

---

# **SAMPL Docker Documentation**

**Mike Henry**

**Oct 19, 2022**



**CONTENTS:**

<b>1</b>	<b>Docker Overview</b>	<b>3</b>
<b>2</b>	<b>Conda</b>	<b>5</b>
<b>3</b>	<b>Admin Tasks</b>	<b>7</b>
3.1	General overview . . . . .	7
<b>4</b>	<b>Indices and tables</b>	<b>9</b>



**Warning:** This documentation is in active development and is subject to change.



## DOCKER OVERVIEW

See the [docker](#) documentation for more details.





## CONDA

---

**Note:** A familiarity of conda is assumed

---

Example docker file

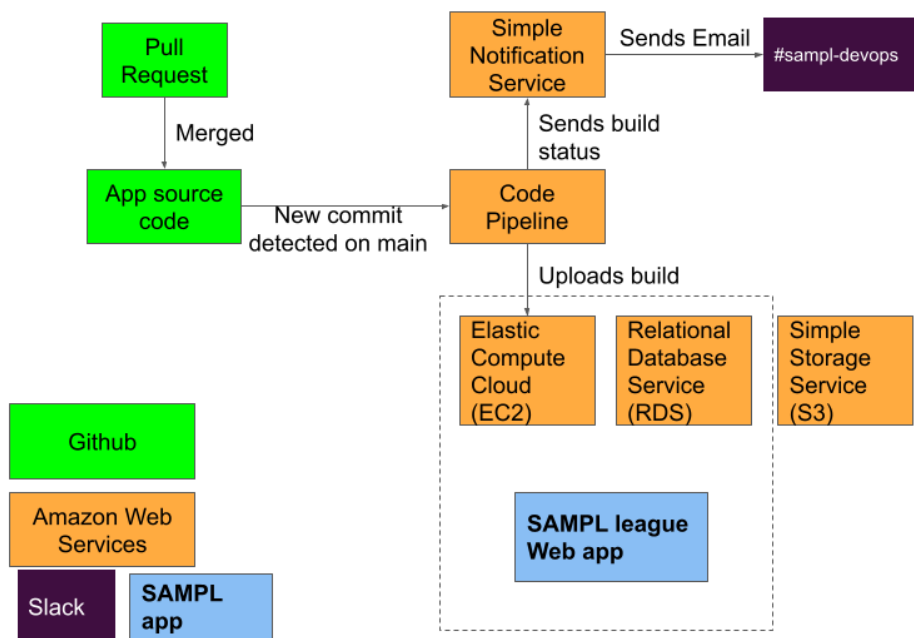
```
FROM continuumio/miniconda3:4.9.2-alpine
WORKDIR /opt/app
COPY environment.yml ./
RUN conda env create -f environment.yml && \
    conda clean --all --yes
ENTRYPOINT ["conda", "run", "-n", "test"]
```

The ENTRYPOINT directive will ensure that commands passed into the container with `run` will use the correct conda environment.



## ADMIN TASKS

### 3.1 General overview



When a commit is made to the main branch (e.g. merging in a pull request that targets the main branch) a build is triggered on **Code Pipeline**. Then a notification is sent (using **SNS**) to the **#saml-devops** channel in the Mobley Lab **Slack workspace**. Once the build completes, the code is sent to the **Elastic Beanstalk** instance and the web service restarted.

#### 3.1.1 Amazon Web Services infrastructure (AWS)

All services are hosted in US East (Ohio) us-east-2.

- Code Deployment: **Code Pipeline**
- Container Registry: **ECR**
- DNS: **Route 53**
- Database: **RDS**

- E-mail: [SES](#)
- Elastic Beanstalk: [Elastic Beanstalk](#)
- Monitoring & Alarms: [Cloud Watch](#)
- Slack E-mail messages: [SNS](#)
- Storage: [S3](#)
- Web Server: [EC2](#)

### 3.1.2 Update ever\_given

You will need an account on [PyPI](#). Starting from the root of the repository (after making changes to the code): ever\_given version number must be incremented in setup.cfg for a new PyPI upload .. code-block:: bash

```
$ cd ever_given $ python -m pip install --upgrade build $ python -m build $ python3 -m twine upload dist/*
```

The new package should show up [here](#).

### 3.1.3 Spin up worker

Generally:

1. Provision an EC2 instance
2. SSH onto instance and clone codebase
3. Use *pipenv* to setup environment
4. Ensure environmental variables are set
5. Launch as many *dask-workers* as needed

### 3.1.4 Expand Local Storage On EC2 Instance

First, follow the guide “[Modify an EBS volume using Elastic Volumes](#)”. Then follow the steps below (adapted from “[Extend a Linux file system after resizing a volume](#)”).

```
$ df -hT # Check disk space, mount points, and file system type
$ lsblk # Check which device needs to be expanded
$ sudo growpart /dev/nvme0n1 1
$ df -hT # Confirm that nothing has changed on the file level
$ lsblk # Check to see that partition has grown
$ sudo xfs_growfs -d / # Be sure to use the correct tool to grow the file system
$ df -hT # Confirm that everything looks correct
```

While doing a backup first is not a bad idea, since the database is part of another service and S3 is used for storage, if a major crash occurs during this operation, the web server can be restarted and things *should* be fine.

## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`