

Podman Network Setup

How does rootless networking work?

- Podman 5.X: `network` as default network driver
 - Behaves like application is running on the host
 - Host IP, correct source IP, but still isolated
 - Port forwarding is required compared to host networking
- Previous default: `bridge` driver
 - Still default when connecting containers via shared network
 - NAT is used -> Source IP is always the NAT IP
 - IPv6: Interesting... Sometimes broken

Q: How can two pods access each other over the network?

- Pod: Share network namespace, reach each other via localhost
- Bridge: Connect containers to bridged network, both gain a NATed IP

Podman: Connecting Services

The following script spawn a Postgres client and server and connects them through a custom network

```
1 #!/bin/sh
2 set -x
3 net=test; client=pgclient; server=pgserver
4 pg_password='p4$$w0rd'; pg_version=18-alpine
5 # Create network and launch detached server
6 podman network create ${net}
7 podman run --rm --name ${server} --network ${net} -e POSTGRES_PASSWORD=${pg_password} \
8   -d postgres:${pg_version}
9 # Wait some time to avoid race condition
10 sleep 3
11 # Run the client (note different password variable)
12 podman run --rm --name ${client} --network ${net} -e PGPASSWORD=${pg_password} \
13   --entrypoint psql postgres:${pg_version} -h ${server} -U postgres -c "SELECT version();"
14 # clean up
15 podman rm -f ${client} ${server} && podman network rm -f ${net}
```

----- [finished] -----

Podman: Connecting Services

```
+ net=test
+ client=pgclient
+ server=pgserver
+ pg_password='p4$$w0rd'
+ pg_version=18-alpine
+ podman network create test
test
+ podman run --rm --name pgserver --network test -e 'POSTGRES_PASSWORD=p4$$w0rd' -d
postgres:18-alpine
18bc73b71cf1ce159240af477181a67e9e2d7d39c664b57ea6deb87798e10b4b
+ sleep 3
+ podman run --rm --name pgclient --network test -e 'PGPASSWORD=p4$$w0rd' --entrypoint psql
postgres:18-alpine -h pgserver -U postgres -c 'SELECT version();'
                version
-----
 PostgreSQL 18.2 on x86_64-pc-linux-musl, compiled by gcc (Alpine 15.2.0) 15.2.0, 64-bit
(1 row)

+ podman rm -f pgclient pgserver
+ podman network rm -f test
test
```

Podman: Connecting Services

Shell Script Approach

Do you like the manual setup?

- Easy setup of two Postgres instances
 - Can communicate over network
 - Includes DNS resolution -> Nice!

But:

- Verbose syntax
- Setup is error-prone
 - Lots of repeated custom logic
- What if one of the commands fails?

Goal

- Meta file containing setup instructions
 - More declarative than shell script
 - Better error handling
 - Cross platform
- Lifecycle handling?
 - Autostart?
 - Restart after crash?
 - Automatic updates??

Declarative Definition

We can use `docker-compose` and Quadlets for declarative definition

This can be seen as a subset of Infrastructure as Code

	docker-compose	Quadlet
Requirements	Docker/Podman + Compose Script	Podman + systemd
Lifecycle Manager	Docker Daemon	systemd
Recommended for	Development (project-level)	Servers (host-level)
File format	Custom (compose)	systemd / Kubernetes

Docker vs Podman Compose

There are two major implementation of the `compose spec` (<https://compose-spec.io>)

Docker Compose

- Reference implementation (Docker 1st party)
 - Always supports full feature set
- Utilizes Docker socket ("REST" API)
- Allows parallel operation (image pulls etc.)

- v1: Standalone Script `docker-compose`
- v2: Plugin for Docker `docker compose`

Podman Compose

- Community implementation for Podman
- Translates compose file into `podman` CLI commands
- Sequential operation (last time I checked)

- Standalone Script `podman-compose`

Docker Compose + Podman

We want to use Docker Compose with Podman

- v2 can still be used as standalone script
 - Many distros only ship v1.X though...
 - Ensure your distro ships `docker-compose` v2.X or install directly from GitHub Releases
 - `docker-compose --version`
- We need to enable the Podman Socket for Docker API support
 - `systemctl --user enable --now podman.socket`
- We need to tell `docker-compose` the path of our API socket!
 - `export DOCKER_HOST=unix://${XDG_RUNTIME_DIR}/podman/podman.sock`
 - Add this to your `~/.bashrc` or similar

All slides in this lecture assume this `docker-compose` + `podman` setup. If you use another setup, you need to replace the `docker-compose` commands (officially legacy) with `docker compose` or `podman-compose`.

Compose: Simple Example

Compose Spec

- Simple YAML file (Minimal Syntax)
- All container under `services`
 - `web`: Implicit container name
- Declarative definition of CLI parameters

```
1 services:
2   web:
3     image: docker.io/nginxinc/nginx-unprivileged
4   ports:
5     - 8080:8080/tcp
6     volumes: # serve generated lecture slides
7     - ../../_site:/usr/share/nginx/html:ro,Z
```

- Launching: `docker-compose -f examples/compose/compose.simple.yml up`
 - Relative paths are read from compose file
 - `compose.yml` and `docker-compose.yml` (and `yml`) are automatically detected --> no `-f`

Compose: Subcommands

Command	Action
<code>config</code>	Parse and show the final config
<code>pull</code>	Pull service images
<code>build</code>	Build or rebuild services if required
<code>create</code>	Creates the containers & networks
<code>start</code>	Start services (must be created first)
<code>restart</code>	Restart services
<code>up</code>	Shortcut for create and <code>start</code>
<code>stop</code>	Shutdown all services (keep containers)
<code>down</code>	Stop and remove containers & networks
<code>logs</code>	View output from containers
<code>exec</code>	Execute command in running container
<code>run</code>	Execute command in a new container
<code>watch</code>	Watch service, refresh when files change

Subcommands can be executed for all containers (default) or for a specific container

- `docker-compose [up/down/logs] myservice`
 - Use `up -d` to detach the logs
 - Use `logs -f` to stream updated logs
 - Use `up --build` for a `build` shortcut
 - Use `run --rm` to prevent zombies

Compose: More Syntax

```
1 services:
2   app:
3     build:
4       context: .
5       dockerfile: Containerfile
6     container_name: demo-app
7     command: ["--http", ":8080"]
8     entrypoint: ["/bin/myapp"]
9     user: "1000:1000"
10    working_dir: /app
11    env_file: .env
12    environment:
13      LECTURE: "SE2"
14    ports:
15      - 127.0.0.1:8080:8080/tcp
```

We can specify any container option in the compose file

- Build the service `app` from local `Containerfile` in directory `.`
- When using `Dockerfile` in same folder, `build` is sufficient
 - Remember: No automatic rebuilds!

Naming

- `COMPOSE_PROJECT_NAME`: Parent folder of the compose file, e.g. `compose`
 - Can be overridden (environment or top-level `name`.)
- `container_name` sets explicit name `demo-app` instead of `compose-app-1`

Compose: Variables

```
1 services:
2   app:
3     image: alpine:${TAG:-latest}
4     env_file: .env
5     environment:
6       OTHERPASSWORD: ${MYPASSWORD}
7     command:
8       - sh
9       - -c
10      - env | grep PASSWORD
```

Variables

- Compose reads current environment and parses `.env` file relative to compose file
- Variable must be passed to containers explicitly

Substitution

Syntax	Evaluates To
<code>\$(VAR)</code>	Value of <code>VAR</code>
<code>\${VAR:-default}</code>	Value of <code>VAR</code> , otherwise <code>default</code>
<code>\${VAR:?error}</code>	Require <code>VAR</code> or throw error

Compose: Variables

```
1 services:
2   app:
3     image: alpine:${TAG:-latest}
4     env_file: .env
5     environment:
6       OTHERPASSWORD: ${MYPASSWORD}
7     command:
8       - sh
9       - -c
10      - env | grep PASSWORD
```

```
# ../examples/compose/.env
MYPASSWORD=s3cr3t
```

```
file=../examples/compose/compose.env.yml
docker-compose -f $file config | grep -A2
environment
```

[finished]

```
environment:
  MYPASSWORD: s3cr3t
  OTHERPASSWORD: s3cr3t
```

Compose: Network

Config

```
1 services:
2   proxy:
3     image: nginxinc/nginx-unprivileged:alpine
4     networks: ["web"]
5   backend:
6     image: python:alpine
7     networks: ["web", "db"]
8   db:
9     image: postgres:alpine
10    networks: ["db"]
11 networks:
12   db:
13   web:
14     enable_ipv6: true
15     ipam:
16       config:
17         - subnet: 2001:db8::/64
```

- Compose creates a default project network
 - Each container has implicit default network
- Custom networks can be created under `networks:`
 - Simple name is sufficient
 - Advanced options like subnet and network driver can be specified
- Bridged networks support DNS for container names
 - E.g. `backend` can reach `db` via `db` hostname
 - Ensure Aardvark DNS plugin is installed!

Exercise

Remember the script from the beginning? Your task is to reproduce this using a compose file!

The script is located in the `git` repo under `./examples/compose/manual.sh`

- Spawn two Postgres containers
 - `client` and `server`
 - `server` should keep running (daemon)
 - `client` should only execute the query and exit

Use a `env` file like this to share variables between the two containers:

```
POSTGRES_VERSION=18-alpine
POSTGRES_PASSWORD=s3cr3t
```

Compose: Advanced Network

```
1 services:
2   proxy_host:
3     image: nginxinc/nginx-unprivileged:alpine
4     network_mode: host
5   proxy_old: # Workaround for Podman 4.X
6     image: nginxinc/nginx-unprivileged:alpine
7     network_mode: "
slirp4netns:port_handler=slirp4netns"
8   proxy_new: # Workaround for Podman 5.X
9     image: nginxinc/nginx-unprivileged:alpine
10    network_mode: "pasta"
```

As mentioned, source IP handling can be difficult with rootless networking

- Other network modes than `bridge` can be used
 - No support for DNS resolution!
- To get source IPs with bridge network:
 - Look into **socket activation** (https://systemd.io/DAEMON_SOCKET_ACTIVATION)

Compose: Sidecars

In K8s (or Pods in general) there is the "Sidecar" pattern:

- Shared network namespace between containers
- Can reach each other via localhost
- Problem: No native Docker support for Pods
 - We can still share the network
 - Ports must be exposed from main service
 - Funny issues when starting/deleting containers

```
1 services:
2   app:
3     image: nginx:alpine
4     ports:
5       - 8001:80 # app port
6       - 8080:8080 # sidecar port
7   sidecar:
8     image: nginxinc/nginx-unprivileged
9     entrypoint:
10      - sh
11      - -c
12      - sleep 1 && nginx -g "daemon off;"
13     network_mode: "service:app"
```

Compose: Healthcheck & Dependencies

```
1 services:
2   app:
3     image: nginx:alpine
4     healthcheck:
5       test: ["CMD", "curl", "[::1]"]
6       interval: 3s
7       retries: 5
8       start_period: 1s
9       timeout: 10s
10  dependend:
11    image: nginx:alpine
12    depends_on:
13      app:
14        condition: service_healthy
```

- We can define a startup order between services
 - `depends_on` with multiple conditions
 - `service_running`: Wait for service to start
 - `service_healthy`: Wait for health check
- Healthchecks...
 - can be specified on image level
 - can be overridden by compose
 - run inside the container
 - use `CMD` or `CMD-SHELL`
 - array or string notation

Compose: Watch

```
1 services:
2   app:
3     build: ./watch
4     entrypoint: ["sleep", "9000"]
5     develop:
6       watch:
7         - action: rebuild
8           path: ./watch/requirements.txt
9
```

Auto update builds

- For development, we can watch and reload project
- Can be done for volume mounts (hot-reload)
 - `action: sync`
- Also for full container builds
 - `action: build`
- Use `docker-compose up -w` to watch

Compose: Anchors

```
1 x-common-env: &common-env
2   TZ: UTC
3   LOG_LEVEL: info
4
5 services:
6   web:
7     environment:
8       <<: *common-env
9       HTTP_PORT: 8080
10    volumes: &app-volumes
11      ./data:/data:ro,z
12   worker:
13     environment: *common-env
14     volumes: *app-volumes
```

- YAML anchors can be used to reuse config
- Can be placed anywhere in the file
 - Use top-level `x-` attributes for config snippets (no invalid syntax errors)
- Advantage: Ensure consistent config
- Disadvantage: Harder to read

Compose: Profiles

```
1 services:
2   app:
3     image: python:alpine
4     profiles: ["app", "web"]
5   web:
6     image: nginx:alpine
7     profiles: ["web"]
8   db:
9     image: postgres:18-alpine
10    profiles: ["app", "db"]
```

■ Separating larger applications

- Profiles can be used to divide complex setups
- Each service can be added to multiple profiles
- Use `docker-compose --profile db up` to spawn db
 - Use multiple `--profile` flags for `web` + `app`
 - Or: Use `COMPOSE_PROFILES=web,app` environment

Compose: File based separation

Include

- The `include` directive can be used to import content from other compose files
 - Simple, declarative way if files get too large

```
1 include:
2   - ./compose.simple.yml
3 services:
4   db:
5     image: postgres:alpine
```

Multiple files via CLI

- We can also specify multiple files via CLI
- Useful if we use tooling around compose
 - E.g. conditionally include files
- `docker-compose -f file1 -f file2`
- Each file overrides contents of previous files
 - Objects (lists, dicts) get merged

```
file=../examples/compose/compose.include.yml
docker-compose -f $file config --services
```

----- [finished] -----

```
db
web
```

Thank you for your attention!

Don't forget the feedback