# wharfrat Documentation

*Release REL*

**Julian Phillips**

**Feb 05, 2023**

# Contents:


**1 Introduction** **1**

**2 Installation** **3**

**3 Basic Usage** **5**

**4 Configuration** **7**
- 4.1 Project Configuration . . . 7
- 4.2 Crate Configuration . . . 7
- 4.3 Local Configuration . . . 9

**5 Accessing Containers** **11**
- 5.1 wharfrat run . . . 11
- 5.2 Exposing Commands . . . 11

**6 Environments** **13**
- 6.1 Purpose . . . 13
- 6.2 Creating an Environment . . . 13
- 6.3 Using an Environment . . . 14

**7 Examples** **15**

**8 Indices and tables** **17**

CHAPTER 1

# Introduction

wharfrat is intended to make it easy and convenient to use a development environment in a docker container. The benefits of this are:

- **Simple:** A new development environment is setup with a single simple command, meaning new team members are ready to go immediately.
- **Shared:** Everyone with access to the project (and docker) has access to the development environment.
- **Controlled:** Everyone gets a development environment created from the same image - no more "works for me" issues.
- **Versioned:** The configuration is version controlled, meaning you get the development environment that matches the code branch you are working on.

CHAPTER 2

# Installation

...

CHAPTER 3

# Basic Usage

…

CHAPTER 4

# Configuration

**Table of Contents**


- *Configuration*
  - *Project Configuration*
  - *Crate Configuration*
  - *Local Configuration*


## 4.1 Project Configuration

## 4.2 Crate Configuration

The table below lists the settings available for each crate, their types and default values (if the default is not empty):

| | | |
|---|---|---|
| cap-add | array of strings | capabilities to enable for the container |
| cap-drop | array of strings | capabilities to disable for the container |
| copy-groups | array of strings | groups to copy from the host to the container |
| env | table of strings | mapping from environment variable name to value |
| env-blacklist | array of strings | host environment variables to drop |
| env-whitelist | array of strings | host environment variables to keep |
| groups | array of strings | groups the user should be in |
| hostname | string | hostname for container (default: "dev") |
| image | string | name of image to create container from |
| image-cmd | string | a script to run to determine the image name (instead of using image). |
| mount-home | bool | should /home be mounted into container (default: true) |
| network | string | the network to connect the container to |
| path-append | array of strings | extra paths to add to end of PATH |
| path-prepend | array of strings | extra paths to add to beginning of PATH |
| ports | array of strings | ports to be exposed from container (-p option to docker) |
| project-mount | string | path to mount project in container |
| setup-post | string | script to run in container after unpacking tarballs |
| setup-pre | string | script to run in container before running tarballs |
| setup-prep | string | script to run locally before the other setup |
| shell | string | shell to use in the container |
| tarballs | table of strings | mapping from tarball location to install location |
| tmpfs | array of strings | paths in the container where tmpfs should be mounted |
| volumes | array of strings | list of volume mounts (-v option to docker) |
| working-dir | string | method to use to set working dir (default: "match") |

**cap-add** Add additional Linux capabilities to the container. The list of possible values can be found in the docker run reference (https://docs.docker.com/engine/reference/run/#runtime-privilege-and-linux-capabilities).

For example to add the ability to use ptrace inside the container:

```
cap-add = ["SYS_PTRACE"]
```

**cap-drop** Drop normally enabled Linux capabilities from the container. The list of possible values can be found in the docker run reference (https://docs.docker.com/engine/reference/run/#runtime-privilege-and-linux-capabilities).

For example to drop the ability to bind to privileged ports:

```
cap-drop = ["NET_BIND_SERVICE."]
```

**copy-groups** TODO . . .

**env** Specify environment variables to be set in the container. This consists of a table, where the keys are the variable names and the values are the variable values. For example to set SOME_VARIABLE to "some value":

```
[crates.demo.env]
    "SOME_VARIABLE" = "some value"
```

## 4.3 Local Configuration

In addition to the shared project configuration each user can have a local configuration. This configuration allows changing the Docker URL, and adding extra steps to the container setup.

On Linux this file can be found at "`$XDG_CONFIG_HOME/wharfrat/config.toml`". If `$XDG_CONFIG_HOME` is not set, then the default path is "`$HOME/.config`", so the default location for the config is "`$HOME/.config/wharfrat/config.toml`".

```
docker-url = "file:///var/run/docker.sock"
auto-clean = true

[[setups]]
    project = ".*/test"
    setup-prep = """
        echo "LOCAL PREP: $*"
        pwd
    """

    setup-pre = """
        echo "LOCAL PRE"
        pwd
    """

    setup-post = """
        echo "LOCAL POST"
    """

    [setups.tarballs]
        "path/to/tarball.tgz" = "/path/in/container/to/unpack"

    [setups.env]
        "LOCAL_CRATE_ENV" = "true"

[[setups]]
    setup-prep = """
        echo "LOCAL PREP: $*"
        pwd
    """

    setup-pre = """
        echo "LOCAL PRE"
        pwd
    """

    setup-post = """
        echo "LOCAL POST"
    """

    [setups.env]
        "LOCAL_CRATE_ENV" = "true"
```

The available settings are:

| | | |
|---|---|---|
| docker-url | The URL to use to connect to Docker | |
| auto-clean | If set to true, then wharfrat run will automatically replace containers that were built from old config, or the wrong image. | |
| setups | project | a regular expression that much match the project path for this setup to be applies. If not specified, then ".*" is used. |
| | crate | a regular expression that must match the crate name for this setup to be applied. If not specified, then ".*" is used. |
| | setup-prep | script to run locally before doing anything else |
| | setup-pre | script to run remotely before unpacking tarballs |
| | setup-post | script to run remotely after unpacking tarballs |
| | tarballs | a table to tarballs to be unpacked into the container, mapping tarball path to target path in the container |
| | env | a table of environment variables to set in the container, mapping name to value |

CHAPTER 5

# Accessing Containers

**Table of Contents**


- *Accessing Containers*
  - *wharfrat run*
  - *Exposing Commands*


## 5.1 wharfrat run

## 5.2 Exposing Commands

It is possible to use the `wr-exec` command to expose commands from inside a container to the host. This is normally done by creating an executable config file with a `#!` line that invokes `wr-exec`. For example, if `./test` contains:

```
#!/usr/bin/env wr-exec

project = "/path/to/project/file"
command = ["command", "arg1"]
```

Then, running `./test arg2` will run the command `command arg1 arg2` in the container for the default crate defined in the wharfrat project file at `/path/to/project/file`.

CHAPTER 6

# Environments

**Table of Contents**


- *Environments*
  - *Purpose*
  - *Creating an Environment*
    - *Exporting Binaries*
    - *Creating the Environment Files*
  - *Using an Environment*


## 6.1 Purpose

Environments are a convienient way to access a set of executables from the host system as if they are on the host.

They are based on approach taken by tools such as virtualenv, where the path is modified to inject executables from the environment earlier into the search path.

## 6.2 Creating an Environment

There are two parts to creating an environment:

- Setting up one or more crates to export binaries
- Creating the actual environment files

The first part is typically done once as part of the project configuration, and checked in. The second is done by each person that wishes to use an environment with that project.

### 6.2.1 Exporting Binaries

To export binaries from a crate into an environment the `export-bin` crate setting. This setting consists of a list of patterns, and any executable that matches any of these patterns will be exported into the environment.

### 6.2.2 Creating the Environment Files

The files for the environment are created using the `wharfrat env create` command. This command takes the path at which to create the environment. You can also optionally select a subset of the crates in the project that should be used to create the environment (the default being all crates in the project).

## 6.3 Using an Environment

Once an environment has been created, then is must be activated to be used. Activating the environment will change the shell's `PATH` such that the exported binaries are first in the search path.

In addition to the binaries exported from the crates, the environment includes a cached copy of the wharfrat binary that was used to create the environment.

Once an environment has been activated, any new exported binaries will be noted against the commands that caused them to appear. This allows the commands to be re-run when the create has to be recreated (which is done automatically when running an exported command by default).

```
~/proj# wharfrat env create wrenv
~/proj# . wrenv/bin/activate
(wr: wrenv) ~/proj#
```

CHAPTER 7

# Examples

...

CHAPTER 8

# Indices and tables

- genindex
- modindex
- search