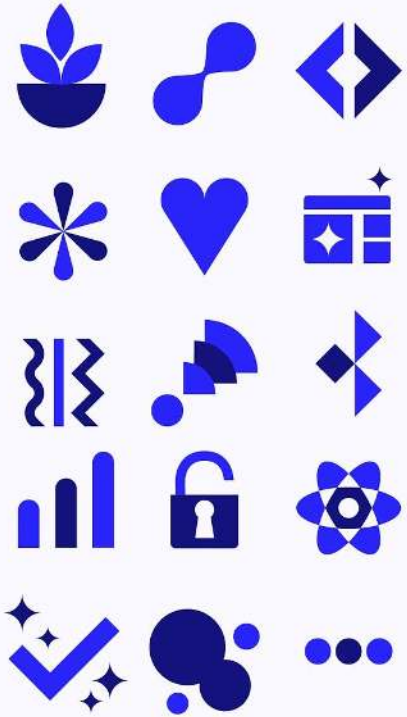

Production Docker Image for Apache Airflow

Polidea

Airflow Summit 2020 - 14.07.2020

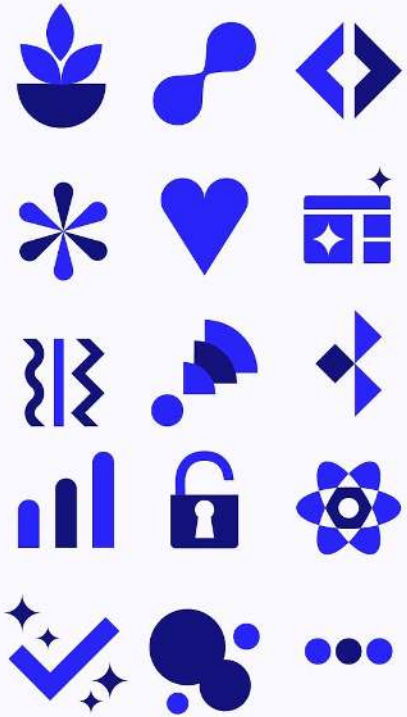

Production Container Image for Apache Airflow

Polidea

Airflow Summit 2020 - 14.07.2020

Hi!

Jarek Potiuk

Apache Airflow:
PMC Member and Committer

Polidea:
Principal Software Engineer
(ex-CTO)

Airflow Summit:
Co-Organizer: Content (Lead)

@higrys

Intro

What questions will be answered?

- Context
  - What container images are and why there are important ?
- Status
  - How it looked like so far ?
  - How it is going to look like now ?
- Internals
  - What is in the image?
  - How we test the image?
- Usage
  - How to extend Airflow Image?
  - How to customize Airflow Image?
  - How you can use the Image?
- Future
  - What’s next?

What this talk is NOT about?

- Basic container image knowledge
  - https://docker-curriculum.com/
- Details of CI container image of Airflow
  - https://github.com/apache/airflow/blob/master/IMAGES.rst
- Details of how Kubernetes Airflow integrate
  - “Airflow on Kubernetes” by Michael Hewitt
    https://www.crowdcast.io/e/airflowsummit/6
- Details on deploying Airflow with the image

Who is the talk for?

- You want to deploy Airflow using container images
- You want to contribute to Airflow in Devops area
- You want to learn about best practices of using Airflow Containers
- You are a curious person that want to learn something new

Container Images

Context

What is a container ?

- **Standard** unit of software.
    - OCI: https://opencontainers.org/
- Packages code and its dependencies
- Lightweight execution package of software
- Container images - binary packages

Container

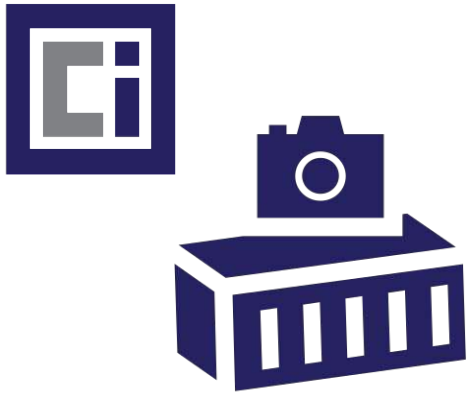

Container image

Container ≠ Docker

- Docker is a command line tool
  - Building, Running, Sharing containers
- Docker Engine runs containers
- Alternatives: **rkt, containerd, runc, podman, lxc, …**
- DockerHub.com is popular container registry
- Alternatives: GitHub, GCR, ECR, ACR

Container management CLI

Container execution engine

Container registry

Context: What is Container file

- Specify base image
- Run commands
- Copy files
- Set working directory
- Define entrypoint
- Define default command

```
FROM ubuntu:18.04
COPY . /app
RUN make /app && make install
WORKDIR /bin/project
ENTRYPOINT ["/bin/project"]
CMD ["--help"]
```

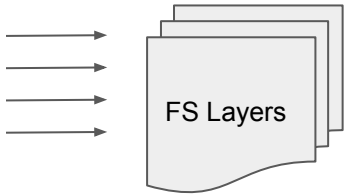

Context: Container Lifecycle: Build

Container execution engine

Build

Container Image file (Dockerfile)

Container image

Container registry

Context: Container Lifecycle: Run

Container execution engine

Run

Container Image file (Dockerfile)

Container image

Container registry

Context: Container Lifecycle: Push

Container execution engine

Push

Container Image file (Dockerfile)

Container image

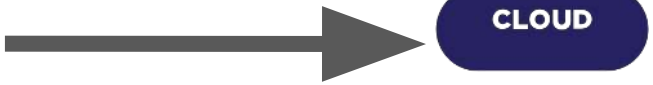

Container registry

Context: Container Lifecycle: Pull

Container execution engine

Pull

Container Image file (Dockerfile)

Container image

Container registry

Why containers are important?

- Predictable, consistent development & test environment
- Predictable, consistent execution environment
- Lightweight but isolated: sandboxed view of the OS isolated from others
- Build once: run anywhere
- Kubernetes runs containers natively
- Bridge: “Development -> Operations”

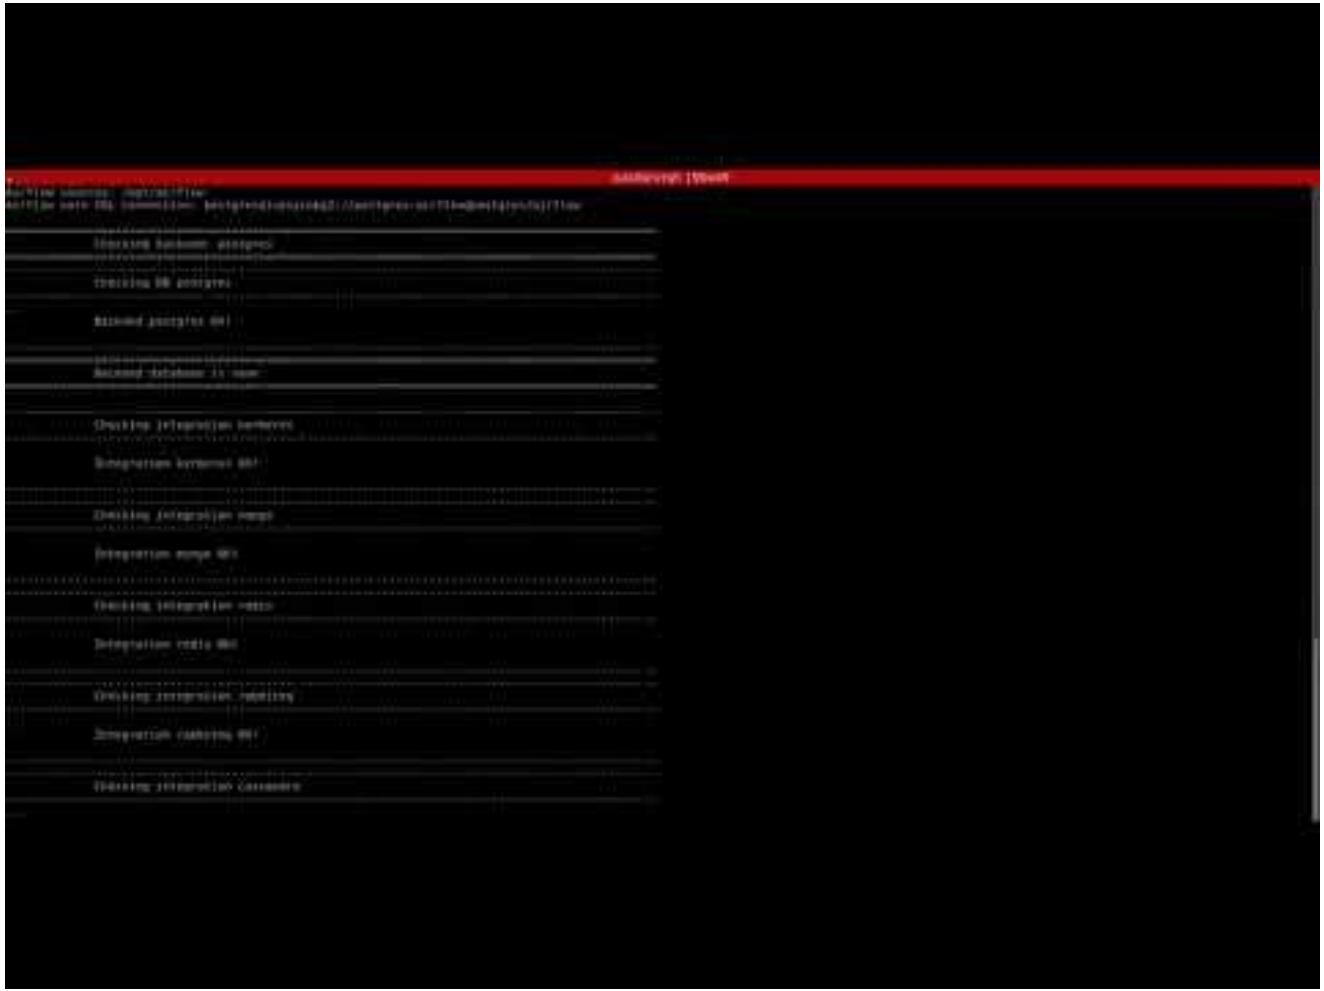

Container Images

Status

History of Containers in Airflow: CI

- Used for CI for > 2 years: Gerardo Curiel
- Optimized and incorporated by Breeze 1.5 years ago or so
- Docker Compose as execution engine
- Slimmed down recently (Thanks Ash!)
- Optimized for development use

History of Containers in Airflow: Prod

- Puckel image created by Matthieu "Puckel_" Roisil (Thanks Matthieu!)
  - Used by many users in production
  - Used by the publicly available Helm Chart (not managed by community)
- Official Production Image (managed by community)
  - Alpha Quality community image in 1.10.10
  - Beta Quality community image in 1.10.11 (now!)

State of the Official Production image

- Beta Quality - usable for production
- Most important feedback incorporated
- Already used in production
- Public Helm Chart switched to the Official Production Image
- Community Helm Chart (donated by Astronomer!) uses it for testing
- Stable version in v1-10-stable, development in master

Container Images Internals

Internals: DockerHub releases

Released image

- ~ 210 MB compressed size
- Python: 2.7, 3.5, 3.6, 3.7, 3.8
- 1.10.11 = Python 3.6
- manually released
- using “1.10.11” tag
- latest = 1.10.11
- docker pull apache/airflow

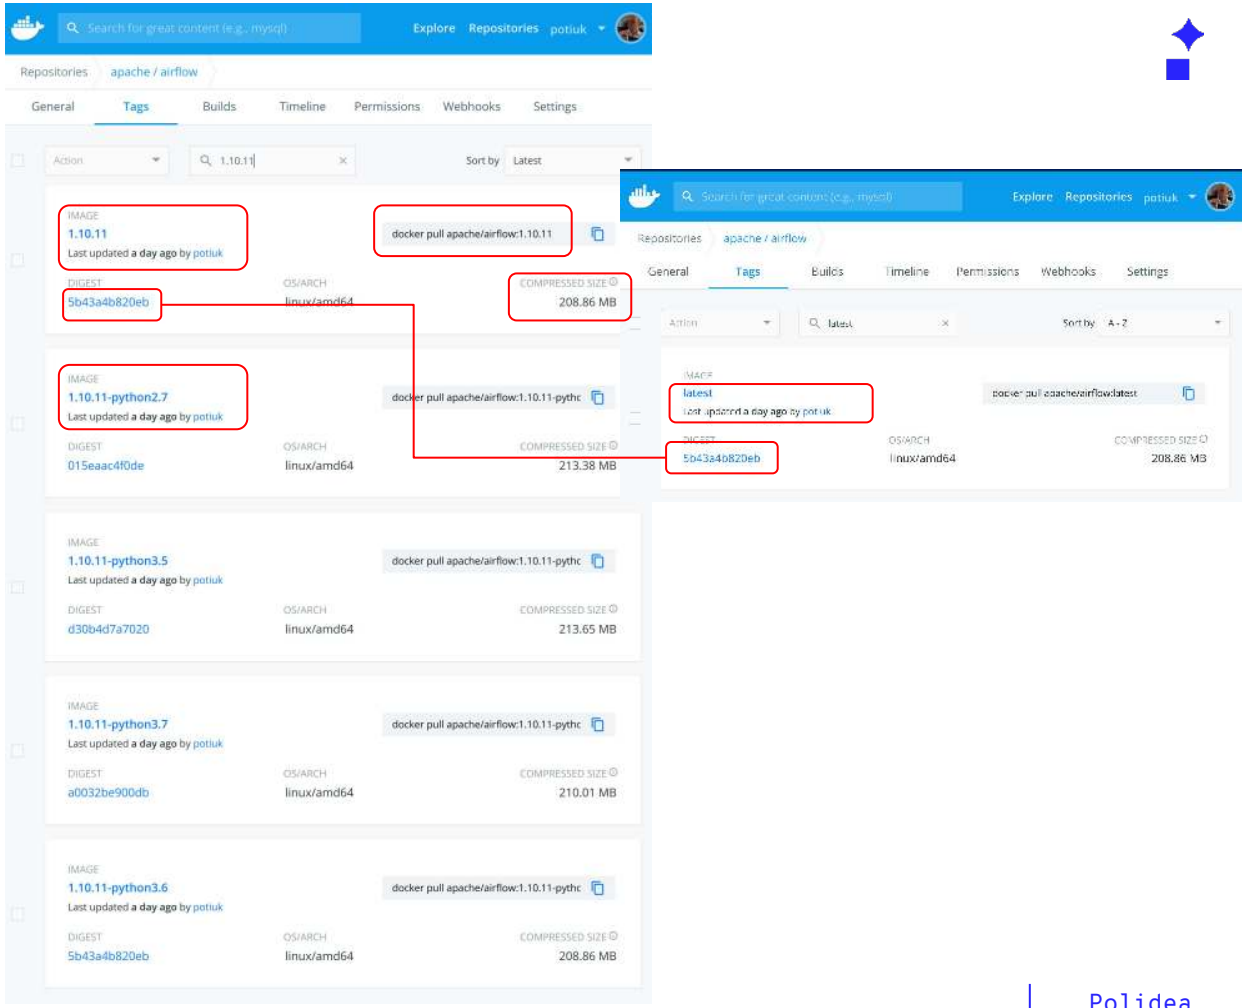

Container Image or Container File ?

- Apache Software Foundation releases sources, not binaries
- Binaries can only be released for convenience of users
- Binaries must be rebuildable from released sources (PyPI, for example)
- Users should be able to build the software they need
- Should we release Container Image, Container File, or both?

Features of the production image

- Optimised for size (Compressed: ~230MB, ~800 MB on disk)
- Python 3.6, 3.7, 3.8 (2.0 and 1.10.*) , 2.7, 3.5 (1.10.*)
- Extras installed:
  - async,aws,azure,celery,dask,elasticsearch,gcp,kubernetes, mysql,postgres,redis,slack,ssh,statsd,virtualenv
- OpenShift compatible (dynamic uid allocation)
- Gunicorn using shared memory (optimised parallelism)

Features of the production image file

- Builds optimised image
- Highly customizable (ARGs)
- Multi segmented (build + main)

Internals: build image

Build image

- Pass arguments
- Define variables
- Install apt dependencies (with dev ones)
- Install airflow (sources, pip, github): `--user`
- Include constraints
- Transpile website (yarn)
- ~700 MB compressed, ~2GB on disk
- Root user


```
ARG AIRFLOW_VERSION="2.0.0.dev0"
ARG AIRFLOW_EXTRAS="async,aws,azure,celery,dask,elasticsearch,gcp,kubernetes,mysql,postgres,redis,slack
# ...
#######################################################################################################
# This is the build image where we build all dependencies
#######################################################################################################
FROM ${PYTHON_BASE_IMAGE} as airflow-build-image
ARG PYTHON_BASE_IMAGE
ENV PYTHON_BASE_IMAGE=${PYTHON_BASE_IMAGE}
# ...
RUN apt-get update \
    && apt-get install -y --no-install-recommends \
           apt-transport-https \
           apt-utils \
           build-essential \
           freetds-bin \
           freetds-dev \
# ...
    && apt-get autoremove -yqq --purge \
    && apt-get clean \
    && rm -rf "/var/lib/apt/lists/*"
#
RUN pip install --user \
    "https://github.com/${AIRFLOW_REPO}/archive/${AIRFLOW_BRANCH}.tar.gz#egg=apache-airflow[${AIRFLOW_EX
        --constraint \
    "https://raw.githubusercontent.com/${AIRFLOW_REPO}/${AIRFLOW_BRANCH}/requirements/requirements-pytho
    && pip uninstall --yes apache-airflow;

ARG CONSTRAINT_REQUIREMENTS="requirements/requirements-python${PYTHON_MAJOR_MINOR_VERSION}.txt"
ENV CONSTRAINT_REQUIREMENTS=${CONSTRAINT_REQUIREMENTS}

WORKDIR /opt/airflow

# hadolint ignore=DL3020
ADD "${CONSTRAINT_REQUIREMENTS}" /requirements.txt

RUN pip install --user "$(AIRFLOW_INSTALL_SOURCES)[${AIRFLOW_EXTRAS}]${AIRFLOW_INSTALL_VERSION}" \
    --constraint /requirements.txt
# ...

RUN AIRFLOW_SITE_PACKAGE="/root/.local/lib/python${PYTHON_MAJOR_MINOR_VERSION}/site-packages/airflow"; \
    if [[ -f "${AIRFLOW_SITE_PACKAGE}/www_rbac/package.json" ]]; then \
        WWW_DIR="${AIRFLOW_SITE_PACKAGE}/www_rbac"; \
    elif [[ -f "${AIRFLOW_SITE_PACKAGE}/www/package.json" ]]; then \
        WWW_DIR="${AIRFLOW_SITE_PACKAGE}/www"; \
    fi; \
    if [[ ${WWW_DIR:=} != "" ]]; then \
        yarn --cwd "${WWW_DIR}" install --frozen-lockfile --no-cache; \
        yarn --cwd "${WWW_DIR}" run prod; \
        rm -rf "${WWW_DIR}/node_modules"; \
    fi
```

Install to ${HOME}/.local

(side comment)

~ 730 modules
~ 360 MB

Internals: main image

Main image

- Pass arguments/ define variables
- Install apt dependencies (without dev!)
- Add user
- Uses root group (OpenShift)
- Copy(!) Airflow
- Copy DAGs (optionally)
- Copy entrypoint and clean-logs
- Access to /etc/passwd
- Embed dags (for tests)
- Optimized Gunicorn parallelism
- Set working dir
- Exposes port
- Set user
- Entrypoint and command
- ~230 MB compressed, ~800MB on disk


```
#######################################################################################################
# This is the actual Airflow image - much smaller than the build one. We copy
# installed Airflow and all it's dependencies from the build image to make it smaller.
#######################################################################################################
FROM ${PYTHON_BASE_IMAGE} as main
SHELL ["/bin/bash", "-o", "pipefail", "-e", "-u", "-x", "-c"]
# ...
ARG PYTHON_BASE_IMAGE
ENV PYTHON_BASE_IMAGE=${PYTHON_BASE_IMAGE}
# ...
RUN apt-get update \
    && apt-get install -y --no-install-recommends \
           apt-transport-https \
           apt-utils \
           ca-certificates \
           curl \
# ...
RUN addgroup --gid "${AIRFLOW_GID}" "airflow" && \
    adduser --quiet "airflow" --uid "${AIRFLOW_UID}" \
        --gid "${AIRFLOW_GID}" \
        --home "${AIRFLOW_USER_HOME_DIR}"
# ...
COPY --chown=airflow:root --from=airflow-build-image /root/.local "${AIRFLOW_USER_HOME_DIR}/.local"
# ...
COPY scripts/prod/entrypoint_prod.sh /entrypoint
COPY scripts/prod/clean-logs.sh /clean-logs

ARG EMBEDDED_DAGS="empty"
COPY --chown=airflow:root ${EMBEDDED_DAGS}/ ${AIRFLOW_HOME}/dags/

# Make /etc/passwd root-group-writeable so that user can be dynamically added by OpenShift
# See https://github.com/apache/airflow/issues/9248
RUN chmod g=u /etc/passwd

ENV PATH="${AIRFLOW_USER_HOME_DIR}/.local/bin:${PATH}"
ENV GUNICORN_CMD_ARGS="--worker-tmp-dir /dev/shm"

WORKDIR ${AIRFLOW_HOME}

EXPOSE 8080

USER ${AIRFLOW_UID}

ENTRYPOINT ["/usr/bin/dumb-init", "--", "/entrypoint"]
CMD ["--help"]
```

Internals: entrypoint

- Creates user dynamically if missing (OpenShift)
- Fallbacks to sqlite metadata
- Waits until metadata DB is up
- Waits until broker DB is up
- If “bash” or “python” -> runs command
- Else execute airflow command

```
# In case the user is not locally created we automatically create it in /etc/passwd
# This is to handle OpenShift case where random UIDs are used
if ! whoami &> /dev/null; then
  if [[ -w /etc/passwd ]]; then
    echo "${USER_NAME:-default}:x:$(id -u):0:${USER_NAME:-default} user:${AIRFLOW_USER_HOME_DIR}:/sbin/nologin" \
        >> /etc/passwd
  fi
  export HOME="${AIRFLOW_USER_HOME_DIR}"
fi

# if no DB configured - use sqlite db by default
AIRFLOW__CORE__SQL_ALCHEMY_CONN="${AIRFLOW__CORE__SQL_ALCHEMY_CONN:="sqlite:///${AIRFLOW_HOME}/airflow.db"}"

verify_db_connection "${AIRFLOW__CORE__SQL_ALCHEMY_CONN}"

AIRFLOW__CELERY__BROKER_URL=${AIRFLOW__CELERY__BROKER_URL:=}

if [[ -n ${AIRFLOW__CELERY__BROKER_URL} ]] && \
       [[ ${AIRFLOW_COMMAND} =~ ^(scheduler|worker|flower)$ ]]; then
    verify_db_connection "${AIRFLOW__CELERY__BROKER_URL}"
fi

if [[ ${AIRFLOW_COMMAND} == "bash" ]]; then
   shift
   exec "/bin/bash" "${@}"
elif [[ ${AIRFLOW_COMMAND} == "python" ]]; then
   shift
   exec "python" "${@}"
fi

# Run the command
exec airflow "${@}"
```

Internals: .dockerignore

- Ignores everything by default
- You must explicitly include what you want by “!”
- You can further exclude specific subdirectories/patterns
- We generate a lot of stuff in airflow sources
- Sending big context to Docker engine takes time
- You avoid accidental inclusion of unneeded artifacts

```
# NOTE! This docker ignore uses recommended technique
# Where everything is excluded by default and you deliberately
# Add only those directories/files you need. This is very useful
# To make sure that Docker context is always the same on any machine
# So that generated files are not accidentally added to the context
# This allows Docker's `COPY .` to behave in predictable way

# Ignore everything
**

# Allow only these directories
!airflow
!common
!dags
!dev
!docs
!licenses
!metastore_browser
!scripts
!tests

#....

# Now - ignore unnecessary files inside allowed directories
# This goes after the allowed directories

# Git version is dynamically generated
airflow/git_version

# Exclude static www files generated by NPM
airflow/www/static/coverage
airflow/www/static/dist
airflow/www/node_modules
# Exclude static www_rbac files generated by NPM in v1-10-test
airflow/www_rbac/static/coverage
airflow/www_rbac/static/dist
airflow/www_rbac/node_modules

# Exclude link to docs
airflow/www/static/docs

# Exclude python generated files
**/__pycache__/
**/*.py[cod]
**/*$py.class
**/.pytest_cache/
**/env/
```

How we test the image ?

- The image and chart are part of Apache Airflow monorepo
- We build the image with every PR (dependencies)
- We use it in the Kubernetes tests for master (Helm Chart integration)
- We will use released images in the Helm Chart (backward compatibility)
- We will add more tests for various Helm configurations

Container Images

Usage

Usage: Extending Airflow image - use released image

```
docker build . -t yourcompany/airflow:1.10.11-BUILD_ID
```

FILE

```
FROM apache/airflow:1.10.11

# change to root user temporarily
USER root

# Optionally install your own apt dependencies
RUN apt-get update \
    && apt-get install -y --no-install-recommends \
    emacs \
    && apt-get autoremove -yqq --purge \
    && apt-get clean \
    && rm -rf "/var/lib/apt/lists/*"

# Change back to the airflow user
USER airflow

# Add extra dependencies
RUN pip install --user numpy

# Embed DAGs (Optionally) - DAGs can be baked in but also
# they can be git-synced or mounted from shared volume
COPY --chown=airflow:root dags-folder ${AIRFLOW_HOME}/dags/
```

apache/airflow:1.10.11

Container registry

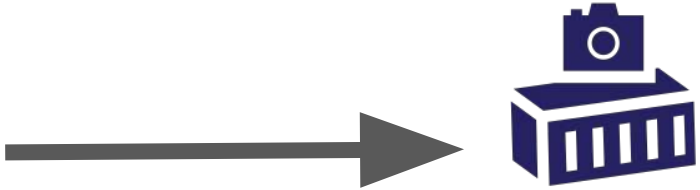

Container image

yourcompany/airflow:1.10.11-BUILD_ID

Extending image - Pros & Cons

Pros

- Use released images
- Simple build command
- Own Dockerfile
- No need for Airflow sources

Cons

- Potentially bigger size
- Predefined extras only
- Installs limited set of python dependencies

Usage: Customising Airflow image - default docker build

```
git clone git@github.com:apache/airflow.git

cd airflow

git checkout v1-10-stable
```

```
docker build .
```

Same as apache/airflow:1.10.11

- Python 3.6
- Default extras
- No additional dependencies

Container image

Usage: Customising Airflow image - use build args

- Installs from PyPi ==1.10.11
- Additional airflow extras, dev, runtime deps …
- Does not use local sources (can be run from master including entrypoint!)

```
docker build . \
  --build-arg PYTHON_BASE_IMAGE="python:3.7-slim-buster" \
  --build-arg PYTHON_MAJOR_MINOR_VERSION=3.7 \
  --build-arg AIRFLOW_INSTALL_SOURCES="apache-airflow" \
  --build-arg AIRFLOW_INSTALL_VERSION="==1.10.11" \
  --build-arg CONSTRAINT_REQUIREMENTS=\
  "https://raw.githubusercontent.com/apache/airflow/1.10.11/requirements/requirements-python3.7.txt" \
  --build-arg AIRFLOW_SOURCES_FROM="empty" \
  --build-arg AIRFLOW_SOURCES_TO="/empty" \
  --build-arg ADDITIONAL_AIRFLOW_EXTRAS="jdbc"
  --build-arg ADDITIONAL_DEV_DEPS="gcc g++"
  --build-arg ADDITIONAL_RUNTIME_DEPS="default-jre-headless"
```

Usage: Image Customization options

- Choose Base image (python)
- Install Airflow from PyPI
- Install from GitHub branch/tag
- Install additional extras
- Install additional python deps
- Install additional apt dev deps
- Install additional apt runtime deps
- Choose different UID/GID
- Choose different AIRFLOW_HOME
- Choose different HOME dir
- Build Cassandra driver concurrently

See IMAGES.rst in the Airflow repo.

e following build arguments (`--build-arg` in docker build command) can be used for production images:

Build argument	Default value	Description
`PYTHON_BASE_IMAGE`	`python:3.6-slim-buster`	Base python image
`PYTHON_MAJOR_MINOR_VERSION`	`3.6`	major/minor version of Python (should match base image)
`AIRFLOW_VERSION`	`2.0.0.dev0`	version of Airflow
`AIRFLOW_ORG`	`apache`	Github organisation from which Airflow is installed (when installed from repo)
`AIRFLOW_REPO`	`apache/airflow`	the repository from which PIP dependencies are pre-installed
`AIRFLOW_BRANCH`	`master`	the branch from which PIP dependencies are pre-installed
`AIRFLOW_GIT_REFERENCE`	`master`	reference (branch or tag) from Github repository from which Airflow is installed (when installed from repo)
`REQUIREMENTS_GIT_REFERENCE`	`master`	reference (branch or tag) from Github repository from which requirements are downloaded for constraints (when installed from repo).
`AIRFLOW_EXTRAS`	(see Dockerfile)	Default extras with which airflow is installed
`ADDITIONAL_AIRFLOW_EXTRAS`		Optional additional extras with which airflow is installed
`ADDITIONAL_PYTHON_DEPS`		Optional python packages to extend the image with some extra dependencies
`ADDITIONAL_DEV_DEPS`		additional apt dev dependencies to install
`ADDITIONAL_RUNTIME_DEPS`		additional apt runtime dependencies to install
`EMBEDDED_DAGS`	`empty`	Folder containing dags embedded into the image in the ${AIRFLOW_HOME}/dags dir
`AIRFLOW_HOME`	`/opt/airflow`	Airflow's HOME (that's where logs and sqlite databases are stored)
`AIRFLOW_UID`	`50000`	Airflow user UID
`AIRFLOW_GID`	`50000`	Airflow group GID. Note that most files created on behalf of airflow user belong to the `root` group (0) to keep OpenShift Guidelines compatibility
`AIRFLOW_USER_HOME_DIR`	`/home/airflow`	Home directory of the Airflow user
`PIP_VERSION`	`19.0.2`	version of PIP to use
`CASS_DRIVER_BUILD_CONCURRENCY`	`8`	Number of processors to use for cassandra PIP install (speeds up installing in case cassandra extra is used).

Usage: It's a Breeze to build images

- Breeze - development and test environment
- Supports building production image
- Auto-complete of options
- New Breeze video showing building production images: https://s.apache.org/airflow-breeze
- `./breeze build-image --help`

```
./breeze build-image --production-image --additional-extras "jira"


./breeze build-image --production-image --python 3.7 \
    --additional-extras "jira"

./breeze build-image --production-image \
    --additional-python-deps "torchio==0.17.10"


./breeze build-image --production-image \
   --additional-dev-deps "libasound2-dev" \
   --additional-runtime-deps "libasound2"


./breeze build-image --production-image \
   --additional-extras "jira" --install-airflow-version="1.10.11"
```

See BREEZE.rst in the Airflow repo

Customising image - Pros & Cons

Pros

- Highly optimized for size
- Build image from sources (security reviews!)
- Can add any extras
- Can add any dependency
- Breeze build commands
- Works from master and 1.10.*

Cons

- Need access to airflow sources
- Complex build command
- Need to understand internals

Why not eat and have cake ?

```
git clone git@github.com:apache/airflow.git

cd airflow

git checkout v1-10-stable
```

```
./breze build-image --production-image --additional-extras "jira" \
  --install-airflow-version "1.10.11"
```

When dependencies change

base-image-for-your-company:1.10.11-2020-07-14

Base Container image

```
FROM base-image-for-your-company:1.10.11-2020-07-14

COPY --chown airflow:root dags-folder "${AIRFLOW_HOME}/dags"
```

When DAGs change

Runtime Container image

How to deploy the images ?

- Docker and Docker-Compose - not recommended for production
- Managed Container Services
  - Managed: Amazon ECS, Google Container on VMs, Azure Container Instances
- Kubernetes on-Prem:
  - Helm Chart
  - Airflow Operator (not recommended yet)
- Managed Kubernetes: Amazon EKS, Google GKE, Azure AKS
- OpenShift (also Kubernetes)

Container Images Future

What is the future for Airflow images?

- It won't change too much !
- Better automated testing via Helm Chart
- Automated releases for 2.0
- ARM support might be the big one. (Apple Mac OS)
- Official Docker Compose
- Smaller features (depends on feedback and expectations):
    - ON BUILD support ?
    - AIRFLOW__CORE__SQL_ALCHEMY_CONN_CMD, AIRFLOW__CELERY__BROKER_URL_CMD support ?
    - Automated user creation ?

Q&A

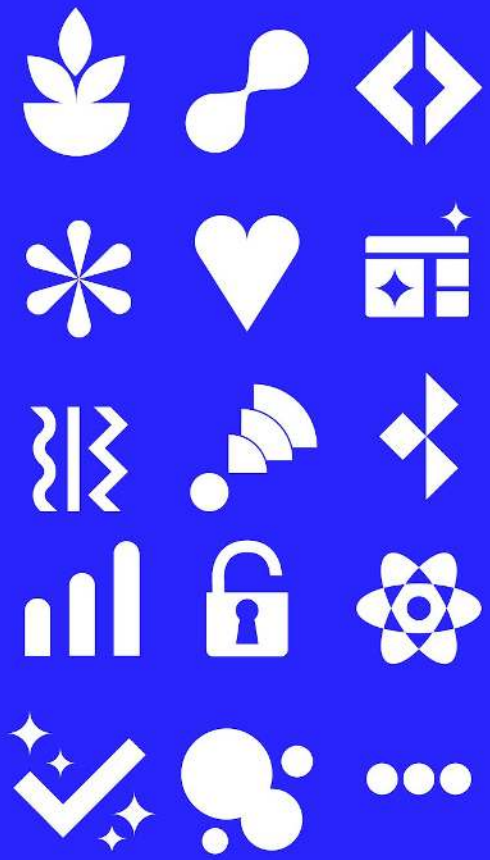

Thanks!

Polidea

hello@polidea.com