

# Patterns of Multi-Container Composition for Service Orchestration with Docker Compose

Kalvin Eng · Abram Hindle · Eleni Stroulia

Received: date / Accepted: date

**Abstract** Software design patterns present general code solutions to common software design problems. Modern software systems rely heavily on containers for running their constituent service components. Yet, despite the prevalence of ready-to-use Docker service images ready to participate in multi-container service compositions of applications, developers do not have much guidance on how to compose their own Docker service orchestrations. Thus in this work, we curate a dataset of successful projects that employ Docker Compose as an orchestration tool to run multiple service containers; then, we engage in qualitative and quantitative analysis of Docker Compose configurations. The collection of data and analysis enables the identification and naming of repeating multi-container composition patterns that are used in numerous successful open-source projects, much like software design patterns. These patterns highlight how software systems are orchestrated in the real-world and can give examples to anybody wishing to compose their own service orchestrations. These contributions also advance empirical research in software engineering patterns as evidence is provided about how Docker Compose is used.

**Keywords** Docker · Docker Compose · Containerization

## 1 Introduction

Following the successful adoption of software design patterns (Gamma et al. 1994), we analyze the patterns of successful self-deployed open-source projects employing Docker Compose as an orchestration tool to run multiple service containers. Software design patterns (Gamma et al. 1994) provide solutions to

---

Kalvin Eng · Abram Hindle · Eleni Stroulia  
University of Alberta  
Edmonton, Alberta, Canada  
E-mail: {kalvin.eng, abram.hindle, stroulia}@ualberta.ca

recurring problems in software development, often at the class level, sometimes intersecting with architecture.

Software systems have become more diverse over time as monolithic architectures and n-tier/n-layered projects transition to containerized deployments. now require developers to be familiar with many tools and services of a software system that are often rapidly changing (Carey 2021). This diversity motivates the need to identify common composition patterns of services that could be reused by projects to improve simplicity. Since many projects use off-the-shelf container images to orchestrate services with different configuration parameters, we can identify composition patterns, such as using a reverse proxy service to route to services hosted internally.

In this work, we curate a dataset of successful projects that employ Docker Compose as a service orchestration tool from the self-hosted GitHub community; then, we engage in qualitative and quantitative analysis of Docker Compose configurations. Successful projects are those that have captured the attention and interest of a wider community, thus being recognized as valuable and significant. The collection of data and analysis enables the identification and naming of repeated composition patterns orchestrated by numerous successful open-source projects, much like software design patterns.

Our work derives patterns from the open coding of Docker Compose multi-container compositions in projects of the self-hosted GitHub community. We identify three dimensions of Docker Compose usage from open coding: the service level, the orchestration level, and the repository level. The service level consists of the images and instructions used in each service of a Docker Compose file. The orchestration level refers to how the services interact among each other in a Docker Compose file when they are orchestrated together. The repository level examines how Docker Compose files are executed among other files in a repository. These dimensions categorize the observed patterns of Docker Compose usage. Patterns are considered to be commonly observed solutions that can address specific problems when using Docker Compose.

By observing these patterns, developers can use them as a point of reference when deciding how to compose their services for orchestration. These patterns are beneficial as Docker Compose’s versatility can make configurations unnecessarily complex (Podviaznikov 2017). Furthermore, previous research (Ibrahim et al. 2021) fails to capture the essence of usage for Docker Compose as data is analyzed at a macro level without considering the minute details of service orchestration usage, such as how developers can use Docker Compose in different contexts like development and testing.

Our work is driven by the desire to allow practitioners to see real-world examples of Docker Compose patterns that encompass multi-container composition. Therefore, this work contributes the following:

- An open-coded dataset of concepts in Docker Compose files;
- A set of Docker Compose patterns for software systems.

These contributions are helpful for anybody who uses Docker Compose to define their multi-container compositions for service orchestration and needs

examples of . Furthermore, these contributions advance empirical research in software engineering patterns based on how Docker Compose is used.

## 2 Background and Related Work

In this section, we outline what multi-container composition and service orchestration is, how service containers are composed and orchestrated in the Docker ecosystem, and introduce related work motivating why we wish to investigate patterns in Docker Compose files.

### 2.1 Multi-Container Composition and Service Orchestration

Our use of the words: *composition* and *orchestration* stems from their usage in service-oriented architectures (SoA) in the past. Composition refers to a “composite relationship between a collection of services” (Erl 2005). While orchestration refers to “a centrally controlled set of workflow logic facilitates interoperability between two or more different applications” (Erl 2005).

Using the definition of composition from SoA, we define *multi-container composition* as a collection of service containers that are associated with each other.

Orchestration has evolved from its SoA definition to be a term that is inferred as *container orchestration* at the time of this paper writing in . Container orchestration is a term that has been widely attributed to Kubernetes where multiple application containers are managed across different host machines (The Kubernetes Authors 2022). The common definition for “container orchestration” has appeared in survey literature as the deployment of containerized applications across computer clusters (Rodriguez and Buyya 2019; Truyen et al. 2019). The Docker alternative to Kubernetes is Docker Swarm, which uses the Docker Compose file format to define the services, networks, and volumes of the containers (Webb 2022).

Docker Compose falls short of being a complete container orchestration solution as it solely manages and deploys containers on a single host instead of a cluster of host machines (Webb 2022). However, Docker Compose encompasses the aspect of managing multiple containers and is considered to be an intermediate solution between running a single container and running multiple containers across multiple hosts (Webb 2022; Miell and Sayers 2019; Kane and Matthias 2023). Docker Compose has been referred to as container orchestration by Raj et al. (2015) and Miell and Sayers (2019). However, for the sake of clarity, we use in this paper the term *service orchestration* to describe Docker Compose which manages and deploys multiple service containers locally on a host machine.

## 2.2 The Docker Ecosystem

Docker is a suite of *Infrastructure as Code* (IaC) tools that can orchestrate containers for deployment on varying environments from a local machine to the cloud (Guerriero et al. 2019). To orchestrate containers on a local machine, one would use the Docker Compose tool. While one would use the Docker Swarm tool to orchestrate containers across multiple machines in the cloud.

Docker is designed to create *images* and run *instances* of these images called *containers* which are defined in a Dockerfile. A Dockerfile is a set of instructions that are defined to create a containerized environment for running an application. In order to orchestrate multiple containers at once, Docker Compose or Docker Swarm can be used to facilitate the orchestration through the use of configurations defined via a YAML based Docker Compose file (Ibrahim et al. 2021). The Docker Compose file describes a composition of services which can span multiple containers. An example can be seen in Figure 1.

```
1  services:
2    web:
3      image: ubuntu
4      ports:
5        - "80:8000"
6      volumes:
7        - web-logs:/var/log/web
8      deploy:
9        replicas: 2
10       placement:
11         constraints:
12           - "node.role==worker"
13
14     db:
15       image: postgres:13
16       hostname: db
17       deploy:
18         placement:
19           constraints:
20             - "node.role==manager"
```

Fig. 1: Example of a Docker Compose file

Figure 1 is an example of a simple Docker Compose file where the services “web”, that is an Ubuntu container, and “db”, that is a pre-built container image, are defined. At a minimum a Docker Compose file will always have at least one service defined with the option to give additional instructions such as mounting volumes, exposing ports, passing environment variables, and what to do in a swarm configuration. A swarm configuration allows for repli-

cas of services to be deployed to a cluster which also uses Docker Compose configurations. As of , a Docker Compose file has 194 different keywords for configuration available , making it highly configurable for different use cases.

In 2020, the Docker Compose file format was released as open source under the name “Compose Specification” (Docker Inc. 2020, 2022a). The “Compose Specification” is used by alternative Docker Compose service orchestration tools such as nerdctl (nerdctl Authors 2023), Okteto (Okteto, Inc. 2023), Podman (Podman Compose Authors 2023), and Rancher (Rancher Labs 2023). The specification has also been used to help translate Docker Compose files to Kubernetes deployments with Kompose (Kompose Authors 2022).

Using Docker Compose files, a developer is quickly able to set up multiple containers for use cases such as development and testing. This makes the use of Docker a versatile choice in a software stack. However, the wide swath of instructions available in Docker Compose makes deployment complex as developers who use Docker may not be aware of possible features or configurations to use or to avoid. The flexibility of instructions makes learning Docker Compose more challenging as there are many options to solve a problem. Therefore, this work seeks to identify patterns of configuration to better support the usage of Docker Compose in terms of simplicity, readability, and reproducibility.

### 2.3 Previous Docker Studies and Prior Patterns Research

Since we wish to investigate patterns in Docker Compose files of the Docker ecosystem, we review prior work on understanding the Docker ecosystem. Zhang et al. (2018, 2019) look at the evolution of Dockerfiles to understand container build time and how images change in projects. They find that different Dockerfile architecture attributes, such as the number of image layers, the size of each layer, and the diversity of instructions, can affect the quality of Dockerfiles with regard to linter (static code analysis) errors and build time. Furthermore, studies have thoroughly investigated Dockerfiles on public repository services and found that developers who build and maintain containers would benefit from utilizing tools for container creation (Eng and Hindle 2021; Lin et al. 2020; Cito et al. 2017). The empirical evidence observed in these studies supports the motivation for developing tools to improve the developer Docker experience, such as: DockerSlim <sup>1</sup> which helps to optimize container images, dive <sup>2</sup> which helps to analyze container images, Buildpacks <sup>3</sup> which helps to build images, and Binci <sup>4</sup> which helps containerize a development workflow. Despite the existence of tools aimed at enhancing the Docker experience, there still needs to be empirical research on patterns that can streamline the orches-

---

<sup>1</sup> <https://github.com/docker-slim/docker-slim>

<sup>2</sup> <https://github.com/wagoodman/dive>

<sup>3</sup> <https://buildpacks.io/>

<sup>4</sup> <https://github.com/binci/binci>

tration of various container contexts. This need motivates our exploration into Docker Compose configuration patterns.

Prior research has seen the potential in investigating service orchestration tools like Docker Compose. Xu and Marinov (2018) suggest that Docker Compose configurations can offer invaluable insights as multiple containers are “glued” together to form a service highlighting why Docker Compose is important to investigate. Ibrahim et al. (2021) investigate the usage of Docker Compose and find that Docker Compose files are infrequently updated, advance Docker Compose operations are infrequently used, and that applications rarely adopt new versions of Docker Compose. Furthermore, they also find that over a quarter of Docker Compose configurations use a single container. This prior research motivates our work as it suggests that there is an interest into how to develop software systems with Docker Compose.

Docker Compose can also be classified under *Infrastructure as Code* (IaC), motivating the review of studies on IaC. In terms of related work for patterns in IaC, the most closely related work is by Rahman et al. (2019) where they qualitatively identify security smells in 1,726 Puppet scripts using descriptive coding. From the descriptive coding, they describe and identify with qualitative analysis seven security smells that developers should be aware about. In a similar vein, Ksontini et al. (2021) identify the refactoring types in Dockerfiles and Docker Compose configurations using manual analysis to suggest best practices in Docker. Guerriero et al. (2019) perform 44 semi-structured interviews with the findings that developers should find patterns for IaC that can be used across different IaC tools. This previous research suggest that a manual analysis of Docker Compose files could offer invaluable insights.

Patterns have also been discovered by analyzing the commits and developers of repositories. Rahman et al. (2020a) look at Puppet file commits in repositories qualitatively to determine a taxonomy of defects. Furthermore, Rahman et al. (2020b) also look at developers and their development activities and how some activities can be anti-patterns towards developing defect-free Puppet code. The discovery of patterns through analyses of development activity suggests that it can help promote more robust codebases.

Patterns on other platforms have been investigated. Shamim et al. (2020) perform a grey literature review of the best security practices for Kubernetes configurations. Burns and Oppenheimer (2016) present cloud patterns observed in distributed systems using containers from anecdotal experiences. However, Burns and Oppenheimer (2016) do not provide any empirical evidence of the cloud patterns used. The presentation of these patterns suggests that there is interest in understanding what patterns other developers are using.

From the Pattern Languages of Programs workshops, numerous general cloud patterns have been derived by the same group of authors: Sousa et al. (2018a,b,c, 2017, 2015). These patterns are further investigated by the authors by performing a survey on 100 practitioners finding that practitioners use at least one of the patterns (Sousa et al. 2021). Sousa et al. (2021) also find that 67% of practitioners use containerization and 69% of practitioners

use container orchestration, and that the “ease of use and security [are] main drivers for adoption” to the cloud which motivates the need for more patterns. Notably, none of the patterns by Sousa et al. (2018a,b,c, 2017, 2015), to the best of our knowledge, have been presented in literature with empirical observations in real-world code.

Prior works on the Docker ecosystem have mainly focused on Dockerfile usage (Zhang et al. 2018, 2019; Eng and Hindle 2021; Lin et al. 2020; Cito et al. 2017). Research into patterns of IaC and development activities have not specifically addressed Docker Compose (Rahman et al. 2019; Ksontini et al. 2021; Rahman et al. 2020a,b). Therefore, our work focuses more on deriving the orchestration patterns of containers using Docker Compose similar to previous research for patterns in Infrastructure as Code. We elaborate on how the patterns we discover relate to the prior IaC works in Section 5.4.

### 3 Methodology

Our methodology consists of three steps, separated into the subsections below. First, we select projects that use Docker Compose in the self-hosted GitHub community (nodiscc et al. 2022) and extract 527 files for analysis from 218 projects. Next, we use open coding to identify categories of Docker Compose file characteristics at a service level, orchestration level, and repository level. Finally, we describe patterns based on the categories extracted from open coding. The results of our analysis can be replicated using our replication package (Eng et al. 2024).

#### 3.1 Project Selection

The data curated for this work was obtained by extracting GitHub repository URLs from the README of the *Awesome-Selfhosted* (nodiscc et al. 2022) GitHub repository. The software projects of the self-hosted community often replace *Software as a service* (SaaS) provider services with a diverse set of self-hosted solutions ranging from analytics software to wiki software.

We choose the list of self-hosted repositories to ensure we are studying projects that people use and deploy as the list is manually curated by the self-hosted community. The intended purpose of a repository — which surrounds many of the perils of mining GitHub (Kalliamvakou et al. 2014) — might introduce noise into an analysis as repositories may not necessarily be used for the purposes of maintaining and writing code. A representative sample of typical Docker Compose usage in software systems can be shown by selecting a list curated by the self-hosted community, which frequently uses Docker Compose to launch software systems. This is beneficial because we are looking for deployable and real examples of patterns in Docker Compose files.

Our project selection process visualized in Figure 2. Initially, 848 GitHub projects were extracted from the README pushed in the “1726a40” commit.

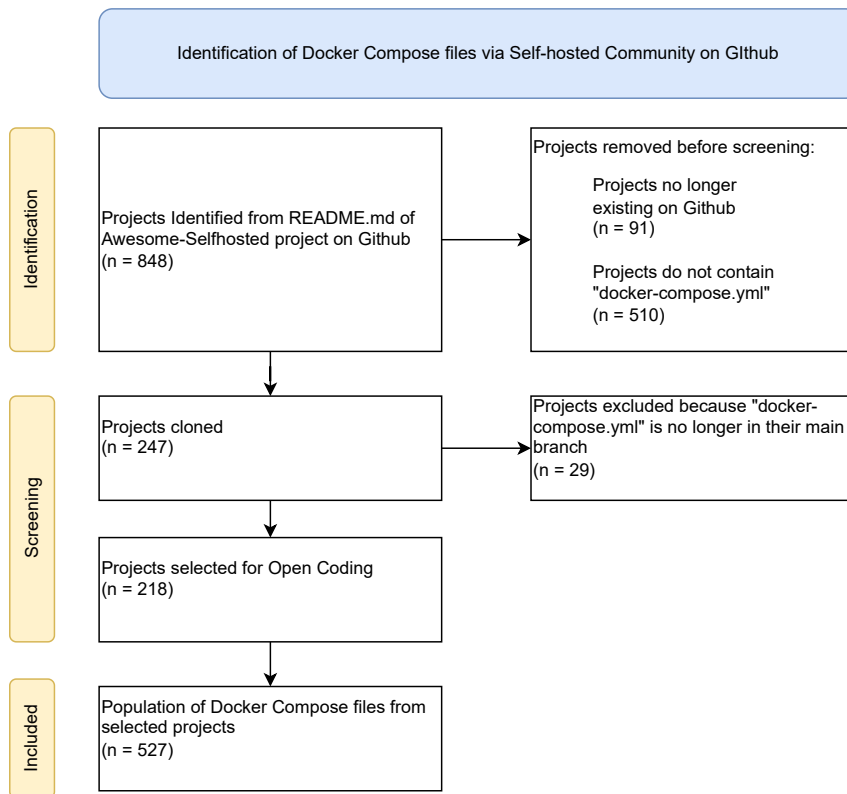


Fig. 2: Project Selection Process

To determine if projects contain at least 1 Docker Compose file, we clone each of the repositories using the mirror option and the git history was queried across all refs and the HEAD for filenames containing “docker-compose” (case-insensitive). From cloning, we obtain the actual commits and blob contents relating to the repositories for each of the repositories. The metadata was extracted using the PyDriller library (Spadini et al. 2018). We were able to obtain 247 successfully cloned projects as of May 9, 2022.

Using the 247 cloned repositories, we find the latest main branch commit that has modified a file containing the name “docker-compose”. We find that 29 projects have removed file(s) containing the name “docker-compose” in their latest main branch commit. This indicates that Docker Compose may have been deprecated or moved elsewhere. Furthermore, we find that 99 projects have multiple files with the name “docker-compose”. We remove the projects that no longer have “docker-compose” in their latest main branch commit resulting in . To gain a better understanding of the diversity in the 218 projects, we outline in Table 1 their respective categories as they appear



Table 1: Project Categories and Counts. Categories that have a count of 1 or 2 have been merged into Other <sup>5,6</sup>.

Category	Count
Communication - Social Networks and Forums	12
Communication - Custom Communication Systems	11
Automation	10
E-commerce	9
Bookmarks and Link Sharing	8
Feed Readers	8
Money, Budgeting & Management	8
Content Management Systems (CMS)	7
Miscellaneous	7
Media Streaming - Video Streaming	6
Task Management & To-do Lists	6
Communication - Email - Complete Solutions	5
File Transfer - Single-click & Drag-n-drop Upload	5
Note-taking & Editors	5
Pastebins	5
Personal Dashboards	5
Software Development - IDE & Tools	5
Software Development - Project Management	5
Archiving and Digital Preservation (DP)	4
Password Managers	4
Photo and Video Galleries	4
Software Development - API Management	4
Wikis	4
Conference Management	3
File Transfer - Web-based File Managers	3
Internet of Things (IoT)	3
Learning and Courses	3
Media Streaming - Audio Streaming	3
Polls and Events	3
Proxy	3
Ticketing	3
URL Shorteners	3
Other	28

in the *Awesome-Selfhosted* (nodiscc et al. 2022) GitHub repository. In Table 1 we combine categories with a count of 1 and a count of 2 into “Other” <sup>5,6</sup>.

<sup>5</sup> Projects with count of 1: *Booking and Scheduling, Office Suites, Search Engines, Resource Planning, Communication - SIP, Human Resources Management (HRM), Maps and Global Positioning System (GPS), Document Management - Integrated Library Systems (ILS), Calendar & Contacts - CalDAV or CardDAV Servers, Read-it-later Lists*

<sup>6</sup> Projects with count of 2: *Resource Planning - Enterprise Resource Planning, Groupware, Software Development - FaaS & Serverless, Gateways and Terminal Sharing, File Transfer - Peer-to-peer Filesharing, File Transfer - Object Storage & File Servers, Document Management, Community-Supported Agriculture (CSA), Document Management - E-books, Document Management - Institutional Repository and Digital Library Software, Communication - Email - Mailing Lists and Newsletters, Communication - Email - Webmail Clients, Software Development - Localization, Self-hosting Solutions, Games, Communication - XMPP - Servers, File Transfer & Synchronization*

From the 218 projects, we extract 527 Docker Compose-like files. The files are “Docker Compose-like” because they might resemble a Docker Compose file, i.e., the file is YAML and is named “docker-compose.yml” or co-exists with a file named “docker-compose.yml”. 404 Docker Compose-like files are parts of projects that have more than one Docker Compose-like file, while 123 Docker Compose-like files belong to projects that have only a single Docker Compose-like file. We use these files for open coding as described in Section 3.2 below.

All the projects were cloned with considerations of the perils when curating git repositories highlighted by Bird et al. (2009) such as: being aware of commits across branches; tracing changes that occur not only in the main-line branch; determining the branches of commits; and tracing the sources of merges. Therefore, in order to gain as much information as possible, all projects were cloned using the mirror option which is a total copy of the repository including all refs to track as many changes as possible.

We also try to verify if these self-hosted projects represent products that developers might use by measuring the GitHub stars of each project. The measure of GitHub stars has been used in previous studies to measure the popularity of a project (Brisson et al. 2020; Blincoe et al. 2016; Borges et al. 2016) — a project is assumed to be more popular if it has more stars, and as a result, by association, it should be used and important.

To support why we should care about the selected self-hosted projects that we analyze, we count the GitHub stars — the proxy for popularity — of each project as of December 20, 2022, as seen in the boxplot of Figure 3a. We compare the count of the stars in the self-hosted projects with the distribution of GitHub stars of Docker Compose projects selected by Ibrahim et al. (2021) from GitHub data in Google BigQuery seen in Figure 3b. .

We use the Mann-Whitney U test, which indicates that the stars in the self-hosted projects are greater than the stars of projects selected by Ibrahim et al. (2021) with  $p = 1.08e-94 < 0.05$ . To ensure that the difference between the distributions are not only statistically significant but also practically significant, we get a large effect size from Cliff’s Delta being  $0.82 > 0.474$ , which indicates that the greater stars in the selected self-hosted projects are practically significant. The greater number of stars demonstrates why the selected self-hosted projects are considered to be successful.

Overall, the star distributions are quite different as seen in Figure 3. The self-hosted projects’ distribution has larger values, while Ibrahim et al.’s (2021) distribution has smaller values suggesting that the sample of self-hosted projects With a median of 1,512 and 25% of the projects falling between 14 stars and 382 stars and 25% having at least 4,674 stars, the self-hosted projects’ have more stars overall. In comparison, Ibrahim et al.’s (2021) distribution has a smaller number of stars overall, with a median of 9 and 25% of the projects falling between 0 and 1 stars and 25% of projects having at least 86 stars. Therefore, the higher number of stars in our chosen self-hosted projects is indicative of projects being used in practical applications compared to the prior Docker Compose usage study (Ibrahim et al. 2021).

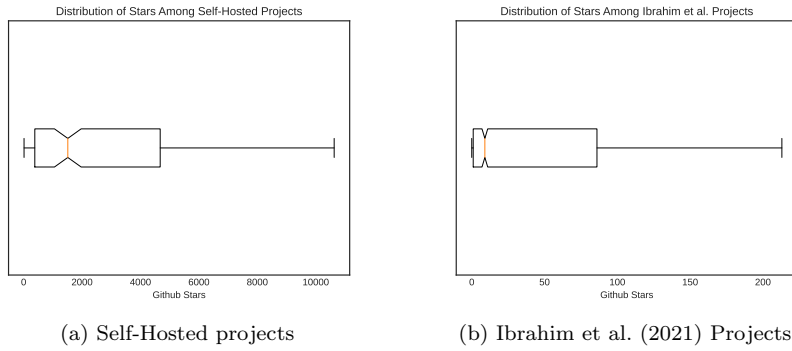


Fig. 3: Distribution of Stars




Fig. 4: Open Coding Process Outline Inspired by Seidel (1998)

Notably, in the previous Docker Compose study (Ibrahim et al. 2021)

### 3.2 Open Coding

Open coding is the process of finding concepts through manual inspection to gain deeper insights within data. Seidel (1998) describes it as a qualitative data analysis process where we *notice*, *collect* and *think* about interesting things. The process is non-linear as observations can be re-iterated through many times. Qualitative methods are useful for coming up with observations that are difficult to derive with quantitative analysis. This process can be visualized in Figure 4.

With the collected Docker Compose files, a rater analyzes the files for what they consider to be “interesting” observations and “important” observations. “Interesting” observations catch the attention of the rater and may appear infrequently while “important” observations emerge from the repeated analysis of files from noticing common characteristics among the different files. Interesting observations within the files are *noticed* and *collected*. Any important observations are also *noticed* and *collected* separately with a label.

Seidel (1998) says that labels of important observations can be *thought* about and abstracted into “in vivo codes” which are observations directly from the data or “constructed codes” where observations are created by the

annotator. From these codes a combination of *objectivist* and *heuristic* codes can emerge. *Objectivist* codes occur when observations require some sort of interpretation as opposed to *heuristic* codes which are rote observations. Codes can also be a mixture of *objectivist* and *heuristic* codes (Seidel 1998).

From the codes and interesting thoughts, discoveries of “categories” and “patterns” can emerge by *thinking* about relations, similarities, and dissimilarities among the codes in relation to the interesting thoughts. This process is iterative and repeated until the annotator deems the categories and patterns satisfactory.

We can define our open coding process in different levels of detail. We examine each Docker Compose file as part of our ad hoc process on three different levels: (1) what and how services are defined, (2) how the services holistically interact when orchestrated together, and (3) how the files function within a repository context. From the coding, we have several codes that emerge and are used to derive patterns as described in Section 3.3.

For each of the three levels, there are guiding questions that emerge for our open coding such as:

1. **Service** - How are the services defined and what are the services?
2. **Orchestration** - How are the services similar and different among each other and how do they interact with each other?
3. **Repository** - How does Docker Compose work in the repository with files and folders?

The codes from these guiding questions help to develop the patterns that we present in this paper. From 218 GitHub projects, we obtain 527 Docker Compose-like files for open coding resulting in 80 codes with: 40 codes at the service level, 35 codes at an orchestration level, and 5 codes at a repository level.

### 3.2.1 Service Level

At a service level, we inspect each service configuration of a Docker Compose file and create new service type codes when an existing service purpose cannot be categorized under existing codes. To identify service types, we consider the image of a service.

If the image is pulled from a registry, we look at the README in the container registry and read the README to determine if there is a purpose for the image that can be used to identify the type of a service. If the purpose cannot be identified, then the repository containing the Docker Compose file is inspected for READMEs to indicate which services may be present in the software system and can help identify the service type. If the service type cannot still be determined, then the string of the image (e.g. “webapp:tag” of the line “image: webapp:tag” in a Docker Compose file) is searched in Google to determine a purpose and identify a service type.

If the image is created from a Dockerfile rather than being pulled, we attempt to determine the purpose of the service by examining the Dockerfile

and README of the base image if it is present in the container registry. We inspect the instructions following the base image to determine the service container image. For example, a base image could be an operating system and the following instructions in the Dockerfile are installing software packages using the base image. Then the purpose of the service can be determined based on the software packages installed and used to categorize the service under a service type. Failing to identify a purpose, then the repository containing the Docker Compose file is inspected for READMEs to indicate which services may be present in the software system and can help identify the service type. If the service type can still not be determined, then the string of the image (e.g. “webapp:tag” of the line “FROM image: webapp:tag”) is searched in Google to determine a purpose and identify a service type.

From the inspection of the service level, codes of service types emerge which are used to identify the purpose of a service in a Docker Compose file. We identify 40 different codes in total with 33 types with a count more than 1 as seen in Table 2 and organize the types into 7 higher categories (Core, Networking, Reliability, Setup, Data, Scheduling, and Miscellaneous).

*Core services* refer to services that contain the central logic for a software system. In this category, there are 2 services: the *Frontend service* which are applications which connect with a database or backend; and the *Backend service* which are APIs of applications that connect with a database.

While *networking services* refer to services that help network other services. This includes: the *Reverse proxy service* which routes web traffic to different servers with examples of services being as **Nginx** and **Traefik**; the *Discovery service* which helps connect microservices together; the *DNS service* which runs a DNS server; and the *HTTP accelerator service* refers to services that can speed up HTTP requests with a cache such as **Varnish**.

*Reliability services* help maintain other services. There are numerous services in this category including: the *Testing service* that aids in testing using tools such as **Selenium** which automates tests; the *Container management service* which polls image registries for updates and automatically update the container image; the *Tracing service* which refers to services that help manage logs of applications; and the *Event monitoring service* which monitors and reports the performance of other services.

*Setup services* are services that help initialize other services such as: the *Database init service* which helps initialize a database by running scripts for a database such as loading data; and the *Setup service* which helps initialize or set up an application.

Moreover, *Data services* refer to services that manage the data of a software system. There are 5 services in this category including: the *Database service* which are services that systematically store data in persistent storage such as **Postgres** and **Mongo**; the *Caching service* that systematically stores data in memory and may not be persistent; the *Database administration service* which acts as a GUI interface to databases; the *Object storage service* that stores unstructured data organized by files and folders; and the *Data streaming service*

which refers to services that allow data to be delivered at high throughput such as **kafka**.

*Scheduling services* are used to automatically perform tasks which include: the *Job scheduling service* that schedules tasks to do on a server; the *Cron service* that schedules tasks that are run on the bash command line using **cron**; and the *Workflow service* that can help with automation of services.

Finally, *Miscellaneous services* encompass a diverse range of offerings that do not fit into the above categories, including but not limited to: the *Mail service* which are mail servers that receive emails via protocols like SMTP; the *Search service* utilizing software such as **Solr** to build large indexes for improved search speed; the *Identity service* managing multiple identities across services with software such as **KeyCloak**; the *Visualization service* for data visualization with tools like **Kibana**; the *Certificate service* for generating SSL certificates; the *Zippping service* for compressing files into a zipped format; the *Message broker service* for passing messages from one service to another; the *Image recognition service* providing image recognition through API; the *Chat service* with platforms such as **Mattermost**; the *Office service* offering tools for text, spreadsheets, and presentations; the *Linux utilities service* running Linux utilities like **busybox**; the *Secrets service* for storing application secrets; and the *Hello world service*, demonstrating the proper functioning of Docker.

### 3.2.2 Orchestration Level

When looking at Docker Compose files with multiple services, it is essential to consider individual services in isolation while also examining the relationships between them. We can gain insights into how services interact when deployed during orchestration by taking a holistic view of services within a Docker Compose file. To perform this examination, we develop some guiding questions, which includes:

- Do any of the services share the same Docker Compose options? (e.g. using the same ENV file)
- How do the services share the same Docker Compose options? (e.g. using YAML aliases)
- Do any of the services use the “links” or “depends.on” instructions to connect each other?
- Does one service expose a port and another service connects to that port with an environment variable?
- Does a service extend from another service? Does a service extend from another file?

It should be noted that the “links” we refer to is the one used in Docker Compose files and is different from the legacy container links used by the Docker Engine. To acquire more insights into the service relationships within a Docker Compose file, we iteratively refine these questions depending on interesting discoveries that emerge during our analysis.

Table 2: Counts of Service Type Codes

Service Type Code	Count
Database service	365
Frontend service	343
Backend service	276
Caching service	111
Testing service	98
Reverse proxy service	92
Mail service	62
Search service	37
Database administration service	26
Object storage service	23
Identity service	17
Visualization service	17
Job scheduling service	15
Container management service	14
Event monitoring service	14
Certificate service	11
Database init service	11
Zippping service	11
Message broker service	8
DNS service	8
Tracing service	8
Image recognition service	6
Cron service	5
Chat service	4
Office service	4
Setup service	3
Workflow service	3
Linux utilities service	3
Secrets service	2
HTTP accelerator service	2
Hello world service	2
Discovery service	2
Data streaming service	2

Based on the holistic view of services, codes of service relationships emerge and can form patterns of Docker Compose options that can be reused. We identify 35 codes at the orchestration level that can be seen in Table 3. The table explicitly lists services instead of defining connections at a meta-level as we wish to determine empirically how different kinds of services are commonly configured and related.

### 3.2.3 Repository Level

In a repository context, we look at how Docker Compose files coexist among other files and folders in the repository. In addition to the Docker Compose files themselves, we carefully examine other relevant files to gain a more complete understanding of how they are used in practice. These files include READMEs, which often provide valuable information on the intended usage and context of a Docker Compose file. Moreover, we also analyze the configurations associated

Table 3: Counts of Codes at a Orchestration Level

Code	Count
Frontend service connects to database service with environment vars	28
Duplicate image reuse with different commands	17
Frontend environment variables map to database service environment variables	17
Uses yaml aliases	15
Service labels are configuration for reverse proxy	15
Backend service connects to database service with env variable	12
Uses <i>extends</i> for additional instructions	11
Frontend service connects to caching service with environment var	11
Reverse proxy service shares volume with frontend service	8
Backend service connects to caching service with env variable	8
Frontend shares volumes with backend service	4
Certificate service shares volume with reverse proxy service	3
Backend services shares the same env file with the database service	3
Database management service connects to database service with env variables	3
Duplicate image reuse	2
Frontend connects to mail service with env var	2
Frontend service shares variable with caching service	2
Frontend service connects to backend service with environment var	2
Cluster environment variables	1
Duplicate image reuse with different env vars	1
Frontend services shares the same env file with the caching service	1
Frontend services shares the same env file with the database service	1
Reverse proxy service shares volume with backend service	1
Container management service shares volume with frontend service	1
Frontend service shares variable with secrets service	1
Frontend service connects to message broker service with env variable	1
Frontend service connects to object storage service with env variable	1
Backend service connects to object storage service with env variable	1
Backend service connects to search service with env variable	1
Reverse proxy maps to frontend service with env variable	1
Testing service connects to database service via env	1
Frontend service connects to search service with environment var	1
Docker swarm instructions	1
Backend service shows relationship to identity service with env var	1
Job scheduling service connects to database service with env war	1

with these files, such as ENV files which can indicate that a Docker Compose file is used for development when the debug flag is set. Execution details of Docker Compose can also be found by looking at Makefiles and bash scripts. In addition to looking at the contents of files, the filenames and file paths are also inspected, i.e., we look for keywords such as “development” and “testing” to indicate a Docker Compose file purpose.

The repository context helps identify the codes of use cases for Docker Compose files and codes of how Docker Compose files are executed. At a repository context level, we identify 5 codes as seen in Table 4.

The codes in Table 4 can be described as follows:

- “Override file” refers to Docker Compose files that contain service configurations that can override other configurations when executed with another Docker Compose file.



Table 4: Counts of Codes at a Repository Level

Code	Count
Override file	53
Auto-generated file	29
Not an actual docker compose file, configuration	21
Not an actual docker compose file, template for generating	2
Development use case from comment	2

- “Auto-generated file” are Docker Compose files that have been automatically generated using a tool.
- “Not an actual docker compose file, configuration” means that files appear to be Docker Compose-like but are actually meant to configure services.
- “Not an actual docker compose file, template for generating” are files that are a template for generating Docker Compose files.
- “Development use case from comment” refers to the use case of a Docker Compose file that is determined to be for the development environment and testing based on comments in the Docker Compose file.
- “Cluster environment variables” are environment variables that define how to connect to other instances of a service.

### 3.3 Data for Pattern Discovery

Through the open coding, we identify patterns that are common or interesting to developers and report ones that occur more than once in Section 4. These patterns stem from interesting observations that we note down and describe in Section 3.3.1. We also quantitatively identify patterns by doing the following among the levels of open coding:

- Frequent itemset mining of the service type codes as described in Section 3.3.2.
- Considering the counts of orchestration level codes as described in Section 3.3.3.

#### 3.3.1 Interesting Observations

Observations are considered “interesting” when certain attributes of a Docker Compose file are considered by the reviewer to be unusual among all other Docker Compose files. For our interesting observations we note the following:

- The `mailu`<sup>7</sup> project generates Docker Compose from templates (they could have possibly used overrides) and there’s other kinds of Docker Compose files in its repo;

<sup>7</sup> <https://github.com/Mailu/Mailu/tree/e8641245>

- Many projects (e.g. `httpplaceholder`<sup>8</sup>) host multiple example Docker Compose files to demonstrate how to run their application with different service stacks;
- Projects using override files are usually highly configurable and are compatible with different service stacks;
- The `lardbit/nefarious`<sup>9</sup> project could have used YAML aliases instead of extends;
- There is an override file to change from pulling an image to building an image in the `mailcow` project;
- YAML aliases are used to repeat services with the same images but different commands;
- The `piqueserver`<sup>10</sup> project “reverses” override usage where an image is defined in the override file, and the base template is empty;
- Docker Compose has services that do not pertain to project stack directly (e.g. Using `letsencrypt pebble`);
- In a Docker Compose file, sometimes the “extends” configuration filename does not contain the string “docker-compose” when it typically does — this can be seen in the `rero-ils`<sup>11</sup> project;
- Docker Compose files for each microservice are automatically generated in the `overleaf`<sup>12</sup> project.

### 3.3.2 Frequent Itemset Mining

Frequent itemset mining of the service types helps find sets of services that frequently co-occur together. We only consider service type codes as they are easily understood to be defined together as opposed to the other orchestration and repository codes that may not intuitively co-occur. These frequent itemsets can signify a pattern of interest. To filter the amount of itemsets of interest, we consider itemsets that have support of at least 0.05 meaning that at least five percent of the Docker Compose files contain these services and are considered to be frequent itemsets. The results of our mining can be seen in Table 5.

In Table 5, we can see that there is high support for services connecting to a database service which are highlighted in bold. This motivates the patterns in Section 4.11 and Section 4.12. Also in Table 5, we can note that reverse proxy services are also very common which motivates the “HTTP Reverse Proxy Service” pattern in Section 4.13. Noting that there are some very frequent items within itemsets is a good indicator as to what patterns can be derived.

<sup>8</sup> <https://github.com/dukeofharen/httpplaceholder/tree/74d0b34/docker>

<sup>9</sup> <https://github.com/lardbit/nefarious/tree/a6f49a1>

<sup>10</sup> <https://github.com/piqueserver/piqueserver/tree/b30dab7d>

<sup>11</sup> <https://github.com/rero/rero-ils/tree/9fcf2e9e>

<sup>12</sup> <https://github.com/overleaf/overleaf/tree/ed66b43/services>

Table 5: Frequent Itemsets of Service Types

Support	Itemsets
0.051233	(testing service, database service)
0.053131	(caching service, reverse proxy service)
0.055028	(backend service, caching service, database service)
0.056926	(reverse proxy service, frontend service, database service)
0.060721	(mail service)
0.062619	(backend service, reverse proxy service, database service)
0.064516	(search service)
0.072106	(backend service, caching service)
0.075901	(backend service, frontend service, database service)
0.079696	(caching service, frontend service, database service)
0.085389	(backend service, reverse proxy service)
0.092979	(testing service)
0.096774	(reverse proxy service, frontend service)
0.102467	(reverse proxy service, database service)
0.121442	(backend service, frontend service)
0.125237	(caching service, database service)
0.129032	(caching service, frontend service)
0.161290	(reverse proxy service)
0.178368	(backend service, database service)
0.197343	(caching service)
0.278937	(backend service)
0.309298	(frontend service, database service)
0.529412	(database service)
0.571157	(frontend service)

### 3.3.3 Observation Counts

At the orchestration level and repository level we use the counts of combinations of observations as seen in Table 3 and Table 4. These counts help corroborate some of our interesting observations in Section 3.3.1.

An example where we used counts to motivate a pattern is the observation “Auto-generated file” found in Table 4 with a count of 29 for identifying the “Automatic Docker Compose File Generation” pattern in Section 4.1.

## 4 Patterns

In this section, we present 14 interesting patterns of Docker Compose usage discovered from interesting observations, frequent itemset mining, and observation counts as described in Section 3.3. Although a pattern may come across as common to some readers, there has been no empirical evidence or prior work to support these claims. More details of how the presented patterns were identified and its supporting evidence is described in this section using

- **Motivation.** Why developers might use this pattern.
- **Context/Applicability.** The situations where using the pattern might be appropriate.

- **Description.** An explanation of how the pattern works.
- **Advantages.** Description of the benefits of using this pattern compared to other patterns.
- **Disadvantages.** Description of the possible negative impacts of using this pattern.
- **Potential Issues.** Issues to be aware of when using this pattern.
- **Real-world Examples.** Examples identified from the Docker Compose configuration files within self-hosted software projects.
- **Supporting Evidence.** How the pattern and counts of the real-world examples were discovered.

Along with the template, we also include a visualization of the pattern to better help understand how the pattern works at a high-level.

#### 4.1 Automatic Docker Compose File Generation




Fig. 5: Pattern of Automatic Docker Compose File Generation.

- **Motivation.** Software systems can be orchestrated in many different ways with different configurations and services. Due to the complexity of creating Docker Compose files with different configurations and services, tools have been created to automatically generate Docker Compose files as opposed to manually creating and editing Docker Compose files.
- **Context/Applicability.** Developers of software systems who wish to simplify their deployments for their software systems may wish to develop automatic Docker Compose generation tools. This leaves less ambiguity as to how their software system should be run and configured.
- **Description.** Docker Compose files are automatically generated using a tool that allows for the setting of the application specific variables and services that are needed to run the software system. The Docker Compose files generated can be run using one command without any other further configuration.
- **Advantages.** As a user of the software system, the setup of the software systems is simplified as you can create the orchestration files with a guided tool. Docker Compose files are also consistent because of the automatic generation.

- **Potential Issues.** To recreate Docker Compose files, you will have to reuse the tool again to avoid any misconfigurations. Hand edits could lead to errors in deployment if modified incorrectly. The tool for automatic generation requires maintenance which can increase the dependencies and technical debt of a project.
- **Real-world Examples.** We found 29 Docker Compose files evidencing automatic Docker Compose file generation from web tools and build scripts. For example, the mail server project `Mailu` has created a web tool<sup>13</sup> to generate Docker Compose files that can be run using Docker Compose or in a Docker Swarm. The `Overleaf` project<sup>14</sup> also automatically generates continuous integration, development, and testing configurations for each of their microservices using build scripts.
- **Supporting Evidence.** These examples were found by counting the observations at a repository context level as seen in Table 4 denoted as “Auto-generated file”. It is motivated by the interesting observation, “Docker Compose files for each microservice are automatically generated in the `overleaf` project”.

#### 4.2 YAML Anchor and Alias

```
1  services:
2    first:
3      image: my-image:latest
4      environment: &env
5        - CONFIG_KEY
6        - EXAMPLE_KEY
7        - DEMO_VAR
8    second:
9      image: another-image:latest
10     environment: *env
```

Fig. 6: Example of how the YAML Anchor and Alias pattern is implemented from Docker Inc. (2023e).

- **Motivation.** The use of YAML anchors and aliases allow for the option of defining common sections of code in YAML files once so that it can be used repeatedly. This allows for Docker Compose file configurations to be neater and more efficient (Docker Inc. 2023e).

<sup>13</sup> <https://web.archive.org/web/20220915164516/https://setup.mailu.io/1.9/>

<sup>14</sup> <https://github.com/overleaf/overleaf/tree/ed66b43/services>

- **Context/Applicability.** When services are re-used multiple times with similar configurations in the same Docker Compose file, YAML anchors and aliases can be used.
- **Description.** Docker Compose uses a YAML parser at its core, and therefore can use YAML and anchors aliases as defined in the YAML spec. An anchor “&” can be denoted for a section in a YAML file, where it is placed after a colon with a keyword as such “&keyword”. The section can then be repeated with an alias “\*” by specifying after a colon “\*keyword”. The repeated section can also be overwritten with “<<” by specifying “<<:\*keyword” under a section.
- **Advantages.** The use of YAML aliases prevents duplicate configuration code.
- **Potential Issues.** The use of YAML aliases may appear non-intuitive at first due to the syntax. Furthermore, although it simplifies the Docker Compose file by less repetition, readability may still become an issue when there are many anchors and aliases to trace through.
- **Real-world Examples.** We found 15 Docker Compose files using this pattern. For example, the `chaskiq` project<sup>15</sup> uses YAML anchors and aliases with overrides to configure their Docker Compose file. The `canvas-lms` project<sup>16</sup> uses YAML anchors and aliases with an override to allow their main application to be run as a job scheduler. The `minio` project also uses this pattern to repeat configurations<sup>17</sup>.
- **Supporting Evidence.** The examples were found by counting the observations among services as seen in Table 3 under “Uses yaml aliases”.
- **Category.**
- **Related Patterns.**

#### 4.3 Docker Compose Service Inheritance




Fig. 7: Example of how the Docker Compose Service Inheritance pattern is implemented from Docker Inc. (2022b).

<sup>15</sup> <https://github.com/chaskiq/chaskiq/blob/8fb33b/docker-compose.yml>

<sup>16</sup> <https://github.com/instructure/canvas-lms/blob/d4ee7b/docker-compose.yml>

<sup>17</sup> <https://github.com/minio/minio/blob/95d1a1/docs/orchestration/docker-compose/docker-compose.yaml>

- **Motivation.** Docker Compose service configurations are often similar with slight modifications. Therefore, the inheritance of configurations from other similar services can help simplify configurations by repeating services that are similar. As a feature, Docker Compose offers the “extends” option which allows the inheritance of service configurations from other Docker Compose files at a service level or from services already within a Docker Compose file.
- **Context/Applicability.** When services are re-used multiple times with similar configurations, the “extends” option can be used to import service configurations from another Docker Compose file or existing services already in the Docker Compose file. This can allow for specialization of services within a Docker Compose file.
- **Description.** Under a defined service in a Docker Compose file, a service configuration can be imported by specifying the Docker Compose file path and the service that is to be extended. A service can also be extended by specifying a service name of the current file. The imported configuration can be overridden by re-specifying Docker Compose options.
- **Advantages.** Prevents duplication of code.
- **Potential Issues.** Readability may become an issue as developers have to trace through Docker Compose files to find the imported configurations. Using extends on an existing service is not ideal when you do not want the base service to be orchestrated, to avoid orchestrating the base service the command line Docker Compose orchestration command needs to change.
- **Real-world Examples.** This pattern was discovered in 11 Docker Compose files. The `hawkpost` project<sup>18</sup> extends their services from a common configuration file and overwrites the common service configuration with service specific options. The `nefarious` project extends a service from within the same file<sup>19</sup> and from an external file<sup>20</sup>. The `mattermost` project<sup>21</sup> also extends from an external file.
- **Supporting Evidence.** This pattern is found by counting the observation “Uses *extends* for additional instructions” as seen in Table 4 among the repository context.
- **Related Patterns.**

#### 4.4 Docker Compose Override Use Case

- **Motivation.** A service may have different use cases whereby only small parts of a service configuration need to be changed.

---

<sup>18</sup> <https://github.com/whitesmith/hawkpost/blob/5d4f017/docker-compose.yml>

<sup>19</sup> <https://github.com/lardbit/nefarious/blob/e8aa423/docker-compose.base.yml>

<sup>20</sup> <https://github.com/lardbit/nefarious/blob/e8aa423/docker-compose.transmission-vpn.yml>

<sup>21</sup> <https://github.com/mattermost/mattermost-server/blob/38d0c2/docker-compose.yml>




Fig. 8: Example of how docker compose can override files in the command line.

- **Context/Applicability.** When you need to run services in another mode such as testing and the Docker Compose file is designed for production deployment.
- **Description.** A different override Docker Compose file is created with the same services and configurations that they wish to override. This override Docker Compose file is called during Docker Compose orchestration in the command line with the “-f” flag along with the main Docker Compose file. The result is that containers can be started up with overridden configurations of the main Docker Compose file.
- **Advantages.** It separates use case scenarios of how Docker Compose is used.
- **Potential Issues.** The command to run override configurations is not evident in the code base and requires explicit or inferred instructions of what to run. Override files may not also meet the minimum specifications for a Docker Compose to run, and will not run by itself.
- **Real-world Examples.** We found 53 examples of override files. The `share` project<sup>22</sup> uses an override file to orchestrate a dev mode. The `board` project<sup>23</sup> also uses an override file for development. The `otobo` project uses override files to change the ports of services with one being “http.yml”<sup>24</sup>.
- **Supporting Evidence.** These examples of the pattern were found by counting the observations among repository context as seen in Table 4 under “Override file”.
- **Related Patterns.**






Fig. 9: High-level overview of services in the Certificate Generation and Mapping pattern.

#### 4.5 Certificate Generation and Mapping

- **Motivation.** Developers wishing to provide secure transport (SSL/TLS) to their services may generate a certificate that is used through a reverse proxy.
- **Context/Applicability.** Certificates help enable encrypted communications between machines and also contain identification information about chains of authority and trust. Generating and managing these certificates can be cumbersome, hence the creation of container services to generate certificates for reverse proxy services.
- **Description.** A service container that generates certificates is used. The volumes containing the certificates of the certificate service are mapped to the same volumes of the reverse proxy service. The reverse proxy is configured to use the mapped volume certificates of the certificate service.
- **Advantages.** The certificate issuance and management process is simplified via service orchestration as no certificate files need to be transferred and no command line operations need to be manually executed.
- **Potential Issues.** It can be more difficult to debug between containers.
- **Real-world Examples.** This pattern was discovered in 9 Docker Compose files. The `PeerTube` project uses `certbot` to issue a certificate and map it to the reverse proxy service with a volume.<sup>25</sup> The `sish` project uses `dnsrobocert` to generate a certificate for its main application container.<sup>26</sup>

<sup>22</sup> <https://github.com/MrDemonWolf/share/blob/b231c/docker-compose.override.yml.example>

<sup>23</sup> <https://github.com/RestyaPlatform/board/blob/5804f/docker-compose.override.yml.example>

<sup>24</sup> <https://github.com/RotherOSS/otobo/blob/24d49/scripts/docker-compose/http.yml>

<sup>25</sup> <https://github.com/Chocobozzz/PeerTube/blob/1606ac2/support/docker/production/docker-compose.yml>

<sup>26</sup> <https://github.com/antoniomika/sish/blob/62035/deploy/docker-compose.yml>

The `azuracast` project generates a certificate and maps it to its reverse proxy service via a volume.<sup>27</sup>

- **Supporting Evidence.** These examples were found by counting the Docker Compose file service types that both contained *certificate service* and *reverse proxy service* as seen in Table 2.
- **Related Patterns.**

#### 4.6 Container Management Services

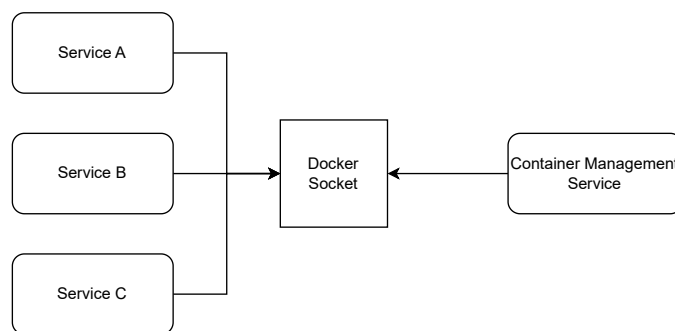


Fig. 10: High-level overview of the Container Management Services pattern.

- **Motivation.** To improve the deployment workflow, container management services are used to maintain the latest image versions of services in a container registry.
- **Context/Applicability.** Instead of re-pulling images and rebuilding the containers of your Docker Compose orchestration, a service can be started to poll a Docker Container registry for container image updates.
- **Description.** A container management service is set up to connect directly with the Docker socket and is configured to watch specific containers for updates in container registries. For example, the watchtower service polls the container registry for any updated to the associated image of a container and automatically updates the container if the associated image is updated.
- **Advantages.** Using a container management service helps reduce maintenance costs as service containers can be upgraded to the latest version without any user intervention.
- **Potential Issues.** There is overhead to manage the management service. As well, since the container management service pulls automatically from

<sup>27</sup> <https://github.com/AzuraCast/AzuraCast/blob/97aaa/docker-compose.multisite.yml>

a container registry, breaking changes to the container registry could break the current software system.

- **Real-world Examples.** We count 14 Docker Compose files using this pattern. The `budibase` project uses `watchtower` to manage container updates.<sup>28</sup> The `egroupware` project also manages container updates with `watchtower`.<sup>29</sup> While the `appsmith` project uses `watchtower` to update containers.<sup>30</sup>
- **Supporting Evidence.** The examples were found by counting the service types as seen in Table 2 under *Container management service*.

#### 4.7 Database Initialization Service with Database Service




Fig. 11: High-level overview of the Database Initialization Service with Database Service pattern.

- **Motivation.** Databases are initialized via an external service that shuts down after the task is run.
- **Context/Applicability.** To import data into an existing database, a database initialization service can be created. This is often seen in development and testing service orchestrations.
- **Description.** The database initialization service is mapped to the database service via environment variables. The database initialization service runs scripts to initialize the database in the database service.
- **Advantages.** Short running initialization container that is separate from the main application service allows for separation of concerns in the software system.
- **Potential Issues.** There is an overhead cost for setting up the initialization service where data must be dumped and scripts need to be created to import the data.
- **Real-world Examples.** We count 9 Docker Compose files using this pattern. The `budibase` project<sup>31</sup> initializes `couchdb` with an initialization service. The `CKAN` project creates their own customer container that loads

<sup>28</sup> <https://github.com/Budibase/budibase/blob/56147/hosting/docker-compose.yaml>

<sup>29</sup> <https://github.com/EGroupware/egroupware/blob/1c053/doc/docker/docker-compose.yml>

<sup>30</sup> <https://github.com/appsmithorg/appsmith/blob/382ea53/app/server/appsmith-server/src/main/resources/docker-compose.yml>

<sup>31</sup> <https://github.com/Budibase/budibase/blob/56147/hosting/docker-compose.yaml>

their database from a tabular file like CSV.<sup>32</sup> The `graphql` engine project initializes their example project using a migration container.<sup>33</sup>

- **Supporting Evidence.** These examples were found by counting the Docker Compose file that contained *database service* and *database init service* service types together as seen in Table 2.

#### 4.8 Database Administration Service with Database Service




Fig. 12: High-level overview of the Database Administration Service with Database Service pattern.

- **Motivation.** Databases are often easier to manage with a graphical user interface (GUI).
- **Context/Applicability.** Developers wishing to access the database visually without installing any tools on their local machine and without using any external online web services.
- **Description.** A database administration service such as `adminer` or `phpmyadmin` are orchestrated into the same network as the database service. Users can then browse the database via the database administration service through a web browser.
- **Advantages.** An orchestrated database can be managed with a graphical interface without any external concerns such as security.
- **Potential Issues.** The database administration service provides an interface into the database that can be exploited if the service is exposed.
- **Real-world Examples.** This pattern was discovered in 22 Docker Compose files. The `libretime` project uses `adminer` to manage a postgres instance.<sup>34</sup> The `personal management system` project uses `adminer` to manage a mariadb instance.<sup>35</sup> The `unmark` project uses `adminer` to manage a mysql instance.<sup>36</sup>
- **Supporting Evidence.** These examples were found by counting the service type combination of *database service* and *database administration service* present in Docker Compose files as seen in Table 2.

<sup>32</sup> <https://github.com/ckan/ckan/blob/e1568fd/contrib/docker/docker-compose.yml>

<sup>33</sup> <https://github.com/hasura/graphql-engine/blob/2325755/community/sample-apps/tic-tac-toe-react/docker-compose.yml>

<sup>34</sup> <https://github.com/LibreTime/libretime/blob/05342/docker-compose.yml>

<sup>35</sup> <https://github.com/Volmarg/personal-management-system/blob/8b377/docker-compose.yml>

<sup>36</sup> <https://github.com/cdevroe/unmark/blob/da271/docker-compose.yml>

## 4.9 Service Labels are used to Configure Reverse Proxy

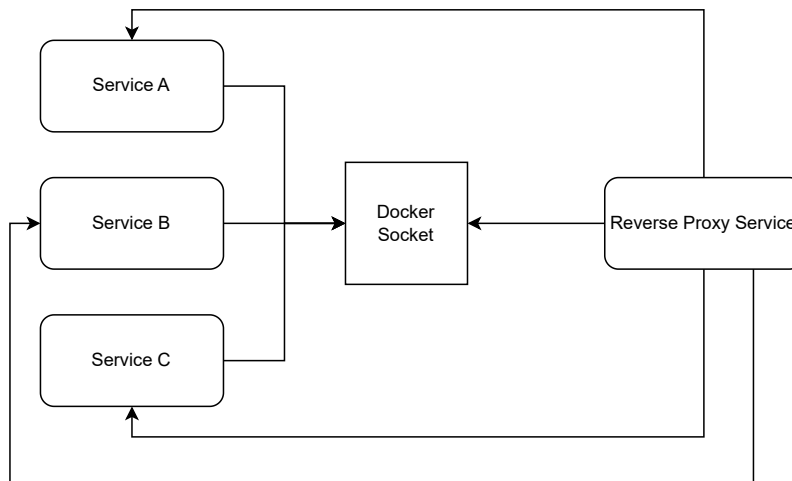


Fig. 13: High-level overview of the Service Labels are used to Configure Reverse Proxy pattern.

- **Motivation.** Reverse proxies such as nginx require a configuration file on how to route its proxy to different servers. Instead of using a configuration file, a reverse proxy can be configured with the Docker service labels. Docker service labels are labels that can be given to services in a Docker Compose file. The use of a reverse proxy as a pattern has been described by Sousa et al. (2018a, 2015).
- **Context/Applicability.** To simplify reverse proxy configuration, all reverse proxy configuration can be centralized to a Docker Compose file.
- **Description.** Services such as the `traefik` reverse proxy uses the Docker API (via the Docker Socket) to obtain the configuration of routing to services via their service labels. The `traefik` reverse proxy routes requests to the relevant containers.
- **Advantages.** Instead of managing configuration files that need to be mounted to a reverse proxy service, the configuration can be centralized into a Docker Compose file.
- **Potential Issues.** It can be difficult to debug if the configuration for `traefik` does not work. There is an issue with Docker API security as the host may be able to be accessed via `traefik`'s access to the Docker API (Traefik 2022).
- **Real-world Examples.** We count 15 Docker Compose files using this pattern. The `isotope mail client project` uses labels to configure their

`traefik` proxy.<sup>37</sup> The `shaarli` bookmarking project uses labels to configure `traefik`.<sup>38</sup> The `zenbot` project has a Docker Compose file that uses `traefik` with services that have `traefik` service labels.<sup>39</sup>

- **Supporting Evidence.** These examples were found by counting the observations among services in Table 3 under “Service labels are configuration for reverse proxy”.

#### 4.10 Mail Service Testing




Fig. 14: High-level overview of the Mail Service Testing pattern.

- **Motivation.** To test the email functionalities of an app, an SMTP service needs to be created. The creation of SMTP services can be complex hence the existence of SaaS providers like Mailtrap, SendGrid, Amazon SES, and Mailgun.
- **Context/Applicability.** Developers wishing to keep development local may wish to orchestrate their own SMTP services.
- **Description.** When the mail service is orchestrated, the mail service may connect with the application service with environment variables explicitly. Otherwise the mail service is connected internally by being hard-coded into the application source code.
- **Advantages.** There is no reliance on external services that require internet connections and mail can be retrieved locally as opposed to email providers that may block adhoc SMTP services that attempt to deliver to a provider’s inbox.
- **Potential Issues.** Mail services are an additional service in the software system that have configuration and debugging costs. The alternative are SaaS providers which reduce these costs by providing SMTP credentials that are guaranteed to work with an Internet connection.
- **Real-world Examples.** This pattern was discovered in 62 Docker Compose files. The `roadiz` project uses `mailhog` to test emails.<sup>40</sup> The `shopware`

<sup>37</sup> <https://github.com/manusa/isotope-mail/blob/37722/deployment-examples/docker-compose.yml>

<sup>38</sup> <https://github.com/shaarli/Shaarli/blob/b7c50a5/docker-compose.yml>

<sup>39</sup> <https://github.com/DeviaVir/zenbot/blob/0cb3541/docker-compose-traefik.yml>

<sup>40</sup> <https://github.com/roadiz/roadiz/blob/ac370b0/docker-compose.yml>

project also uses `mailhog` to test emails in development.<sup>41</sup> The `umbraco` CMS project uses `smtp4dev` to test emails in development.<sup>42</sup>

- **Supporting Evidence.** These examples were found by counting the service types in Table 2 under *mail service*.

#### 4.11 Application Service with Database




Fig. 15: High-level overview of the Application Service with Database pattern.

- **Motivation.** Most software systems contain some sort of backend database service. Therefore, it would be beneficial to know how these backend database services are orchestrated.
- **Context/Applicability.** For any software system that does not store data in memory and stores data in a backend database service such as: MySQL, MariaDB, Postgres, or MongoDB. The backend service is not part of the main application container.
- **Description.** A main application uses the “`depends_on`” instruction to define a relationship between the main application and backend database service. This signifies that the main application will not start up until the backend database service has started. The backend database service mounts a volume for persistent storage, so that the data in the database is not destroyed when the “`docker compose down`” command is run. The backend database service might also bind to a port on the host machine, the bound port can be used to directly access the database from the host machine.
- **Advantages.** This pattern is useful when needing to orchestrate a backend service within your software system especially when testing and developing applications as containers are quite flexible. Spinning up and destroying databases simply requires the definition of a database service in a Docker Compose file.
- **Disadvantages.** If the backend database service is hosted on the same host machine, it can lead to catastrophic failures in a production environment when the host machine is down.
- **Potential Issues.** Using “`depends_on`” does not guarantee that your application will connect to the backend database service on the first try. This

<sup>41</sup> <https://github.com/shopware/platform/blob/d93ca0f/docker-compose.yaml>

<sup>42</sup> <https://github.com/umbraco/Umbraco-CMS/blob/a73c7bb/.devcontainer/docker-compose.yml>

is because although a container may have started, the service application might not be ready to accept applications. Therefore, Docker Inc. (2022c) recommends that a check is built into your application container to ensure the database service is ready to accept connections.

- **Real-world Examples.** We found 121 Docker Compose files using this pattern. For example, we can see this in the `accent` project<sup>43</sup>, `anchr` project<sup>44</sup>, and `stringer` project<sup>45</sup>.
- **Supporting Evidence.** This was found by counting the Docker Compose files that had the service types: *frontend service* and *database service* but not *caching service* as seen in Table 2.

#### 4.12 Application Service with Database and Caching




Fig. 16: High-level overview of the Application Service with Database and Caching pattern.

- **Motivation.** To better support backend database services, often caching layers are used to improve performance when retrieving data from backend database services.
- **Context/Applicability.** This pattern is useful when looking to develop and debug software systems that have a caching mechanism. The most common caching service found is Redis.
- **Description.** The main application uses the “depends\_on” instruction to define a relationship between the main application and backend database service. It uses another “depends\_on” instruction to define a relationship between the caching service. The backend database service mounts a volume, so that the data in the database is not destroyed when the “docker compose down” command is run. The caching service also mounts a volume too, to avoid data being ephemeral. The backend database service might also bind to a port on the host machine, the binded port can be used to directly access the database from the host machine. The caching

<sup>43</sup> <https://github.com/mirego/accent/blob/22d5847/docker-compose.yml>

<sup>44</sup> <https://github.com/muety/anchr/blob/320ecb5/docker-compose.yml>

<sup>45</sup> <https://github.com/swanson/stringer/blob/0b64dc0/docker-compose.yml>



service might also bind to a port on the host machine for direct access to the caching service.

- **Advantages.** The pattern is useful for developing and testing software systems that wish to cache data. Since all the services are orchestrated locally, services can be easily accessed without any connectivity issues.
- **Disadvantages.** Orchestrating all the services on the same host machine can lead to downtime as if any one service fails then the application may fail to run.
- **Potential Issues.** The backend database service and caching service containers may start up, but the services may not be ready to accept connections. As a result, Docker Inc. (2022c) recommends that the main application should implement a check to ensure that the database and caching services are online.
- **Real-world Examples.** This pattern was discovered in 42 Docker Compose files. For example, we can see this in the *miaou* project<sup>46</sup>, *wildduck* project<sup>47</sup>, and *overleaf* project<sup>48</sup>.
- **Supporting Evidence.** This was found by counting the Docker Compose files that had the service types: *frontend service* and *database service* and *caching service* as seen in Table 2.

#### 4.13 HTTP Reverse Proxy Service

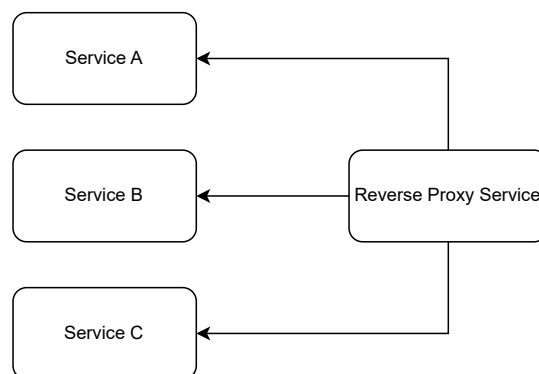


Fig. 17: High-level overview of the HTTP Reverse Proxy Service pattern.

- **Motivation.** Services wishing to be exposed over a single port are combined using a reverse proxy service. Burns and Oppenheimer (2016); Sousa

<sup>46</sup> <https://github.com/Canop/miaou/blob/7b251dc/docker-compose.yml>

<sup>47</sup> <https://github.com/nodemailer/wildduck/blob/ca3a365/docker-compose.yml>

<sup>48</sup> <https://github.com/overleaf/overleaf/blob/fd36c4136/docker-compose.yml>

et al. (2018a, 2015) has described the use of a reverse proxy as a pattern. While Fehling et al. (2014) describes a similar “Application Component Proxy” pattern.

- **Context/Applicability.** This pattern can be used when you wish to expose your service over default ports such as port 80 for HTTP and port 443 for HTTPS. Reverse proxy services such as Caddy and Nginx are often used as reverse proxies.
- **Description.** This pattern works as follows:
  - A reverse proxy service binds to ports such as 80 and 443. A volume is also mounted to the reverse proxy service to keep the configuration files persistent. The folder containing the log files can also be mounted to a volume.
  - The reverse proxy service may use a “depends\_on” instruction to define a relationship to the application(s) that it proxies. However, this is not necessary and can impede the startup time of accepting connections for the software system.
- **Advantages.** Managing the reverse proxy over a container reduces the need to install software on the host machine.
- **Disadvantages.** Since the reverse proxy is hosted within a container, there is additional overhead and complexity to running the reverse proxy software as the container needs to be managed.
- **Potential Issues.** Using a reverse proxy over a container can be more difficult to debug when log files are not mounted to a volume.
- **Real-world Examples.** There are 92 Docker Compose files using this pattern. For example, we can see this in the `goodwork` project<sup>49</sup>, `linkace` project<sup>50</sup>, and `PatrowlManager` project<sup>51</sup>.
- **Supporting Evidence.** These examples were found by counting the service types in Table 2 under *reverse proxy service*.

#### 4.14 Duplicate Service Reuse

- **Motivation.** There are two motivations for duplicate service reuse. First, applications sometimes bundle their functions into one image but would like their functions to be executed on separate instances. Second, application instances can be duplicated and isolated to test software systems that run multiple instances of the same kind of application in a “dev” or “test” mode.
- **Context/Applicability.** This pattern is useful when you wish to run multiple instances of your application to perform different functionalities of your application and to isolate these instances.
- **Description.** Services containing the same base image are defined in one of two ways. First, the service can be defined with multiple of the same

<sup>49</sup> <https://github.com/iluminar/goodwork/blob/f452b7d5/docker-compose.yml>

<sup>50</sup> <https://github.com/Kovah/LinkAce/blob/6589c92/docker-compose.yml>

<sup>51</sup> <https://github.com/Patrowl/PatrowlManager/blob/7567c27/docker-compose.yml>




Fig. 18: High-level example of the Duplicate Service Reuse pattern.

image instances or be built from the same Dockerfile. Or the service can be duplicated using the “YAML Anchor and Alias” pattern (Section 4.2) or “Docker Compose Service Inheritance” pattern (Section 4.3). Next, the services are configured for their specific use cases, this might include changing the “command” option or having different volumes and ports for the service.

- **Advantages.** Multiple processes of the application can be easily managed. To update versions of containers only a single image needs to be pulled. There is reduced complexity in the amount of Docker images that need to be maintained.
- **Disadvantages.** Since the duplicated application can contain all functionalities of a software system, the codebase can be quite large making increasing the time to pull or build the Docker image of the duplicate application. Instead the application image could be separated into microservices.
- **Potential Issues.** If one part of the codebase is vulnerable in the application, the whole application could be compromised since they are all using the same image without separation of functionalities.
- **Real-world Examples.** We found 17 Docker Compose files using this pattern. For example, we can see this in the `mastodon` project<sup>52</sup>, `chaskiq` project<sup>53</sup>, and `Black Candy` project<sup>54</sup>.
- **Supporting Evidence.** These examples were found by counting the service types in Table 2 under “duplicate service reuse”.

## 5 Discussion

Our research is motivated by the popular usage of Docker Compose to deploy and manage software systems. We find that among configurations there are common kinds of service images used such as reverse proxies and databases.

<sup>52</sup> <https://github.com/tootsuite/mastodon/blob/8f7308da8/docker-compose.yml>

<sup>53</sup> <https://github.com/chaskiq/chaskiq/blob/8fb33b/docker-compose.yml>

<sup>54</sup> [https://github.com/blackcandy-org/black\\_candy/blob/1d14fb9/docker-compose.yml](https://github.com/blackcandy-org/black_candy/blob/1d14fb9/docker-compose.yml)

Furthermore, there is a common set of Docker Compose options used as well. This motivates us to investigate the patterns of configuration for the service images used to better understand their implications in self-hosted projects and wider applications.

### 5.1 Identifying Common Patterns

To identify common patterns we use an open coding process where we manually inspect each Docker Compose file and project to determine any interesting observations. As opposed to using strictly quantitative methods to discover patterns, we also use qualitative methods to perform an analysis. Using qualitative methods is beneficial as we can notice patterns that are not easily discoverable by statistics. For example, automatic Docker Compose file generation (Section 4.1) is a pattern that would not have been found from a quantitative analysis due to its low count. As well, the observation cannot be easily discovered by looking at Docker Compose files only because the project repository has to be inspected to understand how the files were created.

### 5.2 Patterns in Self-hosted Projects

Our patterns emerge from open coding software systems that are meant to be self-hosted on any infrastructure that one chooses. The patterns that we present in Section 4 are representative of Docker Compose usage in software systems run by a diverse user base making it beneficial for developers wishing to use Docker Compose for their own projects. We find patterns that can help improve the developer experience such as using container management services described in Section 4.6. The patterns that we discover in self-hosted projects can be used as examples of how to orchestrate containers in future software systems.

### 5.3 Patterns and Wider Applications

Docker Compose users can benefit from these findings by better understanding patterns of what options to use when designing their software systems.

This study helps to empirically validate what common patterns are already being used. Researchers of Docker can use the empirical observations of patterns of Docker Compose usage to motivate further studies on container orchestration in other development communities. The empirical observations of this paper also suggest that there is merit to recovering architecture traces from Docker Compose files since containers that run services are often the building blocks of architecture (e.g. databases).

## 5.4 Comparison with Other Infrastructure as Code Patterns

We compare the patterns we discover with prior work discussed in Section 2.3 and outline how our patterns are similar and different to demonstrate the practicality of our work.

Prior research in *Infrastructure as Code* (IaC), patterns of usage have been investigated by (Fehling et al. 2014; Rahman et al. 2019; Ksontini et al. 2021; Shamim et al. 2020; Burns and Oppenheimer 2016). Rahman et al. (2019) investigates security patterns in Puppet that deal with specific instructions that are used such as using hard-coded passwords. In comparison, our patterns focus on how to orchestrate services at a higher level. Shamim et al. (2020) identify implementation practices that can better secure Kubernetes deployments such as implementing role-based access control authorization, continuously applying security patches, and implementing security policies for Kubernetes pods and networks. By contrast, our work focuses explicitly on how developers might deploy software services such as databases to solve typical software application deployments.

Burns and Oppenheimer (2016) identify patterns for container-based distributed systems. In our paper, we identify the “HTTP Reverse Proxy Service” pattern (Section 4.13) which resembles the ambassador pattern described by Burns and Oppenheimer (2016) — a proxy is used to route to different services from a single server. We also identify the “Duplicate Service Reuse” (Section 4.14) that can be used for the multi-node application patterns described by Burns and Oppenheimer (2016). Therefore, certain patterns we use bear similarities to the patterns found in container-based distributed systems as described in (Burns and Oppenheimer 2016).

Ksontini et al. (2021) identify refactoring types for Dockerfiles and Docker Compose that suggest some best practices. Our patterns support many of their best practices identified from refactoring types. Our “Docker Compose Override Use Case” pattern in Section 4.4 correspond with the refactoring types in (Ksontini et al. 2021): “Extract ports attribute into an override Docker-compose file”; “Extract volume attribute into an override Docker-compose file”; “Extract ENV attribute into an override Docker-compose file or use .env file”; “Extract service into an override Docker-compose file”. While our “Docker Compose Service Inheritance” pattern in Section 4.3 corresponds with the “Add Extends attribute to inherit configuration from an existing service thus avoiding duplication” refactoring type in (Ksontini et al. 2021).

While we do not identify patterns that explicitly set names for services, containers, and volumes, or adding tags to images like the refactoring types identified in (Ksontini et al. 2021), we do identify the “Container Management Services” pattern that can be used to maintain the latest images of services which negates the need for refactoring. We outline in our “Application Service with Database” pattern in Section 4.11 and “Application Service with Database and Caching” in Section 4.12 how dependencies can be an issue in service orchestration which corresponds to the “Order services based on their dependency order” refactoring type in (Ksontini et al. 2021). Finally,

Ksontini et al. (2021) identify the “Add ENV variables to store useful system-wide values” refactoring type, where we make suggestions in the “Database Initialization Service with Database Service” pattern (Section 4.7) and “Mail Service Testing” pattern (Section 4.10) to use environment variables to connect to these services. As a result, our patterns identified with our qualitative analysis align with many of the refactoring types identified by Ksontini et al. (2021).

Sousa et al. (2018a, 2015) presents the idea of a reverse proxy pattern where services are routed to their appropriate network addresses. The most closely related pattern we find is the “Service Labels are used to Configure Reverse Proxy” pattern described in Section 4.9 where service routing can automatically be determined through service labels. Also related is using a reverse proxy in the “HTTP Reverse Proxy Service” pattern (Section 4.13).

## 6 Threats to Validity

In terms of construct validity, the projects that we analyze may not be representative of all real world usage of Docker Compose. As such, the empirical observations seen cannot be assumed to be applied to all projects using Docker Compose. Even so, the projects of the self-hosted community that we do analyze provides some nuance to usage that may not be obtained by analyzing large datasets with unknown contexts. Furthermore, we also assume that the git history of projects has not been rewritten. We also assume that the data is accurately retrieved and that the filtering process for identifying Docker Compose usage is accurate.

With regards to internal validity, we construct our patterns through an open coding process whereby patterns of usage are collected through an iterative process. Since we use an open coding process, some patterns may be overlooked due to the manual nature of the analysis. As well, there may be preconceived notions as to what patterns might be expected when performing the open coding process. To reduce the risk of preconceived notions and overlooking patterns, we review our data with an open mind in three dimensions: (1) at a service level in Docker Compose files, (2) at an orchestration level with how services in a Docker Compose configuration work together, and (3) at a repository level with how Docker Compose files operate within a project repository.

Finally, in consideration of external validity, our dataset is not comprehensive and may not generalize to all projects. Since we only analyze publicly available projects, private projects such as internal company projects may use Docker Compose differently. Furthermore, full architectures may not be reflected since services can run on a single container (e.g., a database and application run in the same container). However, this is unlikely among most projects as it would defeat the purpose of using Docker Compose. We provide a replication package (Eng et al. 2024) outlining our process and scripts for those wishing to analyze more projects.

## 7 Conclusions and Future Work

This paper qualitatively and quantitatively explored a set of successful self-hosted software-based service projects that employ Docker Compose to orchestrate and deploy various software services using Docker containers. Based on these analyses, typical patterns of service architecture, orchestration, and deployment were identified, counted, characterized, and named. Naming and characterizing the patterns allows others to discuss their architectures using the pattern names as shortcuts rapidly. These patterns also describe to developers patterns that have been successfully deployed. We demonstrate these projects are successful based on having more GitHub stars than prior studies like (Ibrahim et al. 2021). The greater number of GitHub stars indicates a project’s attention from a wider community and shows the value and significance of the project.

To encourage further patterns research and replicability, we contribute a replication package (Eng et al. 2024) that includes an open-coded dataset of concepts in Docker Compose files, as well as scripts for analysis and extracting our selected Docker Compose projects.

We find many common patterns one would expect from self-hosted software services such as the use of an “HTTP Reverse Proxy Service” (Section 4.13), or “Application Service with Database and Caching” (Section 4.12). We found interesting patterns that have to do with deploying databases, especially initializing databases before the application can use them, as described by the “Database Initialization Service with Database Service” pattern in Section 4.7. The “Service Labels are used to Configure Reverse Proxy” pattern described in Section 4.9 is also interesting, as it shows how reverse proxies can be dynamically configured by Docker Compose leveraging service labels in the Docker Compose files used to deploy and orchestrate the app.

### Data Availability

The datasets generated and analyzed during this study are available as part of our replication package on <https://zenodo.org/records/10648448>.

### References

- Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison Wesley, Boston, MA, October 1994.
- Gregor Hohpe and Bobby Woolf. *Enterprise integration patterns: Designing, building, and deploying messaging solutions*. Addison-Wesley Professional, 2004.
- Christoph Fehling, Frank Leymann, Ralph Retter, Walter Schupeck, and Peter Arbitter. *Cloud computing patterns: fundamentals to design, build, and manage cloud applications*. Springer, 2014.
- Christian Endres, Uwe Breitenbücher, Michael Falkenthal, Oliver Kopp, Frank Leymann, and Johannes Wettinger. Declarative vs. imperative: two modeling patterns for the automated deployment of applications. In *Proceedings of the 9th International Conference*

- on *Pervasive Patterns and Applications*, pages 22–27. Xpert Publishing Services (XPS), 2017.
- Scott McCarty. Container tidbits: When should i break my application into multiple containers?, March 2016. URL <https://web.archive.org/web/20230820154241/https://www.redhat.com/en/blog/container-tidbits-when-should-i-break-my-application-multiple-containers>.
- Microsoft. Containerizing monolithic applications, April 2022. URL <https://github.com/dotnet/docs/blob/0ceaa6b/docs/architecture/microservices/architect-microservice-container-applications/containerize-monolithic-applications.md>.
- Scott Carey. Complexity is killing software developers, Nov 2021. URL <https://www.infoworld.com/article/3639050/complexity-is-killing-software-developers.html>.
- Anton Podviaznikov. The versatility of docker compose, March 2017. URL <https://web.archive.org/web/20220808200330/https://runnable.com/blog/the-versatility-of-docker-compose>.
- Md Hasan Ibrahim, Mohammed Sayagh, and Ahmed E Hassan. A study of how docker compose is used to compose multi-component systems. *Empirical Software Engineering*, 26(6):1–27, 2021.
- Thomas Erl. *Service-Oriented Architecture: Concepts, Technology, and Design*. Prentice Hall PTR, Upper Saddle River, NJ, 2005.
- The Kubernetes Authors. Kubernetes, June 2022. URL <https://web.archive.org/web/20220601000133/https://kubernetes.io/>.
- Maria A Rodriguez and Rajkumar Buyya. Container-based cluster orchestration systems: A taxonomy and future directions. *Software: Practice and Experience*, 49(5):698–719, 2019.
- Eddy Truyen, Dimitri Van Landuyt, Davy Preuveneers, Bert Lagaisse, and Wouter Joosen. A comprehensive feature comparison study of open-source container orchestration frameworks. *Applied Sciences*, 9(5):931, 2019.
- Jordan Webb. The container orchestrator landscape, August 2022. URL <https://lwn.net/Articles/905164/>.
- Ian Miell and Aidan Sayers. *Docker in practice*. Simon and Schuster, New York, NY, 2019.
- Sean P Kane and Karl Matthias. *Docker: Up & Running*. ”O’Reilly Media, Inc.”, Sebastopol, California, 2023.
- Pethuru Raj, Jeeva S Chelladhurai, and Vinod Singh. *Learning Docker*. Packt Publishing Ltd, Birmingham, UK, 2015.
- Michele Guerriero, Martin Garriga, Damian A Tamburri, and Fabio Palomba. Adoption, support, and challenges of infrastructure-as-code: Insights from industry. In *2019 IEEE International conference on software maintenance and evolution (ICSME)*, pages 580–589. IEEE, 2019.
- Sebastian Van Stijn et al. Code of JSON Schema Compose Specification Found in Docker Compose Source Code, November 2023. URL <https://web.archive.org/web/20231207060306/https://raw.githubusercontent.com/compose-spec/compose-spec/ec342b4/schema/compose-spec.json>.
- Docker Inc. Announcing the compose specification, April 2020. URL <https://web.archive.org/web/20200407131941/https://www.docker.com/blog/announcing-the-compose-specification/>.
- Docker Inc. Docker compose specification, March 2022a. URL <https://web.archive.org/web/20220322124537/https://compose-spec.io/>.
- nerdctl Authors. nerdctl: Docker-compatible cli for containerd, July 2023. URL <https://web.archive.org/web/20230712211127/https://github.com/containerd/nerdctl>.
- Okteto, Inc. Docker compose reference — okteto documentation, June 2023. URL <https://web.archive.org/web/20230605123922/https://www.okteto.com/docs/reference/compose/>.
- Podman Compose Authors. Podman compose, July 2023. URL <https://web.archive.org/web/20230713103640/https://github.com/containers/podman-compose>.
- Rancher Labs. Rancher docs: Rancher compose, August 2023. URL <https://web.archive.org/web/20220826040438/https://rancher.com/docs/rancher/v1.6/>



- en/cattle/rancher-compose/.
- Kompose Authors. Kompose - convert your docker compose file to kubernetes or openshift, January 2022. URL <https://web.archive.org/web/20220119234120/https://kompose.io/>.
- Docker Inc. Ways to set environment variables with compose, October 2023a. URL <https://github.com/docker/docs/blob/6e199de/content/compose/environment-variables/set-environment-variables.md>.
- Docker Inc. Using profiles with compose, October 2023b. URL <https://github.com/docker/docs/blob/e86716e/content/compose/profiles.md>.
- Docker Inc. Extend your compose file, August 2022b. URL <https://github.com/docker/docs/blob/15e9e1e/content/compose/multiple-compose-files/extends.md>.
- Docker Inc. Merge compose files, August 2023c. URL <https://github.com/docker/docs/blob/15e9e1e/content/compose/multiple-compose-files/merge.md>.
- Docker Inc. How to use secrets in docker compose, August 2023d. URL <https://github.com/docker/docs/blob/15e9e1e/content/compose/use-secrets.md>.
- Docker Inc. Control startup and shutdown order in compose, October 2022c. URL <https://github.com/docker/docs/blob/88b4505/compose/startup-order.md>.
- Yang Zhang, Gang Yin, Tao Wang, Yue Yu, and Huaimin Wang. An insight into the impact of dockerfile evolutionary trajectories on quality and latency. In *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, volume 1, pages 138–143. IEEE, 2018.
- Yang Zhang, Huaimin Wang, and Vladimir Filkov. A clustering-based approach for mining dockerfile evolutionary trajectories. *Science China Information Sciences*, 62(1):1–3, 2019.
- Kalvin Eng and Abram Hindle. Revisiting dockerfiles in open source software over time. *arXiv preprint arXiv:2103.12298*, 2021.
- Changyuan Lin, Sarah Nadi, and Hamzeh Khazaei. A large-scale data set and an empirical study of docker images hosted on docker hub. In *2020 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 371–381. IEEE, 2020.
- Jürgen Cito, Gerald Schermann, John Erik Wittern, Philipp Leitner, Sali Zumberi, and Harald C Gall. An empirical analysis of the docker container ecosystem on github. In *2017 IEEE/ACM 14th International Conference on Mining Software Repositories (MSR)*, pages 323–333. IEEE, 2017.
- Tianyin Xu and Darko Marinov. Mining container image repositories for software configuration and beyond. In *Proceedings of the 40th International Conference on Software Engineering: New Ideas and Emerging Results*, pages 49–52, 2018.
- Akond Rahman, Chris Parnin, and Laurie Williams. The seven sins: Security smells in infrastructure as code scripts. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, pages 164–175. IEEE, 2019.
- Emna Ksontini, Marouane Kessentini, Thiago do N Ferreira, and Foyzul Hassan. Refactorings and technical debt in docker projects: An empirical study. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 781–791. IEEE, 2021.
- Akond Rahman, Effat Farhana, Chris Parnin, and Laurie Williams. Gang of eight: A defect taxonomy for infrastructure as code scripts. In *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*, pages 752–764. IEEE, 2020a.
- Akond Rahman, Effat Farhana, and Laurie Williams. The ‘as code’activities: development anti-patterns for infrastructure as code. *Empirical Software Engineering*, 25(5):3430–3467, 2020b.
- Md Shazibul Islam Shamim, Farzana Ahamed Bhuiyan, and Akond Rahman. Xi commandments of kubernetes security: A systematization of knowledge related to kubernetes security practices. In *2020 IEEE Secure Development (SecDev)*, pages 58–64. IEEE, 2020.
- Brendan Burns and David Oppenheimer. Design patterns for container-based distributed systems. In *8th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 16)*, 2016.
- Tiago Boldt Sousa, Hugo Sereno Ferreira, and Filipe Figueiredo Correia. Overview of a pattern language for engineering software for the cloud. In *Proceedings of the 25th*

- Conference on Pattern Languages of Programs*, pages 1–9, 2018a.
- Tiago Boldt Sousa, Hugo Sereno Ferreira, Filipe Figueiredo Correia, and Ademar Aguiar. Engineering software for the cloud: External monitoring and failure injection. In *Proceedings of the 23rd European Conference on Pattern Languages of Programs*, pages 1–8, 2018b.
- Tiago Boldt Sousa, Hugo Sereno Ferreira, Filipe Figueiredo Correia, and Ademar Aguiar. Engineering software for the cloud: Automated recovery and scheduler. In *Proceedings of the 23rd European Conference on Pattern Languages of Programs*, pages 1–8, 2018c.
- Tiago Boldt Sousa, Hugo Sereno Ferreira, Filipe Figueiredo Correia, and Ademar Aguiar. Engineering software for the cloud: Messaging systems and logging. In *Proceedings of the 22Nd European Conference on Pattern Languages of Programs*, pages 1–14, 2017.
- Tiago Boldt Sousa, Filipe Figueiredo Correia, and Hugo Sereno Ferreira. Patterns for software orchestration on the cloud. In *Proceedings of the 22nd Conference on Pattern Languages of Programs*, pages 1–12, 2015.
- Tiago Boldt Sousa, Hugo Sereno Ferreira, and Filipe Figueiredo Correia. A survey on the adoption of patterns for engineering software for the cloud. *IEEE Transactions on Software Engineering*, 48(6):2128–2140, 2021.
- Gregor Hohpe. Modern examples for enterprise integration patterns, February 2017. URL [https://web.archive.org/web/20230815004255/https://www.enterpriseintegrationpatterns.com/ramblings/eip1\\_examples\\_updated.html](https://web.archive.org/web/20230815004255/https://www.enterpriseintegrationpatterns.com/ramblings/eip1_examples_updated.html).
- Michael Falkenthal, Johanna Barzen, Uwe Breitenbücher, Christoph Fehling, Frank Leymann, A Hadjakos, F Hentschel, and H Schulze. Leveraging pattern application via pattern refinement. In *Proceedings of the International Conference on Pursuit of Pattern Languages for Societal Change (PURPLSOC)*, pages 38–61, 2016.
- Beniamino Di Martino and Antonio Esposito. A rule-based procedure for automatic recognition of design patterns in uml diagrams. *Software: Practice and Experience*, 46(7):983–1007, 2016.
- nodisc et al. Awesome-selfhosted, May 2022. URL <https://github.com/awesome-selfhosted/awesome-selfhosted/tree/1726a40>.
- Kalvin Eng, Abram Hindle, and Eleni Stroulia. Replication Package of "Patterns of Multi-Container Composition for Service Orchestration with Docker Compose", February 2024. URL <https://zenodo.org/records/10648448>.
- Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M German, and Daniela Damian. The promises and perils of mining github. In *Proceedings of the 11th working conference on mining software repositories*, pages 92–101, 2014.
- Davide Spadini, Maurício Aniche, and Alberto Bacchelli. Pydriller: Python framework for mining software repositories. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 908–911, 2018.
- Christian Bird, Peter C Rigby, Earl T Barr, David J Hamilton, Daniel M German, and Prem Devanbu. The promises and perils of mining git. In *2009 6th IEEE International Working Conference on Mining Software Repositories*, pages 1–10. IEEE, 2009.
- Scott Brisson, Ehsan Noei, and Kelly Lyons. We are family: analyzing communication in github software repositories and their forks. In *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 59–69. IEEE, 2020.
- Kelly Blincoe, Jyoti Sheoran, Sean Goggins, Eva Petakovic, and Daniela Damian. Understanding the popular users: Following, affiliation influence and leadership on github. *Information and Software Technology*, 70:30–39, 2016.
- Hudson Borges, Andre Hora, and Marco Tulio Valente. Understanding the factors that impact the popularity of github repositories. In *2016 IEEE international conference on software maintenance and evolution (ICSME)*, pages 334–344. IEEE, 2016.
- Nathan Munaiah, Steven Kroh, Craig Cabrey, and Meiyappan Nagappan. Curating Github for Engineered Software Projects. *Empirical Software Engineering*, 22(6):3219–3253, 2017.
- John V Seidel. Qualitative data analysis, 1998.
- Docker Inc. Fragments, November 2023e. URL [https://github.com/docker/docs/blob/8d8014a/\\_vendor/github.com/compose-spec/compose-spec/10-fragments.md](https://github.com/docker/docs/blob/8d8014a/_vendor/github.com/compose-spec/compose-spec/10-fragments.md).

---

Traefik. Traefik docker documentation, December 2022. URL <https://github.com/traefik/traefik/blob/519ed8b/docs/content/providers/docker.md#docker-api-access>.