

Docker image journey

How to shrink a docker image

DevOps Pro Europe 2019
2019-03-20

MOOVEL BECOMES

REACHNOW



pic: © Moovel



MOOVEL BECOMES

REACHNOW ✓

Me



Engineering Manager

github.com/lotharschulz

[@lothar_schulz](https://twitter.com/lothar_schulz)

lotharschulz.info

speakerdeck.com/lothar



pic: <https://pixabay.com/en/turtle-tortoise-swim-sea-turtle-863336>



pic: <https://www.flickr.com/photos/eyecatcherfotosde/41584801594>

Why reduce docker image size ?

- Pulling and pushing docker images from and to remote docker registries is faster
- Security attack surface is often smaller

Proof?

Well ...

Team Journey part 1

```
FROM golang:1.12
```

```
RUN mkdir /app
```

```
WORKDIR /app
```

```
COPY hello.go .
```

```
RUN go build -o hellogo .
```

```
RUN groupadd -g 99 appuser && useradd -r -u 99 -g appuser appuser
```

```
USER appuser
```

```
CMD ["/hellogo"]
```

```
.....  
$ docker build --rm -t ...  
.....
```

<https://github.com/lotharschulz/hellogodocker/blob/master/Dockerfile>

Team Journey part 1

Image size: 780 MB (compressed 299 MB)

CircleCI build time: ~ 0.5 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 21 sec

<https://circleci.com/gh/lotharschulz/hellogodocker/105>

pull to empty nodes, no layers cached

Team Journey part 1.1

```
FROM golang:1.12
```

```
RUN mkdir /app
```

```
WORKDIR /app
```

```
COPY hello.go .
```

```
RUN go build -o hellogo .
```

```
RUN groupadd -g 99 appuser && useradd -r -u 99 -g appuser appuser
```

```
USER appuser
```

```
CMD ["/hellogo"]
```

```
$ docker build --cache-from golang:1.12 -t ...
```

<https://github.com/lotharschulz/hellogodocker/blob/master/Dockerfile>

Team Journey part 1.1

previous step

Image size compressed: 780 MB (c: 299)

CircleCI build time: ~ 0.5 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 21 sec

current step

Image size compressed: 780 MB (c: 299)

CircleCI build time: ~ 0,6 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 18 sec

<https://circleci.com/gh/lotharschulz/hellogodocker/105>

cache warm up in preparation step

Team Journey part 2 - builder pattern

#part one

FROM golang:1.12 as builder

workdir setup

COPY hello.go .

RUN CGO_ENABLED=0 GOOS=linux go build -a -installsuffix cgo -ldflags '-w -extldflags "-static"' -o hellogodocker .

part two

FROM alpine:latest

user & workdir setup ...

COPY --from=builder /go/src/github.com/lotharschulz/hellogodocker/hellogodocker .

CMD ["/hellogodocker"]

\$ docker build --rm -t ...

<https://github.com/lotharschulz/hellogodocker/blob/master/DockerfileBuilder>

please note the static linked binary

Team Journey part 2 - builder pattern

previous step

Image size compressed: **780 MB**

CircleCI build time: ~ 0.6 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ **18 sec**



current step

Image size compressed: **11.9 MB** (c: 5)

CircleCI build time: ~ 0.6 sec

Dockerhub push time: ~ 3.4 sec

Dockerhub pull time: ~ **1.2 sec**

<https://circleci.com/gh/lotharschulz/hellogodocker/105>

Team Journey part 2 - builder pattern - squash & compress

```
#part one
FROM golang:1.12 as builder
# workdir setup
COPY hello.go .

RUN CGO_ENABLED=0 GOOS=linux go build -a -installsuffix cgo -ldflags '-w -extldflags
"-static"' -o hellogodocker .

# part two
FROM alpine:latest
# user & workdir setup ....
COPY --from=builder /go/src/github.com/lotharschulz/hellogodocker/hellogodocker .

CMD ["/hellogodocker"]
.....
$ docker build --squash/--compress --rm -t ...
.....
```

<https://github.com/lotharschulz/hellogodocker/blob/master/DockerfileBuilder>

squash is an experimental feature

Team Journey part 2 - builder pattern - squash & compress

previous step



current step

Image size compressed: 11.9 MB (c: 5)

CircleCI build time: ~ 0.6 sec

Dockerhub push time: ~ 3.4 sec

Dockerhub pull time: ~ 1.2 sec

Image size compressed: 11.9 MB (c: 5)

CircleCI build time: ~ 0.6/1.1 sec

Dockerhub push time: ~ 3.5/4.1 sec

Dockerhub pull time: ~ 1.1/1.3 sec

<https://circleci.com/gh/lotharschulz/hellogodocker/98>

Team Journey part 3 - alpine base

```
FROM alpine:latest
```

```
# ... set up user
```

```
ADD hellogo /
```

```
CMD ["/hellogo"]
```

```
$ docker build [--squash/--compress] --rm -t ...
```

static linked binary created within ci worker node outside of Dockerfile scope

<https://circleci.com/gh/lotharschulz/hellogodocker/105>

Team Journey part 3 - alpine base

previous step



current step

Image size compressed: 11.9 MB (c: 5)

CircleCI build time: ~ 0.6/1.1 sec

Dockerhub push time: ~ 3.5/4.1 sec

Dockerhub pull time: ~ 1.1/1.3 sec

Image size compressed: 12.2 MB (c: 5)

CircleCI build time: ~ 0.6 sec

Dockerhub push time: ~ 3.5 sec

Dockerhub pull time: ~ 1.2 sec

<https://circleci.com/gh/lotharschulz/hellogodocker/105>

Team Journey part 4 - FROM scratch

FROM scratch

```
ADD ca-certificates.crt /etc/ssl/certs/
```

```
ADD hellogo /
```

```
CMD ["/hellogo"]
```

```
$ docker build [--squash/--compress] --rm -t ...
```

static linked binary created within ci worker node outside of Dockerfile scope

<https://github.com/lotharschulz/hellogodocker/blob/master/DockerfileAlpine>

Team Journey part 4 - FROM scratch

previous step

Image size compressed: 12.2 MB

CircleCI build time: ~ 0.6 sec

Dockerhub push time: ~ 3.5 sec

Dockerhub pull time: ~ 1.2 sec



current step

Image size compressed: **4.8 MB** (c: 2)

CircleCI build time: ~ **15 sec**

Dockerhub push time: ~ 3.5 sec

Dockerhub pull time: ~ 0.7 sec

<https://circleci.com/gh/lotharschulz/hellogodocker/105>

What about Ops?



```
$ docker run -p 1234:1234 --name hellogodocker lotharschulz/hello:build.docker-min--0.2.102
```

```
$ docker exec -it hellogodocker sh
```

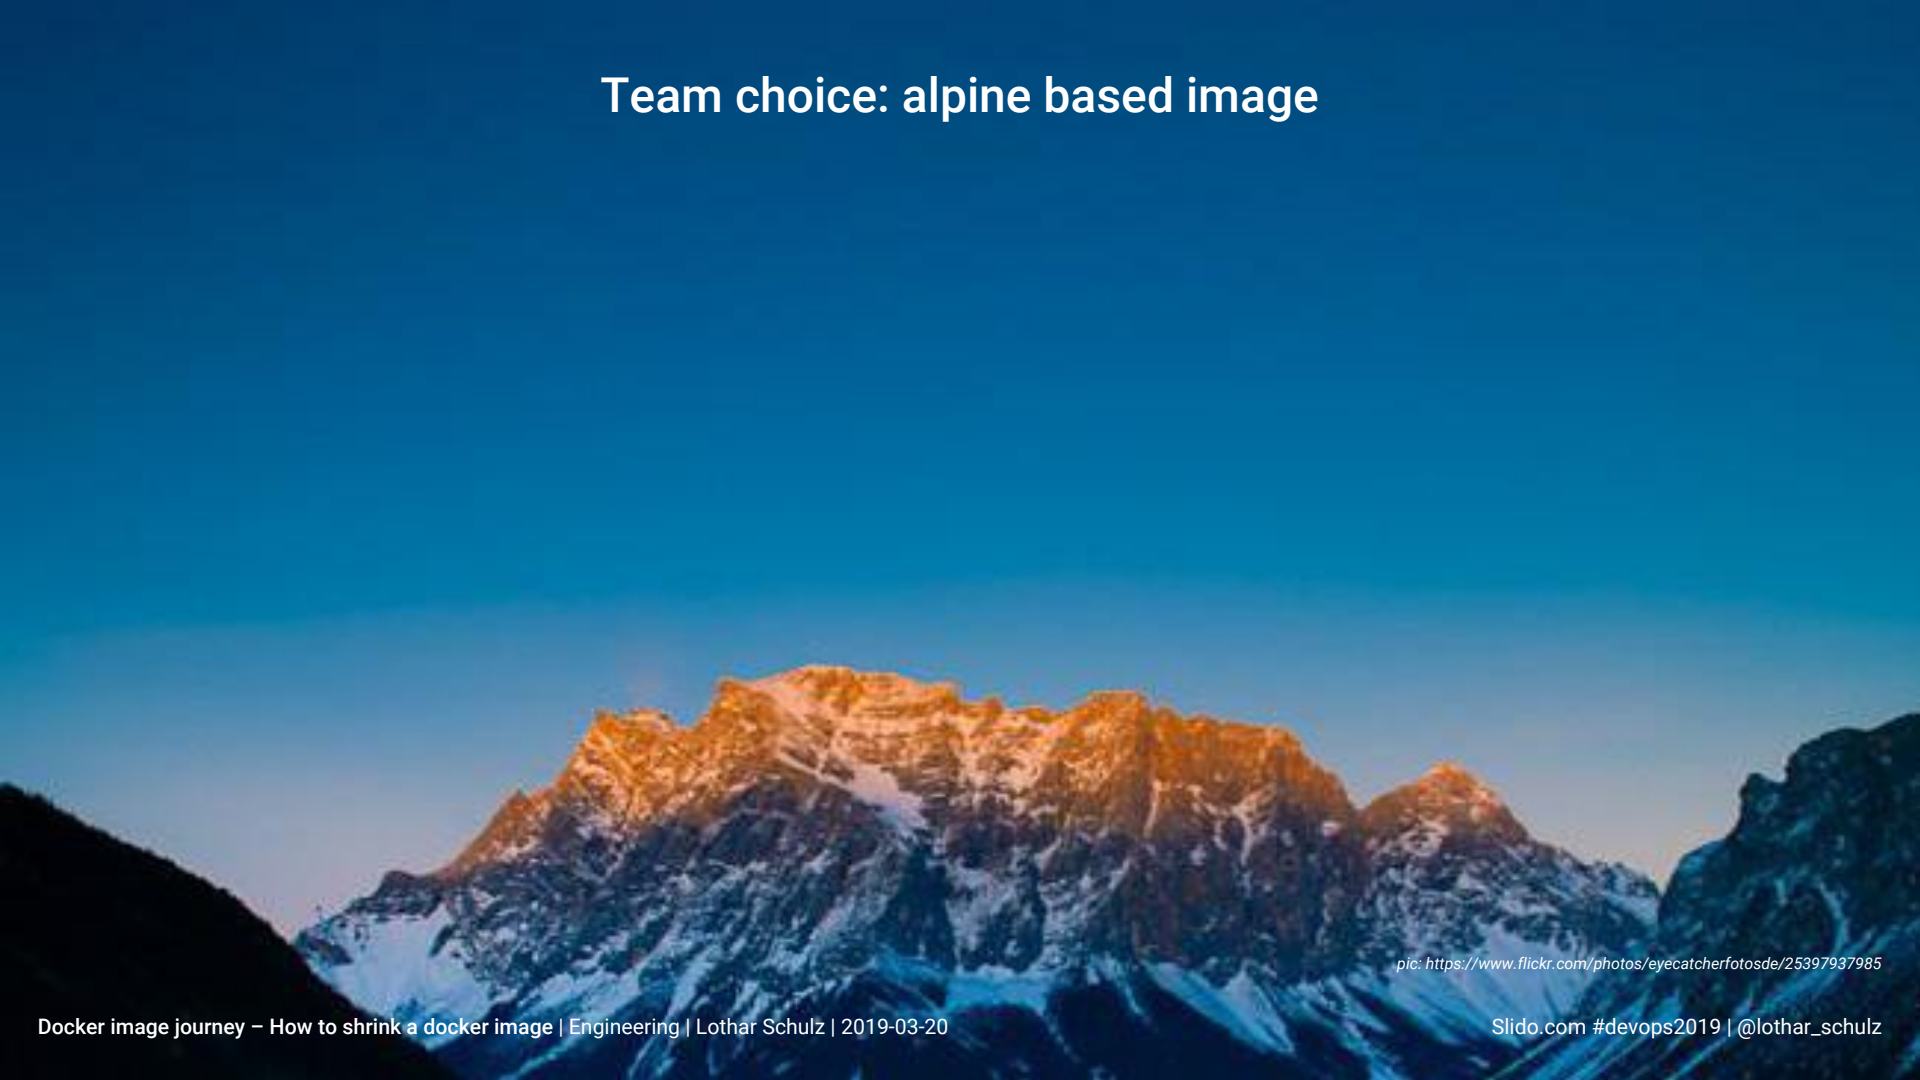
```
OCI runtime exec failed: exec failed: container_linux.go:344:
```

```
starting container process caused "exec: \"sh\":
```

```
executable file not found in $PATH": unknown
```

pic: <https://pixabay.com/en/caution-label-warning-red-mark-943376/>

Team choice: alpine based image



pic: <https://www.flickr.com/photos/eyecatcherfotosde/25397937985>

Some numbers



image size compressed: 780 MB

Dockerhub pull time: ~ 21 sec

0.6%

5%



image size compressed: 4,8 MB

Dockerhub pull time: ~ 0.7 sec

attack surface - FROM golang:1.12

```
CLAIR_OUTPUT=High CLAIR_ADDR=http://192.168.99.100:30060 klar  
lotharschulz/hello:build.docker-cache--0.2.101  
clair timeout 1m0s  
docker timeout: 1m0s  
no whitelist file  
Analysing 11 layers  
Got results from Clair API v1  
Found 300 vulnerabilities  
Unknown: 8  
Negligible: 71  
Low: 95  
Medium: 107  
High: 19
```


attack surface - FROM scratch

```
CLAIR_OUTPUT=High CLAIR_ADDR=http://192.168.99.100:30060 klar  
lotharschulz/hello:build.docker-min--0.2.101  
clair timeout 1m0s  
docker timeout: 1m0s  
no whitelist file  
Analysing 2 layers  
Got results from Clair API v1  
Found 0 vulnerabilities
```

Conclusions

Balance

- docker image size
- security considerations
- infrastructural effort

Conclusions

- potentially reduced container startup time (pull)
- potentially reduced attack surface

"Small Images are always better in terms of absolute performance"

I'm sure you have questions

lotharschulz.info

pic: <https://www.flickr.com/photos/eyecatcherfotosde/25397937985>

Backup

CLAR OUTSIDE CLAR CLAR_ACOB#191.161.100.20666 klar bitnorchublog/build/docker-cache-6.2.101 2019.03.13
clear threat info
clobber threat info
Analyzing 11 files
Get results from Clear API v1
Found 101 items/20 files
Unread: 1
Images: 7
Low: 56, 107
Medium: 13

CVE-2018-16886 [patch]
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
Found in: <https://www.exploit-db.com/exploits/42880/>
A vulnerability in `get_page_owner` of `kernel` allows an attacker to supply arbitrary data across system's execution via `NorthPolar`. This can be used to improperly influence system execution and possibly lead to root privilege escalation. Affected releases are system's versions up to and including 3.18.
<https://security.tracker.debian.org/tracker/CVE-2018-16886>

CVE-2018-16887 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
systemd employs its `systemd` through 227 miscellaneous symlinks present in non-terminal path components, which allows local users to obtain ownership of arbitrary files via vectors involving creation of a directory and a file under that directory, and later replacing that directory with a symlink. This occurs even if the `fs.protected_symlinks` sysctl is turned on.
<https://www.exploit-db.com/exploits/42880/>

CVE-2018-16900 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `git` 2.18 and earlier, there is confusion in the usage of `getcwd()` by `readpath()` which can be used to write before the destination buffer leading to a buffer overflow and potential code execution.
<https://security.tracker.debian.org/tracker/CVE-2018-16900>

CVE-2018-16462 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
A race condition in the implementation of the `posix_memalign` in mergefs functions in the GNU C Library (`glibc` or `libc`) 2.26 and earlier could cause those functions to return a pointer to a heap area that is too small, potentially leading to heap corruption.
<https://security.tracker.debian.org/tracker/CVE-2018-16462>

CVE-2018-16523 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
A race condition in the GNU C Library (`glibc` or `libc`), from version 2.24 to 2.26 on `powerpc`, and only in version 2.26 on `i386`, did not properly handle malloc calls with arguments close to `SOME_MAX` and could return a pointer to a heap region that is smaller than requested, eventually leading to heap corruption.
<https://security.tracker.debian.org/tracker/CVE-2018-16523>

CVE-2018-16186 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In the GNU C Library (`glibc` or `libc`) through 3.25, `process_next_node` in `posix/memalign.c` has a heap-based buffer overflow via an attempted case-insensitive regular-expression match.
<https://security.tracker.debian.org/tracker/CVE-2018-16186>

CVE-2018-2779 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
A vulnerability in `usb_lsm` allows local users to escape to the parent session via a crafted TIOCSTI ioctl call, which pushes characters to the terminal's input buffer.
<https://security.tracker.debian.org/tracker/CVE-2018-2779>

CVE-2017-17438 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In Mercurial before 4.4.4, it is possible that a specially malformed repository can cause OS vulnerabilities to run arbitrary code in the form of a `githubhook=push-update script` checked into the repository. Typical use of Mercurial prevents construction of such repositories, but they can be created programmatically.
<https://security.tracker.debian.org/tracker/CVE-2017-17438>

CVE-2018-13167 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In Mercurial before 4.4, it is possible that a specially malformed repository and subtraction, via CVE-2018-13142.
<https://security.tracker.debian.org/tracker/CVE-2018-13167>

CVE-2018-7460 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
The Direct Rendering Manager (DRM) subsystem in the Linux kernel through 4.8 mishandles requests for Graphics Execution Manager (GEM) objects, which allows constant-dependent attackers to cause a denial of service (memory consumption) via an application that processes graphics data, as demonstrated by JavaScript code that creates many CANVAS elements for rendering by Chrome or Firefox.
<https://security.tracker.debian.org/tracker/CVE-2018-7460>

CVE-2018-12021 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `lib_raid` in the `refs.kh` filesystem driver in the Linux kernel 4.15.0 allows attackers to trigger a stack-based out-of-bounds write and cause a denial of service (kernel oops or panic) or possibly have unspecified other impact via a crafted file filesystem.
<https://security.tracker.debian.org/tracker/CVE-2018-12021>

CVE-2018-7388 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `lib_raid` in the `refs.kh` filesystem driver in the Linux kernel before 4.20.0, performs undesirable out-of-bounds speculation on pointer arithmetic. In various cases, including cases of different states or both to sanitize, leading to side-channel attacks.
<https://security.tracker.debian.org/tracker/CVE-2018-7388>

CVE-2018-6570 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `lib_raid`, `convert`, there is a possible memory corruption due to a use after free. This could lead to local escalation of privilege with System execution privileges needed. User interaction is not needed for exploitation. Product: Android. Vendor: Android kernel. Android ID: A-20180501.
<https://security.tracker.debian.org/tracker/CVE-2018-6570>

CVE-2018-12020 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `lib_raid`, `convert`, there is a possible memory corruption due to a use after free. This could lead to local escalation of privilege with System execution privileges needed. User interaction is not needed for exploitation. Product: Android. Vendor: Android kernel. Android ID: A-20180501.
<https://security.tracker.debian.org/tracker/CVE-2018-12020>

CVE-2018-12019 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `lib_raid`, `convert`, there is a possible memory corruption due to a use after free. This could lead to local escalation of privilege with System execution privileges needed. User interaction is not needed for exploitation. Product: Android. Vendor: Android kernel. Android ID: A-20180501.
<https://security.tracker.debian.org/tracker/CVE-2018-12019>

CVE-2017-100279 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
The Linux kernel's copying an `inode` system will sometimes map the contents of PIE assemblies, the heap or ld so to where the stack is mapped allowing attackers to more easily manipulate the stack. Linux kernel version 4.11.5 is affected.
<https://security.tracker.debian.org/tracker/CVE-2017-100279>

CVE-2018-16888 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
A memory leak in the `kernel_read_file` function in `fs` in the Linux kernel through 4.20.11 allows attackers to cause a denial of service (memory consumption) by triggering `vh_read` failure.
<https://security.tracker.debian.org/tracker/CVE-2018-16888>

CVE-2017-14264 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
In `libata` before 4.1, the `rewarnen` tool could be made to manipulate internal data structures in ways unintended by the authors. Malformed input may lead to crashes (with a buffer overflow or other memory corruption) or other unspecified behaviors. This crosses a privilege boundary in, for example, certain web-hosting environments in which a Control Panel allows an unprivileged user account to create subdomains.
<https://security.tracker.debian.org/tracker/CVE-2017-14264>

CVE-2017-14662 [patch]
Found in: <https://www.exploit-db.com/exploits/42880/>
Fixed By: <https://github.com/SecWiki/linux-kernel-4.15.0-rc4-00100>
Integer overflow in the `decide_nid` function in `purty_decide` in `LibSNI` before 2.5 allows remote attackers to cause a denial of service or possibly have unspecified other impact.
<https://security.tracker.debian.org/tracker/CVE-2017-14662>

Team Journey part 1

```
FROM golang:1.11
```

```
RUN mkdir /app
```

```
WORKDIR /app
```

```
COPY hello.go .
```

```
RUN go build -o hellogo .
```

```
RUN groupadd -g 99 appuser && useradd -r -u 99 -g appuser appuser
```

```
USER appuser
```

```
CMD ["/hellogo"]
```

```
$ docker build --rm -t ...
```

Team Journey part 1

Image size compressed: 299 MB

CircleCI build time: ~ 3 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 21 ses

pull to empty nodes, no layers cached

Team Journey part 1.1

```
FROM golang:1.11
```

```
RUN mkdir /app
```

```
WORKDIR /app
```

```
COPY hello.go .
```

```
RUN go build -o hellogo .
```

```
RUN groupadd -g 99 appuser && useradd -r -u 99 -g appuser appuser
```

```
USER appuser
```

```
CMD ["/hellogo"]
```

```
$ docker build --cache-from golang:1.11 -t ...
```

Team Journey part 1.1

previous step

Image size compressed: 299 MB

CircleCI build time: ~ 3 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 21 ses



current step

Image size compressed: 299 MB

CircleCI build time: ~ **0.5 sec**

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 21 ses

cache warm up in preparation step

Team Journey part 2 - builder pattern

#part one

FROM golang:1.11 as builder

workdir setup

COPY hello.go .

RUN CGO_ENABLED=0 GOOS=linux go build -a -installsuffix cgo -ldflags '-w -extldflags "-static"' -o hellogodocker .

part two

FROM alpine:latest

user & workdir setup

COPY --from=builder /go/src/github.com/lotharschulz/hellogodocker/hellogodocker .

CMD ["/hellogodocker"]

\$ docker build --rm -t ...

use a static linked binary

Team Journey part 2 - builder pattern

previous step

Image size compressed: **299 MB**
CircleCI build time: ~ 0.5 sec
Dockerhub push time: ~ 4 sec
Dockerhub pull time: ~ **21 ses**



current step

Image size compressed: **5 MB**
CircleCI build time: ~ 0.6 sec
Dockerhub push time: ~ 4 sec
Dockerhub pull time: ~ **1.2 ses**



Team Journey part 2 - builder pattern - squash & compress

```
#part one
FROM golang:1.11 as builder
# workdir setup
COPY hello.go .

RUN CGO_ENABLED=0 GOOS=linux go build -a -installsuffix cgo -ldflags '-w -extldflags
"-static"' -o hellogodocker .

# part two
FROM alpine:latest
# user & workdir setup ....
COPY --from=builder /go/src/github.com/lotharschulz/hellogodocker/hellogodocker .

CMD ["/hellogodocker"]
```

```
$ docker build --squash/--compress --rm -t ...
```

squash is an [experimental Docker daemon option](#)

Team Journey part 2 - builder pattern - squash & compress

previous step



current step

Image size compressed: 5 MB

CircleCI build time: ~ 0.6 sec

Dockerhub push time: ~ 4 sec

Dockerhub pull time: ~ 1.2 ses

Image size compressed: 5 MB

CircleCI build time: ~ 0.7/1.1 sec

Dockerhub push time: ~ 4/4.9 sec

Dockerhub pull time: ~ 1.1/1.4 ses

Team Journey part 3 - alpine base

```
FROM alpine:latest  
# ... set up user
```

```
ADD hellogo /  
CMD ["/hellogo"]
```

```
$ docker build [--squash/--compress] --rm -t ...
```

static linked binary created within ci worker node outside of Dockerfile scope

Team Journey part 3 - alpine base

previous step



current step

Image size compressed: 5 MB

CircleCI build time: ~ 0.7/1.1 sec

Dockerhub push time: ~ 4/4.9 sec

Dockerhub pull time: ~ 1.1/1.4 ses

Image size compressed: 5 MB

CircleCI build time: ~ 0.6/1.1 sec

Dockerhub push time: ~ 3.8/4.6 sec

Dockerhub pull time: ~ 1.1 ses

Team Journey part 4 - FROM scratch

```
FROM scratch
```

```
ADD ca-certificates.crt /etc/ssl/certs/
```

```
ADD hellogo /
```

```
CMD ["/hellogo"]
```

```
$ docker build [--squash/--compress] --rm -t ...
```

static linked binary created within ci worker node outside of Dockerfile scope

Team Journey part 4 - FROM scratch

previous step

Image size compressed: 5 MB

CircleCI build time: ~ 0.6/1.1 sec

Dockerhub push time: ~ 3.8/4.6 sec

Dockerhub pull time: ~ 1.1 ses

current step



Image size compressed: 2 MB



CircleCI build time: ~ 16 sec

Dockerhub push time: ~ 3.5/3.9 sec

Dockerhub pull time: ~ 0.8 ses