

# **Docker Compose & Multi-Container Applications**

**Designing, Running, and Managing Containerized Applications with Docker Compose**

# Preface

When Docker revolutionized the way we build and ship software, it gave developers something extraordinary: the ability to package an application and its dependencies into a single, portable container. But it didn't take long for a fundamental truth to emerge—**real-world applications are almost never a single container**.

A modern web application might consist of an application server, a database, a caching layer, a message queue, and a reverse proxy—each running in its own container, each needing to communicate, share data, and start in the right order. Managing all of that with raw `docker run` commands quickly becomes unwieldy, error-prone, and impossible to reproduce reliably. This is precisely the problem that **Docker Compose** was designed to solve.

## Why This Book Exists

*Docker Compose & Multi-Container Applications* was written to fill a gap I've observed repeatedly in the Docker ecosystem. Countless tutorials teach you how to run a single Docker container. Far fewer guide you through the **design, orchestration, and management** of multi-container applications—the kind you'll actually build and deploy in the real world. This book bridges that gap.

Whether you're a developer spinning up a local environment, a DevOps engineer building CI/CD pipelines, or a team lead evaluating container orchestration strategies, this book provides a structured, practical path from your first `docker-compose.yml` file to production-ready multi-container deployments.

# What You'll Learn

The book is organized into a deliberate progression. We begin by examining **why single Docker containers fall short** for complex applications and then dive deep into the architecture of Docker Compose itself. From there, you'll get hands-on experience building your first multi-container application with Docker Compose before exploring the essential building blocks: *services and images, networks and inter-service communication, volumes and persistent data, and environment-based configuration*.

The middle chapters tackle the operational concerns that separate hobbyist setups from professional ones—**service dependencies and startup order, scaling services**, and leveraging Docker Compose for **local development** and **testing within CI pipelines**. We then turn our attention to production, covering best practices, real-world deployment patterns, and the common anti-patterns that trip up even experienced Docker users.

Finally, because no Docker Compose book would be complete without acknowledging the broader ecosystem, we close with a chapter on **transitioning from Docker Compose to Kubernetes**—helping you understand when and how to make that leap. The appendices serve as a lasting reference, including a Docker Compose command cheat sheet, YAML reference patterns, a troubleshooting guide, example multi-container stacks, and a container orchestration learning roadmap.

# Who This Book Is For

If you have a basic understanding of Docker—pulling images, running containers, writing a simple Dockerfile—you're ready for this book. No prior experience with

Docker Compose is required. By the time you finish, you'll be confident designing, running, debugging, and deploying multi-container Docker applications in any environment.

## The Tone and Approach

I've aimed to make this book **practical above all else**. Every concept is grounded in real scenarios, real configurations, and real mistakes I've seen (and made). Expect clear explanations, concrete examples, and honest guidance about what works—and what doesn't—when orchestrating Docker containers at scale.

## Acknowledgments

This book would not exist without the vibrant Docker community, whose open-source contributions, blog posts, forum discussions, and relentless curiosity continue to push containerization forward. I'm grateful to the technical reviewers who challenged my assumptions, the early readers who shaped the clarity of every chapter, and the Docker team itself for building tools that have fundamentally changed how we develop and deliver software.

---

*Containers changed everything. Docker Compose makes those containers work together. Let's get started.*

Dorian Thorne

# Table of Contents

---

<b>Chapter</b>	<b>Title</b>	<b>Page</b>
1	Why Single Containers Are Not Enough	6
2	Docker Compose Architecture Explained	20
3	Installing and Running Docker Compose	36
4	Your First Multi-Container Application	52
5	Services and Images	71
6	Networks and Service Communication	85
7	Volumes and Persistent Data	103
8	Environment Variables and Configuration	120
9	Service Dependencies and Startup Order	137
10	Scaling Services with Docker Compose	160
11	Docker Compose for Local Development	175
12	Testing and CI with Docker Compose	194
13	Production Best Practices	210
14	Docker Compose in Real Deployments	228
15	Common Docker Compose Anti-Patterns	248
16	From Docker Compose to Kubernetes	264
App	Docker Compose Command Cheat Sheet	279
App	docker-compose.yml Reference Patterns	294
App	Common Errors and Troubleshooting Guide	311
App	Example Multi-Container Stacks	326
App	Container Orchestration Learning Roadmap	348

---

# Chapter 1: Why Single Containers Are Not Enough

When you first begin working with Docker, the experience is nothing short of liberating. You write a Dockerfile, build an image, run a container, and suddenly your application is alive, isolated, portable, and reproducible. It feels like you have solved the deployment problem forever. You pull an Nginx image, serve a static website, and think to yourself, "This is all I will ever need." But then reality arrives. Your application grows. It needs a database. It needs a caching layer. It needs a message queue. It needs a reverse proxy sitting in front of multiple services. And suddenly, that single container you were so proud of starts to feel like a studio apartment when you have a family of six. It works, but it does not work well.

This chapter is about understanding why single containers, despite their elegance and power, are fundamentally insufficient for real-world applications. We will explore the limitations of running everything inside one container, examine the architectural philosophy that Docker was built upon, and build the case for why multi-container applications managed by tools like Docker Compose are not just a convenience but a necessity. By the end of this chapter, you will have a deep appreciation for the problems that multi-container orchestration solves, and you will be ready to embrace the patterns that professional teams use every day.

# The Monolithic Container Trap

Let us begin with a scenario that many Docker beginners encounter. You are building a web application. It has a Python Flask backend, a PostgreSQL database, and a Redis cache for session management. You know Docker. You have written Dockerfiles before. So you think, "Why not put everything into one container?" You install Python, PostgreSQL, and Redis inside a single image. You write a startup script that launches all three processes. You build the image, run the container, and everything works. For a while.

This approach is what experienced Docker practitioners call the "monolithic container" or the "fat container" pattern, and it is almost universally considered an anti-pattern. Here is why.

Docker containers are designed around a single-process model. When you run a container, Docker monitors the main process, the one specified by the CMD or ENTRYPOINT instruction in your Dockerfile. If that process exits, the container stops. If that process is healthy, Docker considers the container healthy. This model is elegant in its simplicity. One container, one process, one responsibility.

When you stuff multiple processes into a single container, you break this model in several important ways. First, you lose process isolation. If your PostgreSQL process crashes, your container might keep running because the main process, perhaps your startup script, is still alive. Docker has no way of knowing that a critical component of your application has failed. Second, you lose the ability to scale individual components. If your application is receiving heavy traffic and needs more web server instances, you cannot scale just the web server. You have to duplicate the entire container, database and all, which is wasteful and architecturally unsound. Third, you make updates and maintenance significantly more difficult. If you need to upgrade PostgreSQL from version 14 to version 15, you have to rebuild

the entire image, which means redeploying your application code and your Redis cache as well, even though nothing about them has changed.

Consider this example of a monolithic Dockerfile that tries to do everything:

```
FROM ubuntu:22.04

RUN apt-get update && apt-get install -y \
    python3 \
    python3-pip \
    postgresql \
    redis-server \
    supervisor

COPY ./app /app
COPY supervisord.conf /etc/supervisor/conf.d/supervisord.conf

RUN pip3 install -r /app/requirements.txt

EXPOSE 5000 5432 6379

CMD ["/usr/bin/supervisord"]
```

This Dockerfile installs three completely different services into one image. It uses Supervisor, a process management tool, to keep all three running simultaneously. The resulting image is large, difficult to maintain, and violates every principle of container design. When something goes wrong, and it will, debugging becomes a nightmare because logs from three different services are intermingled, resource usage is opaque, and failure isolation is nonexistent.

The following table summarizes the key differences between the monolithic container approach and the multi-container approach:

<b>Aspect</b>	<b>Monolithic Container</b>	<b>Multi-Container Approach</b>
Process Management	Multiple processes managed by a supervisor tool inside one container	Each container runs a single process managed by Docker

---

Failure Isolation	One crashed process may go undetected; container appears healthy	Each container is independently monitored; failures are immediately visible
Scalability	Cannot scale individual services; must duplicate the entire container	Each service can be scaled independently based on demand
Image Size	Large image containing all dependencies for all services	Smaller, focused images that contain only what each service needs
Update and Maintenance	Any change requires re-building the entire image	Services can be updated independently without affecting others
Logging	Logs from all services are mixed together inside the container	Each container produces its own log stream, making debugging straightforward
Security	Larger attack surface; a vulnerability in one service exposes all others	Smaller attack surface per container; services are isolated from each other
Resource Allocation	Cannot allocate CPU or memory limits to individual services	Docker allows per-container resource constraints
Networking	All services share the same network namespace	Services communicate over defined Docker networks with explicit rules
Reusability	The image is specific to this exact combination of services	Individual service images can be reused across different projects

---

This table makes the case clearly. The monolithic container approach sacrifices nearly every advantage that Docker provides.

# The Unix Philosophy and Docker

Docker's design philosophy is deeply rooted in the Unix philosophy, which can be summarized as: "Do one thing and do it well." In Unix, you have small, focused tools like `grep`, `sed`, `awk`, and `sort`, each of which performs a single function. The power comes from composing these tools together using pipes and redirection. The same principle applies to Docker containers.

A container should do one thing and do it well. A web server container serves web requests. A database container manages data storage and retrieval. A cache container handles in-memory data caching. A message queue container manages asynchronous communication between services. Each container is a specialist, and the application emerges from the collaboration of these specialists.

This philosophy has profound practical implications. When your web server container is a specialist, you can choose the best base image for it, perhaps a slim Alpine-based Python image that is only 50 megabytes. When your database container is a specialist, you can use the official PostgreSQL image maintained by the PostgreSQL community, which is optimized, tested, and regularly updated. You do not have to be an expert in configuring PostgreSQL because the official image handles the complexity for you.

Let us look at what this separation looks like in practice. Instead of one monolithic Dockerfile, you would have three distinct containers:

For the Flask application:

```
FROM python:3.11-slim

WORKDIR /app
COPY requirements.txt .
RUN pip install --no-cache-dir -r requirements.txt
COPY . .

EXPOSE 5000
```

```
CMD ["python", "app.py"]
```

For PostgreSQL, you would simply use the official image:

```
docker run -d --name db -e POSTGRES_PASSWORD=secret postgres:15
```

For Redis, again, the official image:

```
docker run -d --name cache redis:7-alpine
```

Each of these containers is small, focused, and independently manageable. The Flask container knows nothing about how PostgreSQL is configured. The Redis container does not care about your Python dependencies. Each can be updated, scaled, and debugged in isolation.

## The Challenge of Managing Multiple Containers

If the multi-container approach is so clearly superior, why does anyone ever create monolithic containers? The answer is simple: managing multiple containers manually is tedious and error-prone.

Consider what it takes to run our three-container application using only the Docker CLI. You need to create a network so the containers can communicate:

```
docker network create myapp-network
```

Then you need to start each container, connecting it to the network and configuring the appropriate environment variables:

```
docker run -d \
--name db \
--network myapp-network \
-e POSTGRES_PASSWORD=secret \
```

```

-e POSTGRES_DB=myapp \
-v pgdata:/var/lib/postgresql/data \
postgres:15

docker run -d \
--name cache \
--network myapp-network \
redis:7-alpine

docker run -d \
--name web \
--network myapp-network \
-e DATABASE_URL=postgresql://postgres:secret@db:5432/myapp \
-e REDIS_URL=redis://cache:6379 \
-p 5000:5000 \
myapp:latest

```

This is already three long commands, and we have a simple application. Now imagine you need to stop everything:

```

docker stop web cache db
docker rm web cache db

```

And if you want to clean up the network and volumes:

```

docker network rm myapp-network
docker volume rm pgdata

```

Now imagine doing this for an application with ten services. Or twenty. Imagine doing it on a new developer's machine when they join your team. Imagine doing it in a CI/CD pipeline where everything must be automated and reproducible. The manual approach does not scale. It is fragile, undocumented, and prone to human error. Did you remember to create the network before starting the containers? Did you spell the environment variables correctly? Did you start the database before the web server that depends on it?

This is the exact problem that Docker Compose was created to solve. Docker Compose allows you to define your entire multi-container application in a single

YAML file, specifying all the services, networks, volumes, environment variables, and dependencies in one place. Instead of running a dozen Docker commands, you run one:

```
docker compose up
```

And instead of tearing everything down with multiple commands:

```
docker compose down
```

But we are getting ahead of ourselves. The purpose of this chapter is not to teach you Docker Compose yet. It is to make you feel the pain of not having it, so that when we introduce it in the next chapter, you understand exactly why it exists and what problems it solves.

## Real-World Applications Are Inherently Multi-Service

Let us step back and think about the applications we use every day. Consider an e-commerce platform. At a minimum, it has a web frontend, an API backend, a database for product and order data, a search engine for product discovery, a cache for frequently accessed data, a message queue for processing orders asynchronously, and perhaps a separate service for sending emails and notifications. Each of these components has different resource requirements, different scaling characteristics, different update frequencies, and different failure modes.

The web frontend might need to be scaled to handle thousands of concurrent users during a sale event, while the database remains a single, powerful instance. The search engine might need to be rebuilt and reindexed without affecting the rest of the application. The email service might experience temporary failures

when a third-party SMTP provider goes down, but this should not crash the entire application.

These are not edge cases. This is the normal architecture of any non-trivial application. And Docker, with its container-per-service model, is perfectly suited to this architecture. But only if you have a way to manage all these containers as a cohesive unit. That is the role of Docker Compose, and that is what the rest of this book will teach you.

Here is a table that illustrates common services in a typical web application and why each deserves its own container:

<b>Service</b>	<b>Purpose</b>	<b>Why It Needs Its Own Container</b>
Web Server (Nginx)	Serves static files, acts as reverse proxy	Needs to be independently configurable and scalable
Application Server (Flask, Node, etc.)	Handles business logic and API requests	May need multiple instances for load balancing
Relational Database (PostgreSQL, MySQL)	Persistent data storage	Requires dedicated storage volumes and specific resource limits
Cache (Redis, Memcached)	In-memory data caching for performance	Lightweight and stateless; scales differently than other services
Message Queue (RabbitMQ, Kafka)	Asynchronous communication between services	Must be highly available and independently monitored
Search Engine (Elasticsearch)	Full-text search capabilities	Resource-intensive; needs its own memory and CPU allocation

---

Background Worker (Celery, Sidekiq)	Processes long-running tasks asynchronously	Scales based on queue depth, not user traffic
Monitoring (Prometheus, Grafana)	Observability and alerting	Should be isolated from application services to remain operational during failures

---

Each row in this table represents a service that has unique characteristics. Trying to combine even two or three of these into a single container would result in a brittle, unmanageable system. Combining all of them would be madness.

## Exercise: Experiencing the Pain First-hand

The best way to understand why single containers are not enough is to experience the limitations yourself. This exercise will guide you through creating a monolithic container and then breaking it apart into separate containers, so you can feel the difference.

### Step 1: Create a project directory

```
mkdir single-vs-multi && cd single-vs-multi
```

### Step 2: Create a simple Python application that uses Redis

Create a file called `app.py`:

```
from flask import Flask
import redis
import os

app = Flask(__name__)
cache = redis.Redis(host=os.environ.get('REDIS_HOST',
'localhost'), port=6379)
```

```

@app.route('/')
def hello():
    count = cache.incr('hits')
    return f'This page has been visited {count} times.\n'

if __name__ == '__main__':
    app.run(host='0.0.0.0', port=5000)

```

Create a requirements.txt:

```

flask==3.0.0
redis==5.0.1

```

### Step 3: Try the monolithic approach

Create a file called Dockerfile.monolithic:

```

FROM python:3.11-slim

RUN apt-get update && apt-get install -y redis-server && rm -rf /var/lib/apt/lists/*

WORKDIR /app
COPY requirements.txt .
RUN pip install --no-cache-dir -r requirements.txt
COPY app.py .

COPY start.sh /start.sh
RUN chmod +x /start.sh

EXPOSE 5000

CMD ["/start.sh"]

```

Create the start.sh script:

```

#!/bin/bash
redis-server --daemonize yes
python app.py

```

Build and run:

```
docker build -f Dockerfile.monolithic -t myapp-monolithic .
docker run -d -p 5000:5000 --name monolithic myapp-monolithic
```

Test it:

```
curl http://localhost:5000
```

Now, try to see the Redis logs:

```
docker logs monolithic
```

Notice that you only see the Flask logs. The Redis logs are hidden because Redis is running as a background daemon. If Redis crashes, you will not know about it from Docker's perspective. The container will keep running, but your application will start throwing errors.

#### **Step 4: Clean up the monolithic container**

```
docker stop monolithic && docker rm monolithic
```

#### **Step 5: Try the multi-container approach**

Create a file called Dockerfile:

```
FROM python:3.11-slim

WORKDIR /app
COPY requirements.txt .
RUN pip install --no-cache-dir -r requirements.txt
COPY app.py .

EXPOSE 5000
CMD ["python", "app.py"]
```

Now run two separate containers:

```
docker network create myapp-net
```

```
docker run -d --name redis --network myapp-net redis:7-alpine
```

```
docker run -d --name web --network myapp-net \
```

```
-e REDIS_HOST=redis \
-p 5000:5000 \
myapp-multi
```

Wait, we need to build the multi-container image first:

```
docker build -t myapp-multi .
docker run -d --name web --network myapp-net \
-e REDIS_HOST=redis \
-p 5000:5000 \
myapp-multi
```

Test it:

```
curl http://localhost:5000
```

Now check the logs for each service independently:

```
docker logs web
docker logs redis
```

Notice how each container has its own clean log output. You can monitor each service independently. If Redis crashes, its container stops and Docker reports it immediately. You can restart Redis without touching the web application. You can replace Redis with a newer version without rebuilding your application image.

### **Step 6: Clean up**

```
docker stop web redis
docker rm web redis
docker network rm myapp-net
```

**Note:** Even in this simple two-container example, you had to run six commands just to start the application and six more to clean it up. You had to remember to create the network first, set the correct environment variables, and use the right container names for DNS resolution. This is manageable for two containers but be-

comes impractical for larger applications. This is precisely the gap that Docker Compose fills, and we will explore it in depth starting in the next chapter.

## Looking Ahead

You now understand the fundamental problem. Single containers are powerful but limited. Real applications require multiple services working together. Managing those services manually with Docker CLI commands is tedious, error-prone, and does not scale. What we need is a declarative way to define our entire application stack, a way to describe all our services, their configurations, their networks, and their dependencies in a single file, and then bring the entire stack up or down with a single command.

That tool is Docker Compose, and it is the subject of the rest of this book. In the next chapter, we will introduce Docker Compose from the ground up, starting with the YAML file format, the structure of a Compose file, and the basic commands you need to manage multi-container applications. Everything you have learned in this chapter, the pain of monolithic containers, the elegance of the single-process model, the complexity of manual multi-container management, will serve as the foundation for understanding why Docker Compose is designed the way it is.

The journey from a single container to a fully orchestrated multi-container application is one of the most important progressions in modern software development. It mirrors the broader industry shift from monolithic architectures to microservices, from manual operations to infrastructure as code, from fragile deployments to reproducible environments. Docker Compose is your first step on this journey, and by the end of this book, you will be equipped to design, build, and manage containerized applications of any complexity with confidence and precision.

# Chapter 2: Docker Compose Architecture Explained

Understanding the architecture behind Docker Compose is essential for anyone who wants to build reliable, scalable, and maintainable multi-container applications. In the previous chapter, we introduced the concept of Docker Compose and explored why it exists. Now, we are going to peel back the layers and examine how Docker Compose actually works under the hood. This chapter will walk you through the internal architecture, the relationship between Compose and the Docker Engine, the role of the Compose file, the lifecycle of services, and the networking and storage models that make everything function seamlessly. By the end of this chapter, you will have a deep, professional-level understanding of the architectural decisions that drive Docker Compose and how each component fits together to orchestrate multi-container applications.

## The Big Picture: Where Docker Compose Fits in the Docker Ecosystem

Before diving into the specifics, it is important to understand where Docker Compose sits in the broader Docker ecosystem. Docker itself is a platform that provides the ability to package, distribute, and run applications inside containers. The Docker Engine is the core runtime that manages containers, images, networks, and volumes on a single host. Docker Compose, on the other hand, is a tool that sits on

top of the Docker Engine and provides a declarative way to define and manage multi-container applications.

Think of it this way: if the Docker Engine is the engine of a car, Docker Compose is the dashboard and steering wheel. You could operate the engine directly by issuing individual commands, but Compose gives you a unified interface to control everything from a single place. Docker Compose reads a YAML configuration file, interprets the desired state of your application, and then communicates with the Docker Engine API to create and manage the necessary containers, networks, and volumes.

The following table summarizes the relationship between key Docker components and Docker Compose:

<b>Component</b>	<b>Role</b>	<b>Relationship to Compose</b>
Docker Engine	Core container runtime that builds, runs, and manages containers	Compose communicates with the Engine API to execute all container operations
Docker CLI	Command-line interface for interacting with the Docker Engine	Compose extends the CLI experience by adding multi-container orchestration commands
Docker Images	Read-only templates used to create containers	Compose references images in the YAML file and can trigger builds using Dockerfiles
Docker Containers	Running instances of Docker images	Compose creates, starts, stops, and removes containers as defined in the service configuration
Docker Networks	Virtual networks that allow containers to communicate	Compose automatically creates and manages networks for inter-service communication

---

Docker Volumes	Persistent storage mechanisms for container data	Compose defines and manages volumes to ensure data persistence across container restarts
Docker Compose File	YAML file that declares the entire application stack	This is the central configuration artifact that Compose reads and interprets

---

Docker Compose does not replace the Docker Engine. It is not a separate runtime. Every action that Compose performs is ultimately translated into Docker Engine API calls. This is a critical architectural point because it means that anything Compose does, you could theoretically do manually with individual docker commands. Compose simply automates and orchestrates these commands based on your declarative configuration.

## The Docker Compose File: The Architectural Blueprint

At the heart of Docker Compose architecture is the Compose file, typically named `docker-compose.yml` or `compose.yaml`. This file serves as the single source of truth for your entire application stack. It is written in YAML format and describes every aspect of your application: which services to run, how they should be configured, what networks they should be connected to, and what volumes they should use for persistent storage.

The Compose file follows a well-defined schema that has evolved over multiple versions. The modern Compose specification (used by Docker Compose V2) no longer requires a `version` key at the top of the file, though you may still see it in older configurations. The file is organized into several top-level sections, each responsible for a different architectural concern.

Here is a comprehensive example that illustrates the structure:

```
services:
  web:
    build:
      context: ./web
      dockerfile: Dockerfile
    ports:
      - "8080:80"
    environment:
      - DATABASE_HOST=database
      - CACHE_HOST=cache
    depends_on:
      database:
        condition: service_healthy
      cache:
        condition: service_started
    networks:
      - frontend
      - backend
    volumes:
      - web-static:/var/www/static
  restart: unless-stopped
  deploy:
    resources:
      limits:
        cpus: "0.50"
        memory: 512M

  database:
    image: postgres:15
    environment:
      POSTGRES_DB: myapp
      POSTGRES_USER: admin
      POSTGRES_PASSWORD: secretpassword
    volumes:
      - db-data:/var/lib/postgresql/data
      - ./init-scripts:/docker-entrypoint-initdb.d
    networks:
      - backend
  healthcheck:
```

```

  test: ["CMD-SHELL", "pg_isready -U admin"]
  interval: 10s
  timeout: 5s
  retries: 5

  cache:
    image: redis:7-alpine
    command: redis-server --maxmemory 256mb --maxmemory-policy
    allkeys-lru
  networks:
    - backend
  volumes:
    - cache-data:/data

  networks:
    frontend:
      driver: bridge
    backend:
      driver: bridge
      internal: true

  volumes:
    db-data:
      driver: local
    cache-data:
      driver: local
    web-static:
      driver: local

```

Let us break down each top-level section and understand its architectural significance:

---

<b>Top-Level Key Purpose</b>	<b>Architectural Role</b>
services	Defines the containers that make up your application

---

---

networks	Defines custom networks for inter-service communication	Controls the network topology of your application. You can isolate services, create internal-only networks, and control which services can communicate with each other.
volumes	Defines named volumes for persistent data storage	Ensures that data survives container restarts and removals. Named volumes are managed by Docker and can be backed by different storage drivers.
configs	Defines configuration objects (used primarily in Swarm mode)	Allows you to manage non-sensitive configuration data outside of images.
secrets	Defines secret objects for sensitive data	Provides a secure mechanism for handling passwords, API keys, and certificates.

---

The Compose file is not just a convenience. It is an architectural artifact that captures the entire topology of your application. When you share this file with another developer or deploy it to a different environment, the entire application stack can be reproduced exactly as defined.

**Note:** The Compose file is parsed and validated before any operations are executed. If there is a syntax error or an invalid configuration, Compose will report the error and refuse to proceed. This fail-fast behavior is an important architectural safeguard that prevents partially deployed stacks.

## The Project Concept: Namespacing and Isolation

One of the most important architectural concepts in Docker Compose is the "project." A project is the logical grouping of all resources (containers, networks,