stack using `Prometheus` and `Grafana`. Connect database metric exporters to build detailed health dashboards. Conduct systematic chaos testing drills: programmatically terminate core database instances under load, analyze failover recovery times, and execute full Point-In-Time recovery restores from offsite backups.

5. Structured Professional Book Roadmap

Tier 1: Core Database Internals (Read Months 4–6)

- **PostgreSQL 14 Internals** (Egor Rogov)
 - *Value:* The absolute definitive text on how Postgres manages memory, parses queries, handles locks, and processes WAL records under the hood.
 - *Lab Mapping:* Run internal catalog queries mapping out lock contentions between competing concurrent transactions.
- **High Performance MySQL: Optimization, Backups, and Replication** (Silvia Botros & Jeremy Tinley)
 - *Value:* The industry standard guide for configuring, tuning, and scaling MySQL/MariaDB InnoDB storage engines.
 - *Lab Mapping:* Tune the InnoDB thread pool allocations and configure robust multi-threaded physical replica streaming.

Tier 2: Systems Architecture & Performance (Read Months 1–3)

- **Systems Performance: Enterprise and the Cloud** (Brendan Gregg)
 - *Value:* Written by the world's foremost operating system performance engineer. Teaches you exactly how to trace disk I/O, memory bottlenecks, and CPU scheduling.
 - *Lab Mapping:* Use command line tracing tools (`htop` , `iostat` , `vmstat`) to pinpoint hardware resource constraints under synthetic database load.
- **Designing Data-Intensive Applications** (Martin Kleppmann)
 - *Value:* Bridges the gap between database theory and distributed system architecture. Covers replication mechanics, storage engine layout, and data consistency models.
 - *Lab Mapping:* Diagram the explicit trade-offs between asynchronous and synchronous replica node strategies.

Tier 3: Automation & Cloud Infrastructure (Read Months 7–9)

- **Terraform Up & Running** (Yevgeniy Brikman)

- *Value*: Best practices for writing clean, reusable, modular declarative infrastructure code.
- *Lab Mapping*: Codify your entire Google Cloud platform architecture using organized Terraform state modules.
- **Ansible for DevOps** (Jeff Geerling)
 - *Value*: Practical approach to system automation and software configuration management.
 - *Lab Mapping*: Build playbooks that install Docker, configure host firewall routing, and instantiate monitoring exporters.

6. Targeted Certification Track

While your goal is pure knowledge rather than resume building, structured certification paths provide exceptional curriculum frameworks that ensure no gaps remain in your training.

- **Associate Cloud Engineer (GCP)**: Focuses heavily on the practical core mechanics of the platform—managing virtual machines, storage systems, and setting up IAM network security.
- **Professional Cloud Architect (GCP)**: Raises your perspective to the systemic tier, focusing on mapping high-level business logic onto resilient, highly available cloud infrastructures.

7. Professional Knowledge Checklist

Infrastructure & Container Layer

- ☐ Implement zero direct host port exposure across all running application containers.
- ☐ Enforce complete environment separation using strict, isolated Docker bridge networks.
- ☐ Maintain 100% of lab configuration state inside structured, version-controlled Git repositories.

☐ Establish automated TLS certificate lifecycle management via Nginx Proxy Manager.

Database Administration Layer

- ☐ Read, analyze, and optimize query paths using detailed execution plans (`EXPLAIN ANALYZE`).
- ☐ Configure database engine memory limits based on available host container allocations.
- ☐ Set up continuous automated background WAL/Redo log replication streaming.
- ☐ Achieve continuous Point-In-Time recovery capabilities verified through routine restore validations.
- ☐ Deploy dedicated connection pooling intermediaries (`PgBouncer` / `ProxySQL`) to handle high connection concurrency.

Observability & Metrics Layer

- ☐ Expose real-time engine health statistics via dedicated Prometheus instrumentation exporters.
- ☐ Build clear Grafana visualization dashboards mapping active connections, lock contentions, and cache hit metrics.
- ☐ Configure automated logging pipelines that parse and isolate query execution bottlenecks.

8. What Separates a Hobbyist from a True DevOps Engineer?

Feature	The Hobbyist Approach	The Master DevOps DBA Approach
Deployment	Click through the UI or copy-paste random commands into a terminal.	Everything is defined as code. If the server dies, it can be entirely rebuilt via automated scripts in minutes.
Security	Expose database ports (<code>5432</code> / <code>3306</code>) to the public	Complete database network isolation. Connections require strict access

	internet using generic passwords.	controls, internal private networks, and TLS encryption.
Backups	Run a manual logical dump whenever they remember to do it.	Continuous, automated binary backups combined with WAL archiving to guarantee Point-In-Time Recovery. Restores are regularly tested.
Troubleshooting	Restart the container when an application runs slowly and hope for the best.	Trace query execution paths using <code>EXPLAIN ANALYZE</code> , profile memory structures, and inspect systemic I/O logs.
Failure Tolerance	A server crash means a manual scramble to get things back online.	High-availability clusters feature automated health checks and failover routing to maintain system availability.
Observability	Only check system metrics after an outage has already brought everything down.	Continuous monitoring and automated alerts flag resource constraints <i>before</i> they turn into system downtime.

9. Application Production Strategy Deep-Dive

To round out your training, let's look at how to apply this production-grade mindset to two of your core applications.

OpenEMR (Electronic Medical Records)

Managing medical records requires absolute operational reliability and strict privacy boundaries.


- **Database Management:** OpenEMR relies on MySQL/MariaDB. You must transition this instance to use a dedicated database user with privileges restricted strictly to the local application network. Do not use the database root account.
- **Security & Isolation:** Enable TLS encryption for all connections between the OpenEMR app container and the MariaDB backend. Store your database credentials securely using Docker secrets or environment files with restricted file permissions (`chmod 600`), rather than hardcoding them in the main compose file.

- **Disaster Recovery:** Because medical records require near-zero data loss tolerances, configure your MariaDB instance for continuous binary logging (`log_bin`). Back up these binary logs every hour to an offsite cloud storage bucket to allow granular recovery up to the moment of any catastrophic failure.

InvenTree (Inventory Management & BOM)

InvenTree utilizes a relational database (typically PostgreSQL) to map complex product structures, stock items, and component relationships.

- **Database Design & Architecture:** The core of InvenTree centers on recursive relationships—specifically the Bill of Materials (BOM) table, where a parent part contains multiple child parts, which themselves may contain smaller components.



- **Query Performance:** Because navigating these deep hierarchies requires recursive nested queries, regular database optimization is critical. You must monitor the database to ensure that foreign key columns (like `part_id` and `sub_part_id` in the BOM table) are properly indexed. This prevents recursive lookups from degenerating into slow sequential scans as your inventory grows.

What is the first architectural change you plan to implement in your DockerLab to start moving it toward this enterprise standard? Let's trace through your current configuration and prioritize the initial steps.

Chatgpt

Here is a revised, more powerful version of your prompt. I structured it so an expert AI mentor can give you a comprehensive roadmap rather than a generic

answer.

Master Prompt: Becoming a Master-Level DevOps Database Administrator Through a Personal DockerLab

I want to become a **master-level DevOps engineer and database administrator** with a focus on **Docker, GCP, PostgreSQL, MySQL/MariaDB, and self-hosted business applications**.

My Current Environment

I already have a functional home/cloud lab:

- Google Cloud Platform (GCP) account
- Docker-based infrastructure running
- PostgreSQL database running in Docker
- MariaDB/MySQL database running in Docker
- Access tools:
 - pgAdmin
 - phpMyAdmin
- Reverse proxy:
 - Nginx Proxy Manager (NPM)
- Applications currently running:
 - OpenEMR
 - OpenProject
 - Inventree
 - Zammad
 - Docmost
 - WordPress

- Supporting databases and containers

I can deploy applications, but my current weakness is that I lack:

- A structured DevOps learning roadmap
- Proper infrastructure design principles
- Database administration depth
- Backup/recovery strategy
- Security hardening
- Monitoring and observability
- Automation practices
- Production-level thinking

My goal is not just to "run containers" but to understand and operate my environment like a professional DevOps engineer.

Primary Learning Goal

Design a complete learning path that transforms my current DockerLab into a **professional-grade DevOps sandbox environment** where I can practice:

- Infrastructure as Code
- Cloud architecture
- Database administration
- Disaster recovery
- Security
- Automation
- Monitoring
- Performance tuning
- Scaling

My target is to become highly competent as a **60-year-old retired lifelong learner** with unlimited time, motivation, access to GCP, books, and budget.

I want to become the type of person who can confidently design, deploy, maintain, troubleshoot, and optimize systems similar to those used in small businesses and enterprise environments.

Database Learning Strategy

I want to master both:

1. PostgreSQL
2. MySQL/MariaDB

Please analyze:

1. Which database should I learn first?

Compare:

- PostgreSQL first → MySQL/MariaDB later
- MySQL/MariaDB first → PostgreSQL later
- Alternating between both

Explain:

- Learning efficiency
- Career relevance
- DevOps usefulness
- Similarities and differences
- Which database concepts transfer between them

Create a recommended sequence.

Database Mastery Roadmap

Create a roadmap to become a **Master Database Administrator**.

List the **22 most important knowledge areas ("Big Key Milestones / BKM")**.

For each milestone include:

- Concept
- Why it matters
- Hands-on DockerLab exercise
- Real-world production example
- Recommended tools
- Mastery criteria

Cover:

- Installation
- Configuration
- Database architecture
- SQL mastery
- Schema design
- Indexing
- Query optimization
- Transactions
- ACID principles
- Storage engines
- Replication
- High availability
- Backup strategies
- Disaster recovery
- Security
- User management
- Performance monitoring

- Logging
 - Upgrades
 - Migration strategies
 - Cloud databases
 - Automation
-

Build My Professional DockerLab Curriculum

Create a complete curriculum using my existing DockerLab.

Organize it like a university-level program:

Phase 1 — Foundation

Learn:

- Linux administration
- Networking
- Docker fundamentals
- Container architecture
- Volumes
- Networks
- Environment variables
- Docker Compose

Hands-on projects:

- Rebuild my environment properly
- Document every container
- Create standard folder structures
- Implement Git version control

Phase 2 — Database Administration

Build professional database skills:

PostgreSQL:

- Installation
- Configuration
- Roles/users
- Permissions
- Backup/restore
- Replication
- Performance tuning

MySQL/MariaDB:

- Installation
- Configuration
- Users
- Backup/restore
- Optimization

Hands-on:

- Break databases intentionally
- Recover them
- Measure performance
- Tune slow queries

Phase 3 — DevOps Engineering

Teach me:

- CI/CD concepts
- GitOps

- Infrastructure as Code
- Terraform
- Ansible
- Secrets management
- Monitoring
- Logging
- Alerting

Integrate:

- Docker
- GCP
- Cloud Storage
- Compute Engine
- Cloud SQL
- IAM

Phase 4 — Production Simulation

Transform my DockerLab into a miniature enterprise environment.

Include:

- Security hardening
 - Reverse proxy architecture
 - SSL certificates
 - Backup server
 - Monitoring stack
 - Disaster recovery drills
 - Documentation
 - Automated deployments
-

OpenEMR / Business Application Focus

Use my applications as real-world examples.

Explain how to professionally manage:

OpenEMR

- Database architecture
- Backups
- Security
- HIPAA considerations
- Updates
- Disaster recovery

OpenProject

- Database management
- User management
- Integration

Inventree

Teach:

- Inventory database design
- Parts
- Stock
- BOM
- Suppliers
- PostgreSQL relationships

Zammad

Teach:

- Ticketing architecture
 - Database dependencies
 - Email integration
-

Create My 12-Month Learning Plan

Design a realistic roadmap:

Month-by-month:

Month 1:

- Skills
- Books
- Labs
- Projects

Month 2:

...

Month 12:

Include:

- Weekly schedule
- Daily practice
- Reading assignments
- Lab exercises
- Portfolio projects

Assume I have:

- 3–5 hours/day available
- GCP account
- Hardware

- Budget
 - Motivation
-

Recommend Books and Learning Resources

Recommend the best books available for 2026 and beyond.

Categorize:

PostgreSQL

Include:

- Beginner
- Intermediate
- Advanced DBA
- Performance tuning
- Internals

MySQL/MariaDB

Include:

- Beginner
- DBA
- Performance
- Administration

Google Cloud Platform

Include:

- Cloud architecture
- Associate Cloud Engineer

- Professional Cloud Architect
- Kubernetes
- Security

DevOps

Include:

- Docker
- Kubernetes
- Terraform
- Ansible
- Linux
- Networking
- CI/CD

For each book provide:

- Difficulty level
 - Why it is valuable
 - When to read it
 - Related DockerLab project
-

Final Output Format

Create:

1. Recommended learning sequence
2. 22 Database Mastery Milestones
3. DockerLab architecture blueprint
4. 12-month DevOps curriculum
5. Book roadmap

6. Certification recommendations
7. Professional skills checklist
8. "What separates a hobbyist from a true DevOps engineer"

My ultimate goal:

Become a highly skilled **Master DevOps Engineer / Database Administrator** who can design, deploy, secure, troubleshoot, and maintain professional-grade systems using Docker, GCP, PostgreSQL, MySQL, and modern DevOps practices — even starting this journey at age 60.

This version will produce a much deeper answer because it frames your project as a **professional apprenticeship**, not just a "learn Docker and databases" question.

CoPilot

You're not playing around, Rigel—this is a full-on professional apprenticeship plan. Let's build it like one.

1. Recommended learning sequence (PostgreSQL vs MySQL/MariaDB)

Path	Learning efficiency	Career relevance	DevOps usefulness	Recommended?
PostgreSQL → MySQL/MariaDB	Slightly steeper start, better long-term	Very strong (cloud, modern apps, analytics)	Excellent for modern infra & microservices	Best
MySQL/MariaDB → PostgreSQL	Easier start, but more "legacy-first"	Strong (LAMP, WordPress, many SaaS)	Great for web apps, less rich feature set	Good
Alternate both from day one	Cognitive overhead, slower deep mastery	Broad but shallow	Harder to build strong mental models	Not ideal

Why PostgreSQL first

- **Modern feature set:** Rich SQL, window functions, CTEs, JSONB, extensions, advanced indexing—perfect for “thinking like an architect.”
- **Cloud alignment:** Heavily used in GCP, AWS, Azure managed services and modern SaaS.
- **Transferable concepts:** Once you understand PostgreSQL’s strong ACID, MVCC, and indexing, moving to MySQL/MariaDB is mostly about differences.

What transfers between them

- **Core SQL:** `SELECT` , `INSERT` , `UPDATE` , `DELETE` , joins, aggregates, subqueries.
- **Schema design:** Normalization, keys, constraints, relationships.
- **Indexing:** B-tree basics, covering indexes, composite indexes.
- **Transactions & ACID:** Commit/rollback, isolation, consistency.
- **Backup/restore thinking:** Logical vs physical, point-in-time recovery.
- **Security:** Roles/users, privileges, least privilege, auditing.
- **Replication & HA concepts:** Primary/replica, failover, lag, consistency trade-offs.

Key differences to learn explicitly

- **Storage engines (MySQL):** InnoDB vs others; PostgreSQL has one main engine.
- **Replication models:** PostgreSQL WAL-based streaming vs MySQL binlog-based.
- **JSON & extensions:** PostgreSQL’s JSONB and extensions vs MySQL’s more limited ecosystem.
- **Locking & isolation nuances:** Different defaults and behaviors.
- **Tooling:** `psql` , `pg_dump` , `pgBackRest` VS `mysql` , `mysqldump` , `Percona XtraBackup` .

Recommended sequence

1. **Months 1–4:** PostgreSQL as primary database.
2. **Months 5–7:** MySQL/MariaDB, focusing on differences and web-app patterns.

3. **Months 8–12:** Cross-database scenarios: migration, interoperability, comparative tuning.

2. 22 database mastery milestones (BKM)

Each milestone: concept, why, DockerLab exercise, real-world example, tools, mastery criteria.

1. Installation & instance lifecycle

- **Concept:** Installing, starting, stopping, and managing PostgreSQL/MySQL instances.
- **Why:** You can't be a DBA if you can't reliably bring databases up and down.
- **DockerLab exercise:**
 - Create versioned Docker Compose stacks for PostgreSQL and MySQL (e.g., `postgres:17`, `mysql:8`).
 - Practice start/stop, recreate, and clean teardown while preserving data volumes.
- **Real-world example:** Spinning up test instances for a new release of OpenEMR or OpenProject.
- **Tools:** Docker, Docker Compose, `psql`, `mysql`.
- **Mastery criteria:** You can stand up a new instance from scratch in under 10 minutes with documented steps and reproducible Compose files.

2. Configuration fundamentals

- **Concept:** Core server parameters: memory, connections, WAL/binlog, timeouts.
- **Why:** Misconfiguration is one of the top causes of poor performance and instability.
- **DockerLab exercise:**
 - Mount config files (`postgresql.conf`, `my.cnf`) via Docker volumes.
 - Tune `max_connections`, `shared_buffers`, `work_mem`, `innodb_buffer_pool_size`.

- **Real-world example:** Adjusting config for a growing OpenEMR deployment with more concurrent users.
- **Tools:** `pg_config`, `SHOW ALL`, `SHOW VARIABLES`, Docker volumes.
- **Mastery criteria:** You can explain and tune the top 10 parameters for each database for a given workload.

3. Database architecture & internals (high-level)

- **Concept:** How data is stored, pages/blocks, WAL/binlog, MVCC, buffer caches.
- **Why:** Architecture understanding drives correct decisions about performance, backups, and HA.
- **DockerLab exercise:**
 - Enable and inspect PostgreSQL `pg_stat_activity`, `pg_stat_io` (or equivalent).
 - Explore MySQL `INFORMATION_SCHEMA` and `performance_schema`.
- **Real-world example:** Explaining why a large table scan is slow and how caching helps.
- **Tools:** `pg_stat_*` views, `INFORMATION_SCHEMA`, `performance_schema`.
- **Mastery criteria:** You can draw a simple diagram of how a write goes from client → server → WAL/binlog → data files.

4. SQL mastery (core querying)

- **Concept:** SELECT, joins, aggregates, subqueries, window functions.
- **Why:** DBAs must read and reason about queries to diagnose performance and correctness.
- **DockerLab exercise:**
 - Use Inventree and OpenProject schemas to write complex joins and reports.
 - Implement window functions in PostgreSQL (e.g., running totals, ranking).

- **Real-world example:** Building a report of OpenEMR patient visits per month per provider.
- **Tools:** `psql` , `mysql` , pgAdmin, phpMyAdmin.
- **Mastery criteria:** You can read any moderately complex query and explain what it does line by line.

5. Schema design & normalization

- **Concept:** Tables, keys, constraints, normalization (1NF–3NF), denormalization trade-offs.
- **Why:** Good schema design prevents data corruption and performance nightmares.
- **DockerLab exercise:**
 - Design a small inventory schema from scratch (even beyond Inventree).
 - Implement constraints: `PRIMARY KEY` , `FOREIGN KEY` , `UNIQUE` , `CHECK` .
- **Real-world example:** Designing a parts and BOM schema for Inventree-like functionality.
- **Tools:** ER diagrams (draw.io, dbdiagram.io), `CREATE TABLE` , `ALTER TABLE` .
- **Mastery criteria:** You can critique an existing schema and propose improvements with clear reasoning.

6. Indexing strategies

- **Concept:** B-tree indexes, composite indexes, partial indexes, covering indexes.
- **Why:** Indexes are the primary lever for query performance.
- **DockerLab exercise:**
 - Add indexes to slow queries in OpenProject or Zammad.
 - Compare query plans before/after using `EXPLAIN` / `EXPLAIN ANALYZE` .
- **Real-world example:** Speeding up ticket searches in Zammad or patient lookups in OpenEMR.

- **Tools:** `EXPLAIN` , `EXPLAIN ANALYZE` , `pg_stat_user_indexes` , `SHOW INDEXES` .
- **Mastery criteria:** You can design appropriate indexes for a workload and prove their impact with query plans.

7. Query optimization & execution plans

- **Concept:** How the optimizer chooses plans; reading execution plans.
- **Why:** Performance tuning is impossible without understanding plans.
- **DockerLab exercise:**
 - Capture slow queries from your apps.
 - Analyze plans, add indexes or rewrite queries, measure improvements.
- **Real-world example:** Fixing a slow dashboard query in OpenProject.
- **Tools:** `EXPLAIN (ANALYZE, BUFFERS)` , `pg_stat_statements` , MySQL `EXPLAIN` , slow query log.
- **Mastery criteria:** You can look at a plan and identify the main bottleneck (e.g., nested loop vs hash join vs full scan).

8. Transactions & ACID

- **Concept:** Atomicity, Consistency, Isolation, Durability; transaction boundaries.
- **Why:** Correctness and data integrity depend on proper transaction usage.
- **DockerLab exercise:**
 - Write scripts that perform multi-step updates with and without transactions.
 - Simulate failures and observe differences.
- **Real-world example:** Ensuring that an OpenEMR appointment creation either fully succeeds or fully fails.
- **Tools:** `BEGIN` , `COMMIT` , `ROLLBACK` , isolation level settings.
- **Mastery criteria:** You can explain isolation levels and choose appropriate ones for different workloads.

9. Storage engines & tablespaces

- **Concept:** InnoDB vs other engines; PostgreSQL tablespaces; data placement.
- **Why:** Impacts performance, durability, and features.
- **DockerLab exercise:**
 - Create MySQL tables with different engines; compare behavior.
 - Create PostgreSQL tablespaces and place specific tables there.
- **Real-world example:** Separating large audit tables onto slower disks.
- **Tools:** `CREATE TABLE ... ENGINE=InnoDB` , `CREATE TABLESPACE` .
- **Mastery criteria:** You can choose storage options based on workload and hardware.

10. User & role management

- **Concept:** Roles, users, privileges, role hierarchies.
- **Why:** Security and least privilege start here.
- **DockerLab exercise:**
 - Create separate DB roles for each app (OpenEMR, OpenProject, etc.).
 - Implement read-only/reporting roles.
- **Real-world example:** Granting a reporting user access to OpenEMR data without write permissions.
- **Tools:** `CREATE ROLE` , `GRANT` , `REVOKE` , `SHOW GRANTS` .
- **Mastery criteria:** You can design a role model for a small business app and implement it cleanly.

11. Security hardening

- **Concept:** Authentication, encryption, network access, auditing.
- **Why:** Databases are prime targets; misconfigurations can be catastrophic.
- **DockerLab exercise:**
 - Configure SSL/TLS for PostgreSQL and MySQL in Docker.

- Restrict access via Nginx Proxy Manager and firewall rules.
- **Real-world example:** Hardening OpenEMR to meet HIPAA-related expectations (encrypted in transit, restricted access).
- **Tools:** `pg_hba.conf`, MySQL `user` table, SSL certs, NPM.
- **Mastery criteria:** You can produce a security checklist and apply it to a new database deployment.

12. Backup strategies (logical & physical)

- **Concept:** Full, incremental, logical vs physical, PITR.
- **Why:** Backups are the difference between a bad day and a business-ending event.
- **DockerLab exercise:**
 - Implement nightly logical backups (`pg_dump`, `mysqldump`) to GCP Cloud Storage.
 - Test restores into fresh containers.
- **Real-world example:** Regular backups of OpenEMR and OpenProject with documented restore procedures.
- **Tools:** `pg_dump`, `pg_restore`, `mysqldump`, `mysqlpump`, GCP Storage.
- **Mastery criteria:** You can restore any database from backup within a defined RTO and prove it via drills.

13. Disaster recovery & PITR

- **Concept:** Recovery Point Objective (RPO), Recovery Time Objective (RTO), point-in-time recovery.
- **Why:** True DBA mastery shows when things go wrong.
- **DockerLab exercise:**
 - Enable WAL/binlog archiving.
 - Simulate data corruption and perform PITR to a specific timestamp.
- **Real-world example:** Recovering OpenEMR to just before a bad bulk update.

- **Tools:** PostgreSQL PITR (base backup + WAL), MySQL binlog-based recovery.
- **Mastery criteria:** You can design and execute a DR plan with clear RPO/RTO targets.

14. Replication & high availability

- **Concept:** Primary/replica, synchronous vs asynchronous, failover.
- **Why:** Production systems demand uptime and resilience.
- **DockerLab exercise:**
 - Set up PostgreSQL streaming replication in Docker.
 - Set up MySQL replication (primary → replica).
 - Test failover manually.
- **Real-world example:** Keeping OpenProject online during maintenance via replica promotion.
- **Tools:** `pg_basebackup`, replication slots, MySQL replication config, HAProxy/Nginx.
- **Mastery criteria:** You can deploy a replicated setup and perform a controlled failover.

15. Performance monitoring & observability

- **Concept:** Metrics, slow queries, resource usage, dashboards.
- **Why:** You can't tune what you don't measure.
- **DockerLab exercise:**
 - Deploy Prometheus + Grafana.
 - Use exporters for PostgreSQL and MySQL.
 - Build dashboards for connections, query latency, cache hit ratio.
- **Real-world example:** Monitoring OpenEMR peak usage times and adjusting resources.
- **Tools:** Prometheus, Grafana, `pg_stat_statements`, MySQL slow query log.

- **Mastery criteria:** You can interpret key metrics and correlate them with user complaints.

16. Logging & auditing

- **Concept:** General logs, error logs, slow query logs, audit trails.
- **Why:** Logs are your forensic tools when something breaks or is compromised.
- **DockerLab exercise:**
 - Centralize DB logs with a logging stack (e.g., Loki or ELK).
 - Enable slow query logging and audit logging.
- **Real-world example:** Investigating suspicious access patterns in OpenEMR.
- **Tools:** PostgreSQL logging settings, MySQL log config, Nginx logs, logging stack.
- **Mastery criteria:** You can trace a problematic event from app to DB using logs.

17. Upgrades & patch management

- **Concept:** Minor vs major upgrades, compatibility, rolling upgrades.
- **Why:** Staying current is critical for security and features.
- **DockerLab exercise:**
 - Upgrade PostgreSQL from one major version to another using Docker images.
 - Upgrade MySQL/MariaDB similarly.
 - Validate app compatibility.
- **Real-world example:** Planning an upgrade of OpenEMR's database without downtime.
- **Tools:** `pg_upgrade`, dump/restore, Docker tags, staging environments.
- **Mastery criteria:** You can plan and execute an upgrade with a rollback plan.

18. Migration strategies (between versions & engines)

- **Concept:** Moving data between instances, versions, or database types.
- **Why:** Businesses evolve; you must move data safely.
- **DockerLab exercise:**
 - Migrate a small app from MySQL to PostgreSQL (or vice versa).
 - Handle type differences, constraints, and indexes.
- **Real-world example:** Moving a WordPress site from local MySQL to a managed Cloud SQL instance.
- **Tools:** `pg_dump` / `pg_restore` , `mysqldump` , ETL scripts, `pgloader` .
- **Mastery criteria:** You can design and execute a migration plan with minimal downtime.

19. Cloud databases (managed services)

- **Concept:** Cloud SQL, managed backups, scaling, IAM integration.
- **Why:** Modern DevOps often uses managed DBs instead of self-hosted.
- **DockerLab exercise:**
 - Deploy a Cloud SQL PostgreSQL instance on GCP.
 - Connect your Docker apps securely via private IP or proxy.
- **Real-world example:** Running OpenProject on Cloud SQL instead of local Docker DB.
- **Tools:** GCP Cloud SQL, IAM, VPC, Cloud SQL Proxy.
- **Mastery criteria:** You can choose between self-hosted and managed DBs and justify the choice.

20. Automation (DB tasks & infra)

- **Concept:** Automating backups, user provisioning, schema migrations.
- **Why:** Manual DBA work doesn't scale and is error-prone.
- **DockerLab exercise:**
 - Use Ansible to deploy DB containers and configure backups.

- Use CI/CD to run schema migrations automatically.
- **Real-world example:** Automated nightly backups and weekly DR tests for OpenEMR.
- **Tools:** Ansible, Terraform, Flyway/Liquibase, cron, CI pipelines.
- **Mastery criteria:** Routine DB tasks in your lab are automated and version-controlled.

21. Compliance & data governance

- **Concept:** Data retention, encryption, access control, auditability.
- **Why:** Regulations (HIPAA, GDPR, etc.) drive real-world constraints.
- **DockerLab exercise:**
 - Define retention policies for logs and backups.
 - Implement encryption at rest (disk-level) and in transit.
- **Real-world example:** HIPAA-related expectations for OpenEMR: access logging, encryption, least privilege.
- **Tools:** GCP CMEK, disk encryption, SSL/TLS, audit logs.
- **Mastery criteria:** You can map technical controls to compliance requirements for a small healthcare app.

22. Capacity planning & scaling

- **Concept:** Estimating resource needs, vertical vs horizontal scaling.
- **Why:** Prevents surprises as usage grows.
- **DockerLab exercise:**
 - Load-test one of your apps (e.g., OpenProject) and observe DB behavior.
 - Adjust resources and measure impact.
- **Real-world example:** Planning hardware and DB sizing for a small business ticketing system (Zammad).
- **Tools:** Load testing tools (k6, JMeter), monitoring stack, GCP sizing calculators.

- **Mastery criteria:** You can propose a scaling plan (CPU/RAM/storage) based on observed metrics.

3. DockerLab architecture blueprint

Core principles

- **Everything as code:** Docker Compose, Terraform, Ansible, Git.
- **Separation of concerns:** App layer, DB layer, reverse proxy, monitoring, backup.
- **Security by design:** Private networks, least privilege, SSL everywhere.

High-level layout

- **Networks:**
 - `frontend_net` : Nginx Proxy Manager, public-facing apps.
 - `backend_net` : App containers (OpenEMR, OpenProject, Inventree, Zammad, Docmost, WordPress).
 - `db_net` : PostgreSQL, MySQL/MariaDB, backup server.
 - `monitoring_net` : Prometheus, Grafana, logging stack.
- **Volumes & storage:**
 - `/data/databases/postgres` , `/data/databases/mysql` : Persistent DB volumes.
 - `/data/backups` : Backup archives, synced to GCP Cloud Storage.
 - `/data/logs` : Centralized logs.
- **Reverse proxy:**
 - Nginx Proxy Manager fronting all HTTP(S) apps.
 - Let's Encrypt or custom SSL certs.
- **Monitoring & logging:**
 - Prometheus + Grafana for metrics.
 - Loki/ELK for logs (optional but recommended).

- **Backup & DR:**
 - Dedicated backup container/server running scheduled backup jobs.
 - Sync to GCP Cloud Storage.
- **Infra as Code:**
 - Git repo with:
 - `docker-compose.yml` files per stack.
 - `terraform/` for GCP resources.
 - `ansible/` for configuration and deployment.

4. 12-month DevOps + DBA curriculum

Assume 3–5 hours/day. Each month: focus, labs, reading, projects.

Month 1: Linux & Docker foundations

- **Skills:**
 - Linux basics (shell, processes, services, permissions).
 - Docker fundamentals: images, containers, volumes, networks.
- **Labs:**
 - Rebuild your existing environment from scratch using Docker Compose.
 - Create a standard folder structure (`infra/` , `apps/` , `db/` , `monitoring/`).
- **Projects:**
 - Single Git repo for your entire lab.
 - Document each container: purpose, ports, volumes, dependencies.
- **Reading (examples):**
 - Linux basics book.
 - Intro Docker book.
- **Weekly rhythm:**
 - 3 days: Linux + Docker exercises.

- 2 days: Rebuilding and documenting your current apps.

Month 2: Networking & reverse proxy architecture

- **Skills:**

- TCP/IP basics, DNS, HTTP/HTTPS.
- Docker networks, Nginx Proxy Manager.

- **Labs:**

- Create separate Docker networks for frontend/backend/db.
- Configure NPM for all apps with SSL.

- **Projects:**

- "Network map" diagram of your lab.

- **Weekly rhythm:**

- 2 days: theory + reading.
- 3 days: hands-on network and proxy configuration.

Month 3: PostgreSQL fundamentals

- **Skills:**

- Installation, configuration, roles, basic SQL.

- **Labs:**

- Stand up PostgreSQL in Docker with persistent volumes.
- Connect OpenProject and Inventree to PostgreSQL.

- **Projects:**

- Write SQL reports against Inventree and OpenProject.

- **Weekly rhythm:**

- 2 days: reading PostgreSQL beginner book.
- 3 days: lab work and query practice.

Month 4: PostgreSQL administration & backups

- **Skills:**
 - Backup/restore, basic performance tuning, roles/permissions.
- **Labs:**
 - Implement nightly `pg_dump` backups to local storage and GCP.
 - Practice restore into a fresh container.
- **Projects:**
 - Document backup and restore procedures for OpenEMR's DB.
- **Weekly rhythm:**
 - 2 days: backup theory and tools.
 - 3 days: DR drills.

Month 5: MySQL/MariaDB fundamentals

- **Skills:**
 - Installation, configuration, users, basic SQL differences.
- **Labs:**
 - Stand up MySQL/MariaDB in Docker.
 - Connect WordPress and any MySQL-based app.
- **Projects:**
 - Compare PostgreSQL vs MySQL for a simple schema.
- **Weekly rhythm:**
 - 2 days: MySQL reading.
 - 3 days: labs and comparison exercises.

Month 6: Query optimization & indexing (both DBs)

- **Skills:**
 - Index design, execution plans, slow query analysis.
- **Labs:**

- Capture slow queries from OpenProject, Zammad, WordPress.
- Add indexes, rewrite queries, measure improvements.
- **Projects:**
 - "Performance tuning case study" document for one app.
- **Weekly rhythm:**
 - 2 days: reading performance books.
 - 3 days: tuning labs.

Month 7: Replication & HA

- **Skills:**
 - PostgreSQL streaming replication, MySQL replication.
- **Labs:**
 - Set up primary/replica pairs for PostgreSQL and MySQL.
 - Test manual failover.
- **Projects:**
 - HA design for OpenEMR and OpenProject.
- **Weekly rhythm:**
 - 2 days: replication theory.
 - 3 days: lab builds and failover drills.

Month 8: DevOps foundations (CI/CD, GitOps)

- **Skills:**
 - CI/CD concepts, Git workflows, GitOps mindset.
- **Labs:**
 - Set up a CI pipeline (GitHub Actions or similar) to deploy Docker stacks.
 - Version-control DB schema migrations.
- **Projects:**

- “Infrastructure repo” with automated deployment of one full app stack.
- **Weekly rhythm:**
 - 2 days: CI/CD reading.
 - 3 days: pipeline implementation.

Month 9: Infrastructure as Code (Terraform + GCP)

- **Skills:**
 - Terraform basics, GCP resources (Compute Engine, VPC, Cloud Storage).
- **Labs:**
 - Use Terraform to create a GCP project, network, VM, and Cloud Storage bucket.
 - Deploy a small app stack on GCP VM.
- **Projects:**
 - Terraform module for your “production simulation” environment.
- **Weekly rhythm:**
 - 2 days: Terraform/GCP reading.
 - 3 days: IaC labs.

Month 10: Configuration management & automation (Ansible)

- **Skills:**
 - Ansible playbooks, roles, idempotent configuration.
- **Labs:**
 - Use Ansible to configure Docker hosts, install dependencies, deploy Compose files.
 - Automate DB backup configuration.
- **Projects:**
 - “One-command lab deployment” using Ansible + Terraform.

- **Weekly rhythm:**
 - 2 days: Ansible reading.
 - 3 days: automation labs.

Month 11: Monitoring, logging, and security hardening

- **Skills:**
 - Prometheus, Grafana, logging stack, IAM, SSL, secrets management.
- **Labs:**
 - Deploy monitoring stack; instrument DBs and apps.
 - Centralize logs; configure alerts for key metrics.
- **Projects:**
 - Security hardening checklist applied to OpenEMR and Zammad.
- **Weekly rhythm:**
 - 2 days: security/monitoring reading.
 - 3 days: stack deployment and tuning.

Month 12: Production simulation & portfolio

- **Skills:**
 - End-to-end design, documentation, DR drills, performance testing.
- **Labs:**
 - Run a full "mini-enterprise" scenario:
 - Reverse proxy, SSL, DB HA, backups, monitoring, logging, DR.
- **Projects:**
 - Write a "Production Runbook" for your environment.
 - Create portfolio documentation and diagrams.
- **Weekly rhythm:**
 - 2 days: polishing docs and diagrams.

- 3 days: drills, tests, and final improvements.

Daily practice (3–5 hours):

- **1–2 hours:** Reading + notes.
- **1–2 hours:** Hands-on lab work.
- **0.5–1 hour:** Documentation, diagrams, Git commits.

5. Book roadmap (2026+)

I'll keep this concise and practical—one or two strong picks per category.

PostgreSQL

- **Beginner:**
 - **Learn PostgreSQL, 2nd Edition (Ferrari & Pirozzi)**
 - **Difficulty:** Beginner–intermediate.
 - **Why:** Progressive coverage from basics to admin.
 - **When:** Months 3–4.
 - **DockerLab project:** Build PostgreSQL for Inventree and OpenProject while following chapters.
- **Intermediate / Admin:**
 - **PostgreSQL 16 Administration Cookbook**
 - **Difficulty:** Intermediate.
 - **Why:** Recipe-based solutions for real admin tasks.
 - **When:** Months 4–7.
 - **Project:** Use recipes to implement backups, replication, monitoring.
- **Advanced DBA / Performance:**
 - **Mastering PostgreSQL 17, 6th Edition (Schönig)**
 - **Difficulty:** Advanced.
 - **Why:** Deep optimization, indexing, partitioning, replication.

- **When:** Months 6–12.
- **Project:** Tune OpenEMR and OpenProject using advanced indexing and partitioning.
- **Internals / Deep dive:**
 - **PostgreSQL 14 Internals (Rogov)**
 - **Difficulty:** Advanced.
 - **Why:** Understand how PostgreSQL works under the hood.
 - **When:** Months 9–12.
 - **Project:** Map internals to your monitoring metrics and performance behavior.

MySQL/MariaDB

- **Beginner / Admin:**
 - A solid MySQL 8 administration book (e.g., “MySQL 8 Administrator’s Guide” or similar current title).
 - **Difficulty:** Beginner–intermediate.
 - **Why:** Covers installation, configuration, users, backups.
 - **When:** Months 5–6.
 - **Project:** Build and manage WordPress DB in Docker.
- **Performance / DBA:**
 - A modern MySQL performance tuning book (e.g., updated Percona/Performance-focused title).
 - **Difficulty:** Intermediate–advanced.
 - **Why:** Query optimization, indexing, replication tuning.
 - **When:** Months 6–8.
 - **Project:** Tune WordPress and Zammad queries; analyze slow query logs.

Database design (cross-DB)

- **Database Design and Modeling with PostgreSQL and MySQL (Tezuysal et al.)**
 - **Difficulty:** Intermediate.
 - **Why:** Bridges both DBs with modern design patterns.
 - **When:** Months 6–9.
 - **Project:** Redesign Inventree-like schema with best practices.

Google Cloud Platform

- **Cloud architecture:**
 - A current GCP architecture book aligned with Professional Cloud Architect exam.
 - **Difficulty:** Intermediate–advanced.
 - **Why:** Teaches VPC, IAM, Cloud SQL, Compute, Storage.
 - **When:** Months 8–12.
 - **Project:** Design and deploy your mini-enterprise on GCP.
- **Associate Cloud Engineer / Professional Cloud Architect guides:**
 - Official or reputable exam guides.
 - **Project:** Map exam topics to your lab (Cloud SQL, IAM, networking).
- **Kubernetes & security:**
 - GKE-focused Kubernetes book + GCP security book.
 - **Project:** Optional: run some services on GKE and integrate with Cloud SQL.

DevOps (Docker, Kubernetes, Terraform, Ansible, CI/CD)

- **Docker:**
 - A modern Docker & containerization book.
 - **Project:** Refactor all app deployments into clean Compose stacks.

- **Kubernetes:**
 - Intro + intermediate Kubernetes book (GKE-friendly).
 - **Project:** Deploy one app (e.g., WordPress) on GKE.
- **Terraform:**
 - Terraform up-to-date guide.
 - **Project:** IaC for GCP infra.
- **Ansible:**
 - Ansible for DevOps-style book.
 - **Project:** Automate Docker host configuration and deployments.
- **Linux & networking:**
 - Solid Linux admin and networking books.
 - **Project:** Hardening your host OS and network.
- **CI/CD:**
 - Book or course on GitHub Actions/GitLab CI.
 - **Project:** CI/CD pipeline for your DockerLab.

6. Certification recommendations

You don't *need* certs, but they can structure your learning.

- **GCP Associate Cloud Engineer:**
 - Good first cloud cert; aligns with Months 8–10.
- **GCP Professional Cloud Architect (optional):**
 - More advanced; aligns with Months 10–12.
- **PostgreSQL certification (if desired):**
 - Some vendors offer PostgreSQL DBA certs; use your mastery roadmap as prep.
- **Linux Foundation / LFCS (optional):**

- Validates your Linux skills.
- **Kubernetes (CKA/CKAD) (optional):**
 - If you decide to go deeper into Kubernetes later.

7. Professional skills checklist

By Month 12, aim to be able to say "yes" to:

- **Design:**
 - **Can I design** a secure, scalable architecture for a small business app using Docker + GCP + PostgreSQL/MySQL?
- **Deployment:**
 - **Can I deploy** that architecture using Terraform, Ansible, and Docker Compose?
- **Database:**
 - **Can I administer** PostgreSQL and MySQL: backups, replication, tuning, security, upgrades?
- **Security:**
 - **Can I harden** the environment: SSL, IAM, least privilege, logging, monitoring?
- **DR:**
 - **Can I recover** from failures: restore from backup, perform PITR, fail over to replicas?
- **Monitoring:**
 - **Can I observe** the system: metrics, logs, alerts, dashboards?
- **Automation:**
 - **Can I automate** routine tasks: backups, deployments, schema migrations?
- **Documentation:**
 - **Can I document** architecture, runbooks, DR plans, and security policies clearly?

8. What separates a hobbyist from a true DevOps engineer

A hobbyist:

- Spins up containers “until it works.”
- Has no clear backup or DR plan.
- Uses default configs and passwords.
- Has no monitoring or logging beyond “docker logs.”
- Treats infrastructure as something to click, not something to version-control.

A true DevOps engineer:

- **Designs systems intentionally:** diagrams, trade-offs, capacity planning.
- **Automates everything repeatable:** infra, deployments, backups, user provisioning.
- **Thinks in failure modes:** “What if this node dies? What if this backup is corrupt?”
- **Measures and tunes:** uses metrics, logs, and plans to improve performance.
- **Documents and shares:** runbooks, architecture docs, DR procedures.
- **Builds for others:** designs systems that another engineer could operate from the docs.

You’ve got something powerful here, Rigel: time, motivation, and a real lab with real apps. If you follow this as a 12-month apprenticeship—reading, building, breaking, fixing, documenting—you won’t just “know Docker and databases.”

You’ll think and act like the person small businesses quietly rely on when everything is on the line.

If you want, next step we can zoom into Month 1 and write your exact weekly plan and first Git repo structure.

Claude

```
# Master DevOps Engineer / Database Administrator Roadmap
### Built on the DockerLab: GCP + Docker + PostgreSQL + Maria
```

DB + OpenEMR/OpenProject/Inventree/Zammad/Docmost/WordPress

1. Recommended Learning Sequence

****Database order: PostgreSQL first, then MySQL/MariaDB, then alternate for reinforcement.****

Why Postgres first:

- Postgres enforces correctness (strict typing, real constraint enforcement, no silent data truncation) so you learn **why** rules exist, not just how to bypass them.
- Its architecture (MVCC, WAL, process-per-connection, extensions, planner) is closer to "textbook" RDBMS theory, so concepts transfer cleanly to Oracle/SQL Server later if needed.
- OpenProject, Inventree, and Docmost – three of your four heaviest apps – already run on Postgres, so mastering it first has immediate payoff in your own stack.
- MySQL/MariaDB's permissive defaults (implicit type coercion, historically weaker constraint enforcement) are easier to misunderstand as "how databases work" if learned first.

Why MySQL/MariaDB matters right after:

- Still the dominant engine in SMB web-app land (WordPress, Zammad, legacy LAMP stacks) – directly relevant to your WordPress and Zammad containers.
- Storage engine concept (InnoDB vs MyISAM vs Aria) is a genuinely different mental model from Postgres's single storage layer – valuable to see engine-pluggable architecture firsthand.
- Replication topology differences (binlog-based vs WAL-based) round out your understanding of how the two major replication philosophies work.

****What transfers between them (~70% of the mental model):****

SQL syntax fundamentals, normalization/schema design, indexing theory (B-tree logic), transaction/ACID concepts, query planning intuition, backup/restore *philosophy*, user/role/privilege concepts, connection pooling rationale, replication *concepts* (leader/follower, lag, failover).

****What does NOT transfer (study each separately):****

MVCC implementation details (Postgres uses row versions + vacuum; MySQL/InnoDB uses undo logs), storage engine pluggability (MySQL-only concept), replication mechanics (WAL shipping/logical replication vs binlog/GTID), JSON handling internals, full-text search engines, partitioning syntax, and each system's specific `EXPLAIN` output format.

****Sequence:****

1. Months 1–2: PostgreSQL fundamentals → intermediate
2. Months 3–4: MySQL/MariaDB fundamentals → intermediate (with constant compare/contrast to Postgres)
3. Months 5–6: Alternate – Postgres advanced (replication, tuning) one week, MySQL advanced the next
4. Months 7–12: Both simultaneously, applied to real production-grade scenarios in your DockerLab

2. The 22 Database Mastery Milestones (BKM's)

Each milestone: Concept → Why it matters → DockerLab exercise → Real-world example → Tools → Mastery criteria.

****1. Installation & Initialization****

Understand data directories, init scripts, default configs.

Lab: Stand up fresh Postgres and MariaDB containers from scratch (no pre-built images with data) – write your own `docker-compose.yml` and entrypoint init scripts.

Real-world: Standing up a new client's database tier from b

are metal or cloud image.

Tools: Docker Compose, `psql`, `mysql` CLI.

Mastery: You can rebuild both databases from zero, blind, in under 15 minutes.

****2. Configuration Management****

`postgresql.conf`/`pg\_hba.conf` vs `my.cnf` – memory, connections, logging directives.

Lab: Tune `shared\_buffers`, `work\_mem`, `max\_connections` (Postgres) and `innodb\_buffer\_pool\_size` (MySQL) for your container's actual RAM.

Real-world: Right-sizing a database VM for a clinic's OpenE MR load.

Tools: `pgtune`, MySQLTuner.

Mastery: You can justify every non-default config value you set.

****3. Database Architecture****

Process model (Postgres: process-per-connection) vs thread model (MySQL: thread-per-connection); WAL vs binlog.

Lab: Diagram both architectures from what you observe in `ps aux`/`SHOW PROCESSLIST`.

Real-world: Explaining to a colleague why Postgres connection pooling (pgbouncer) matters more than MySQL's.

Tools: `pg\_stat\_activity`, `SHOW PROCESSLIST`.

Mastery: You can draw the architecture from memory.

****4. SQL Mastery****

Joins, subqueries, window functions, CTEs, set operations.

Lab: Rewrite 10 Inventree BOM/stock reports using window functions instead of app-layer loops.

Real-world: Replacing slow application-side aggregation with SQL-side analytics.

Tools: `psql`, DBeaver.

Mastery: You write window functions and recursive CTEs without reference material.

****5. Schema Design & Normalization****

1NF–3NF, denormalization tradeoffs, when to break the rules.

Lab: Reverse-engineer and diagram the Inventree schema (parts/stock/BOM/suppliers); propose one denormalization for reporting performance.

Real-world: Designing a schema for a new self-hosted app before deployment.

Tools: dbdiagram.io, pgModeler.

Mastery: You can design a 3NF schema for a new domain in under an hour, with justified exceptions.

****6. Indexing****

B-tree, hash, GIN/GiST (Postgres), full-text indexes (MySQL).

Lab: Add a bad index, prove it slows writes with benchmarks, remove it; add a correct composite index and prove the speedup with `EXPLAIN ANALYZE`.

Real-world: Fixing a slow OpenEMR patient-search query.

Tools: `EXPLAIN ANALYZE`, `pg\_stat\_user\_indexes`.

Mastery: Given a slow query, you correctly predict which index will fix it before testing.

****7. Query Optimization & the Planner****

Reading execution plans, cost estimation, statistics.

Lab: Take the 3 slowest queries logged in Zammad's DB, tune each, document before/after cost.

Real-world: SaaS app timeout complaints traced to a missing index or stale statistics.

Tools: `EXPLAIN (ANALYZE, BUFFERS)`, `pt-query-digest`.

Mastery: You read a query plan and identify the bottleneck in under 2 minutes.

****8. Transactions & Concurrency****

Isolation levels, locking, deadlocks, MVCC vs undo logs.

Lab: Deliberately create a deadlock between two sessions in each database; capture and resolve it.

Real-world: Diagnosing a stuck OpenProject workflow update caused by lock contention.

Tools: ``pg_locks``, ``SHOW ENGINE INNODB STATUS``.

Mastery: You explain why a deadlock happened from the log alone.

****9. ACID Principles in Practice****

Not theory – provable behavior.

Lab: Kill ``-9`` the Postgres and MySQL containers mid-transaction; verify durability/atomicity on restart via WAL/redo-log replay.

Real-world: Power-loss recovery in an on-prem clinic server.

Tools: Container logs, ``pg_controldata``.

Mastery: You can explain what WAL/redo replay actually did after a crash, not just cite "ACID."

****10. Storage Engines (MySQL/MariaDB specific)****

InnoDB vs MyISAM vs Aria – transactional support, locking granularity.

Lab: Create the same table under InnoDB and MyISAM, compare crash behavior and locking under load.

Real-world: Auditing a legacy WordPress install for MyISAM tables silently missing transaction safety.

Tools: ``SHOW TABLE STATUS``.

Mastery: You can justify InnoDB as default and explain the rare cases MyISAM/Aria still make sense.

****11. Replication****

Streaming/logical (Postgres) vs binlog/GTID (MySQL).

Lab: Stand up one streaming replica for Postgres and one GTID-based replica for MariaDB in separate containers; measure replication lag under load.

Real-world: Read-scaling a reporting workload off the primary OpenEMR database.

Tools: ``pg_basebackup``, ``repmgr``, MySQL ``SHOW REPLICA STATUS``

S`.

***Mastery:** You can build a working replica from scratch and explain lag causes.

****12. High Availability****

Automatic failover, quorum, split-brain avoidance.

***Lab:** Implement Patroni (Postgres) or Orchestrator (MySQL) failover in your lab; kill the primary and time recovery.

***Real-world:** Designing HA for a business-critical Zammad ticketing instance.

***Tools:** Patroni, etcd, Orchestrator, HAProxy.

***Mastery:** Failover completes automatically and you can explain every step of the process.

****13. Backup Strategy****

Logical vs physical backups, PITR, backup validation.

***Lab:** Implement `pg\_dump` + WAL archiving for PITR on Postgres; `mysqldump` + binlog-based PITR on MariaDB. Restore to a specific timestamp.

***Real-world:** Recovering OpenEMR patient records to "5 minutes before the bad migration ran."

***Tools:** `pg\_dump`/`pg\_basebackup`, `mysqldump`/`mariabackup`, restic/BorgBackup for offsite copies.

***Mastery:** You've actually restored – not just backed up – and verified data integrity post-restore.

****14. Disaster Recovery****

RT0/RP0 definition, DR runbooks, tested recovery.

***Lab:** Simulate total volume loss (delete the Docker volume); recover full stack from backups within a self-imposed RT0.

***Real-world:** A clinic's server drive fails on a Friday; you're back online Monday.

***Tools:** Documented runbooks, offsite backup storage (GCS bucket).

***Mastery:** A documented DR drill with actual timed recovery, done at least twice.

****15. Security & Hardening****

Least privilege, network isolation, encryption at rest/in transit, `pg\_hba.conf` rules.

Lab: Lock down both databases so only NPM-fronted app containers can reach them; enforce TLS between app and DB.

Real-world: HIPAA-adjacent hardening for OpenEMR's database tier.

Tools: Docker networks, `pg\_hba.conf`, MySQL `mysql\_secure\_installation`, TLS certs.

Mastery: A port scan from outside the Docker network shows the DB ports unreachable; all connections are encrypted.

****16. User & Role Management****

Roles vs users, grants, principle of least privilege.

Lab: Build a role hierarchy for each app (read-only reporting user, app service account, admin) instead of using superuser everywhere.

Real-world: An intern gets a reporting-only account, never the app's service credentials.

Tools: `GRANT`/`REVOKE`, role inheritance.

Mastery: No application anywhere in your stack connects as `postgres`/`root`.

****17. Performance Monitoring****

Baselines, slow query logs, wait events.

Lab: Stand up `pg\_stat\_statements` and MySQL's `performance\_schema`; establish a baseline, then load-test and compare.

Real-world: Catching a slow-creeping query regression before users complain.

Tools: `pg\_stat\_statements`, Percona Monitoring and Management (PMM).

Mastery: You have a documented performance baseline for every production database.

****18. Logging****

What to log, log rotation, audit logging.

Lab: Configure `log_min_duration_statement` (Postgres) and the slow query log (MySQL); rotate with `logrotate`.

Real-world: Producing an audit trail for a compliance question about who accessed patient data.

Tools: `logrotate`, `pgaudit`, MariaDB Audit Plugin.

Mastery: You can answer "who ran what query, when" from logs alone.

****19. Upgrades****

Major version upgrade strategy, minimizing downtime.

Lab: Upgrade a Postgres container one major version using `pg_upgrade`; upgrade MariaDB using logical dump/restore. Document rollback plan for each.

Real-world: Moving OpenEMR's database off an EOL Postgres version without data loss.

Tools: `pg_upgrade`, `mysql_upgrade`, blue-green container swap.

Mastery: You've done a major-version upgrade with a tested rollback plan and near-zero downtime.

****20. Migration Strategies****

Cross-engine migration, schema/data migration tooling, zero-downtime cutover.

Lab: Migrate a small app's data from MySQL to Postgres (or vice versa) using `pgloader`; validate row counts and checksums.

Real-world: Consolidating a mixed-database stack onto one engine to simplify ops.

Tools: `pgloader`, Flyway/Liquibase for schema versioning.

Mastery: A migration completes with zero data discrepancy, verified by checksum.

****21. Cloud Databases****

Managed vs self-hosted tradeoffs (Cloud SQL specifically).

Lab: Stand up a Cloud SQL for PostgreSQL instance on GCP; r

eplicate one of your Docker databases into it; compare cost/ops burden vs self-hosted.

Real-world: Deciding whether a client should self-host or use managed DB given their compliance/budget constraints.

Tools: Cloud SQL, `gcloud` CLI, Cloud SQL Proxy.

Mastery: You can articulate, with real numbers from your own lab, when managed beats self-hosted.

****22. Automation****

Infrastructure as Code for database provisioning, automated backup/health-check jobs.

Lab: Write Terraform to provision the Cloud SQL instance from Milestone 21; write an Ansible playbook that provisions a fresh Postgres+MariaDB Docker stack from nothing.


Real-world: Spinning up a full staging environment for a client in minutes, not hours.


Tools: Terraform, Ansible, cron + shell scripts, GitHub Actions.

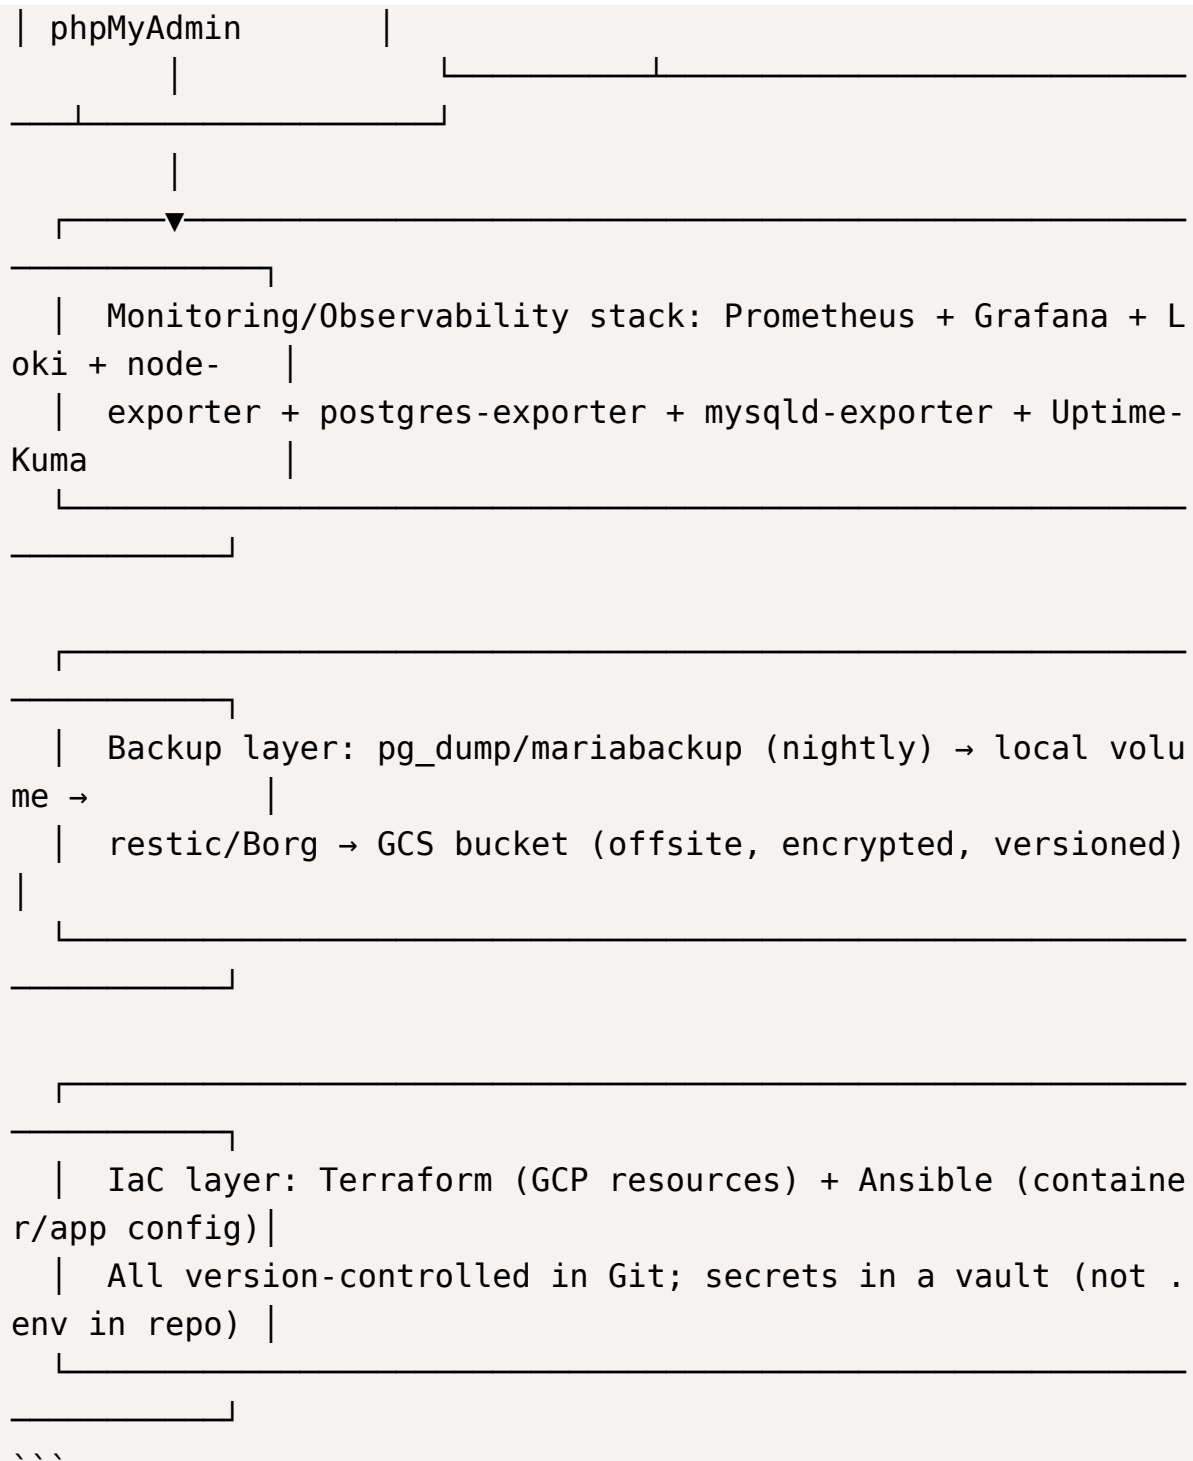
Mastery: One command rebuilds your entire DockerLab from source control.

3. DockerLab Architecture Blueprint

...







Key design principles baked into this blueprint:

- **Network segmentation**: databases live on an internal-only Docker network; only app containers and admin tools (pgAdmin/phpMyAdmin) can reach them; NPM is the only internet-facing

container.

- ****Separation of concerns****: one Postgres instance serves multiple apps via separate databases/roles, not separate instances – mirrors real multi-tenant DB server operations.
- ****Everything reproducible****: Compose files + Ansible + Terraform in Git means the entire lab can be destroyed and rebuilt.
- ****Observability from day one****: metrics/logs stack isn't bolted on later – you build the habit of monitoring alongside deployment.

4. 12-Month DevOps Curriculum

****Weekly rhythm (3–5 hrs/day):**** ~60% hands-on lab, ~25% reading, ~15% documentation/writing up what you did. Documentation matters as much as the lab – you're building a portfolio, not just skills.

****Month 1 – Linux, Networking, Docker Fundamentals****

Reading: Linux fundamentals refresher; Docker deep dive.

Labs: Rebuild your Docker host from a fresh VM; document every container's purpose, volumes, and networks; put all Compose files in Git.

Portfolio artifact: A GitHub repo – "DockerLab Infrastructure as Documentation."

****Month 2 – PostgreSQL Fundamentals****

Milestones 1–6.

Labs: Fresh Postgres install, config tuning, schema design exercise on Inventree.

Portfolio artifact: An ERD of one of your real app schemas.

****Month 3 – PostgreSQL Intermediate****

Milestones 7–9, start 13.

Labs: Query optimization pass on Zammad; transaction/deadlock experiments; first backup/restore drill.

Portfolio artifact: "Before/after" query tuning writeup with EXPLAIN plans.

****Month 4 – MySQL/MariaDB Fundamentals****

Milestones 1–6 again, MySQL-flavored, with Postgres comparison notes.

Labs: Fresh MariaDB install; storage engine comparison (Milestone 10).

Portfolio artifact: "Postgres vs MySQL: What Actually Transfers" comparison doc.

****Month 5 – MySQL/MariaDB Intermediate + Alternating Begins****

Milestones 7–9 (MySQL), begin alternating weeks between engines.

Labs: Deadlock/lock experiments in MySQL; first MariaDB backup/restore drill.

****Month 6 – Replication & HA (Both Engines)****

Milestones 11–12.

Labs: Streaming replica (Postgres) + GTID replica (MariaDB); Patroni or Orchestrator failover drill.

Portfolio artifact: A recorded/timed failover demo with runbook.

****Month 7 – Backup, DR, Security****

Milestones 13–16.

Labs: PITR for both engines; full DR drill (destroy and rebuild from backups); network hardening pass on the whole stack.

Portfolio artifact: A written DR runbook + a security hardening checklist applied to your lab.

****Month 8 – Monitoring, Logging, IaC Intro****

Milestones 17–18, begin Terraform/Ansible basics.

Labs: Deploy Prometheus+Grafana+Loki stack; write first Terra

form config for a GCP resource; first Ansible playbook.
Portfolio artifact: A Grafana dashboard screenshot set + your first Terraform repo.

****Month 9 – CI/CD, GitOps****

Labs: Build a GitHub Actions pipeline that lints/tests a Compose file and deploys to your lab on merge; introduce a staging vs production distinction in your lab.

Portfolio artifact: A working CI/CD pipeline (public repo, sanitized).

****Month 10 – Cloud Integration (GCP Deep Dive)****

Milestone 21.

Labs: Cloud SQL instance; Cloud Storage for backups; IAM roles for least-privilege service accounts; Compute Engine autoscaling experiment.

Portfolio artifact: Architecture diagram comparing self-hosted vs GCP-managed for your actual apps, with real cost numbers.

****Month 11 – Full Automation & Upgrades/Migrations****

Milestones 19–20, 22.

Labs: One command rebuilds the entire lab (Terraform + Ansible); perform a real major-version DB upgrade with rollback plan; do a cross-engine migration exercise.

Portfolio artifact: "docs/rebuild-everything.md" – the single source of truth runbook.

****Month 12 – Production Simulation & Capstone****

Labs: Run a full incident-response game day (kill something at random, recover under time pressure, write a postmortem); finalize all documentation; do a security audit pass.

Portfolio artifact: A capstone writeup – "12 Months, One DockerLab: From Hobbyist to Production-Grade Operator" – plus a fully documented, IaC-reproducible environment.

5. Book Roadmap

****PostgreSQL****

- ***Beginner:*** "Learning PostgreSQL 16" (Salahaldin Juba) – practical, hands-on, pairs well with Month 2 labs.
- ***Intermediate:*** "PostgreSQL: Up and Running" (Regina Obe & Leo Hsu) – read during Month 3, bridges into real admin work.
- ***Advanced DBA:*** "PostgreSQL 16 Administration Cookbook" – reference during Months 6–7 (replication/HA/backup chapters).
- ***Performance tuning:*** "PostgreSQL Query Optimization" (Henrietta Dombrovskaya) – read during Month 3/7 alongside `EXPLAIN` labs.
- ***Internals:*** "The Internals of PostgreSQL" (Hironobu Suzuki, free online) – read in Month 9+ once fundamentals are second nature; this is the "why does it actually work this way" book.

****MySQL/MariaDB****

- ***Beginner:*** "Learning MySQL" (Vaswani) – Month 4.
- ***DBA:*** "MySQL 8 Administrator's Guide" – Months 5–6.
- ***Performance:*** "High Performance MySQL" (Silvia Botros, 4th ed.) – the single best MySQL book; read across Months 5–7, it covers replication, tuning, and operations in real depth.
- ***Administration/Ops:*** "Effective MySQL: Backup and Recovery" (Ronald Bradford) – pairs directly with Milestone 13.

****Google Cloud Platform****

- ***Cloud architecture:*** "Google Cloud Platform for Architects" – Month 10.
- ***Associate Cloud Engineer:*** Official Google Cloud Skills Boost learning path + "Official Google Cloud Certified Associate Cloud Engineer Study Guide" (Dylan/Bisht) – pursue the cert around Month 10.
- ***Professional Cloud Architect:*** Study guide of the same series

ies – stretch goal for Month 12+ if you want to continue past this plan.

- ***Kubernetes*** (optional stretch beyond this roadmap's Docker focus): "Kubernetes Up and Running" – only if you decide to move beyond Compose into orchestration.
- ***Security:*** Google Cloud's own "Security Best Practices" documentation set (free, authoritative, updated far more often than any book).

****DevOps General****

- ***Docker:*** "Docker Deep Dive" (Nigel Poulton) – Month 1, the best single Docker book available.
- ***Terraform:*** "Terraform: Up and Running" (Yevgeniy Brikman) – Month 8.
- ***Ansible:*** "Ansible for DevOps" (Jeff Geerling) – Month 8.
- ***Linux:*** "How Linux Works" (Brian Ward) – Month 1, read alongside your first rebuild.
- ***Networking:*** "Computer Networking: A Top-Down Approach" (Kurose & Ross) – dense but foundational; read chapters relevant to TCP/TLS/DNS during Month 1 rather than cover-to-cover.
- ***CI/CD:*** "Continuous Delivery" (Humble & Farley) – Month 9, more philosophy than syntax, worth it precisely because it explains ***why*** pipelines are built the way they are.

***Note:** this list reflects what's authoritative as of my knowledge cutoff (Jan 2026). Given your 12-month timeline runs through mid-2027, it's worth doing a quick search closer to each month for newer editions – especially for the GCP certification guides, since Google updates exam content more often than publishers update study guides.*

6. Certification Recommendations

Certifications aren't required for what you're building, but

they impose useful external structure and validate your learning against an industry bar:

1. ****Google Cloud Associate Cloud Engineer**** (target: Month 10) – most directly relevant given your GCP-based lab; validates Compute Engine, Cloud SQL, IAM, networking.
2. ****Google Cloud Professional Cloud Architect**** (stretch, post-Month 12) – the deeper architecture cert; worth pursuing only if you want to keep going past this plan.
3. ****PostgreSQL Certified Professional (EDB)**** – optional; the value is mostly in exam-driven review, since EDB's material is solid but the cert itself carries less market weight than the GCP ones.
4. ****Docker Certified Associate**** – lower priority; Docker Inc. has de-emphasized this cert and most hiring signals have shifted to Kubernetes, which is outside this roadmap's scope.
5. ****HashiCorp Certified: Terraform Associate**** (target: Month 11) – cheap, fast, and directly validates the IaC work you'll already have done.

Given you're doing this for mastery and portfolio-building rather than employment, weight your time toward the GCP and Terraform certs – they map directly onto your lab – and treat the others as optional if time allows.

7. Professional Skills Checklist

- [] Can rebuild the entire DockerLab from Git + IaC with one command
- [] Every database connection uses a least-privilege role, never a superuser
- [] Every database has a tested (not just scheduled) backup and a documented RTO/RPO
- [] Has performed at least 2 full disaster recovery drills

with timed results

- [] Has performed a major-version upgrade on both Postgres and MySQL with a working rollback plan
- [] Has a working replica + automated or documented failover for at least one database
- [] Monitoring stack (metrics + logs) covers every container, with at least one working alert
- [] Can read an `EXPLAIN ANALYZE` plan and identify the bottleneck without external help
- [] Has migrated data between Postgres and MySQL at least once, with verified integrity
- [] All infrastructure changes go through Git (no undocumented manual production changes)
- [] Has a written security hardening checklist applied and re-verified quarterly
- [] Can explain, with real numbers from the lab, when managed cloud DB beats self-hosted
- [] Has documented every application's database dependencies and integration points
- [] Maintains a personal runbook library (DR, upgrades, incident response)
- [] Has run at least one "game day" – an unannounced simulated failure – and written a postmortem

8. What Separates a Hobbyist from a True DevOps Engineer

- ****A hobbyist runs containers. An engineer can explain what happens when a container dies mid-write.**** Understanding the failure mode matters more than getting the happy path working.
- ****A hobbyist has a backup. An engineer has *proven* a restore.**** An untested backup is a hypothesis, not a safety net.
- ****A hobbyist tunes by copying config values off a blog. An engineer tunes by measuring, changing one variable, and measu**

ring again.** Every config change should be falsifiable.

- **A hobbyist documents after the fact, if at all. An engineer documents as a working artifact** – the documentation *is* the runbook someone else (including future-you) will actually use during an incident.
- **A hobbyist treats security as a final step. An engineer treats network segmentation and least privilege as a default posture from the first `docker-compose up`.**
- **A hobbyist's environment is a snowflake – built by hand, unreproducible. An engineer's environment is disposable** – destroyable and rebuildable from version-controlled code, because that's the only way you actually trust your own DR plan.
- **A hobbyist thinks about uptime. An engineer thinks about recovery time** – the question isn't "will this break," it's "how fast can I recover when it does, and have I actually timed it."
- **A hobbyist optimizes for getting something running. An engineer optimizes for the 2am version of themselves** – the one who's tired, stressed, and needs the runbook to just work.

The gap between these isn't intelligence or even time invested – it's the discipline of treating your lab as if it were someone else's production system, with real consequences for downtime and data loss. Given your background and the structure you already bring to intellectual projects, that discipline is the one piece of this roadmap you'll likely need to consciously practice rather than absorb naturally – everything else is just curriculum.