

Werkzeuge zum Bauen von Container-Images

Tools to Build Container Images

Patrick Harböck / Martin Höfling

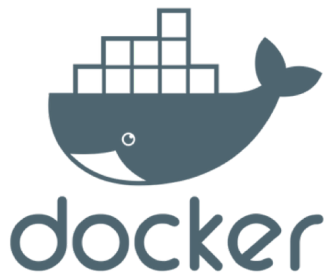
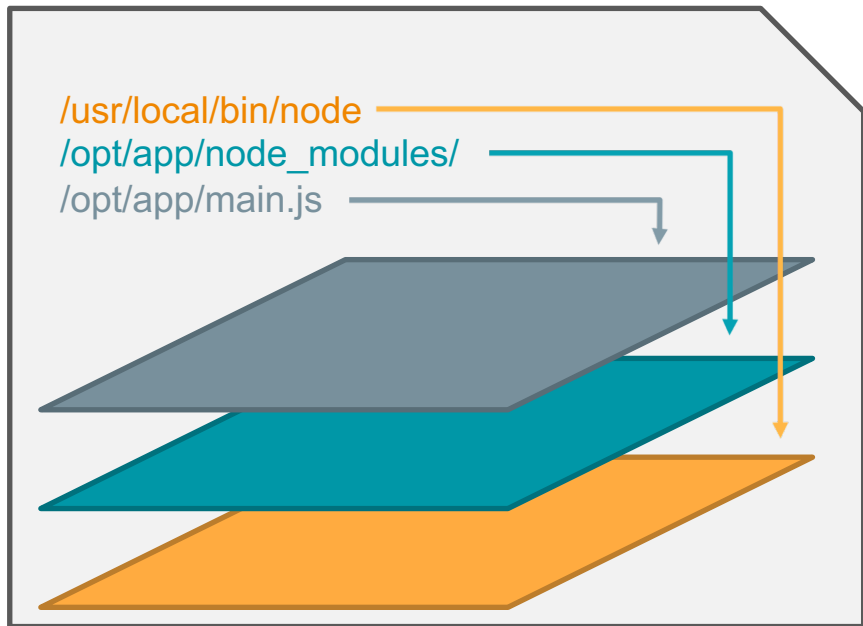


November 13th, 2019
(ContainerConf)

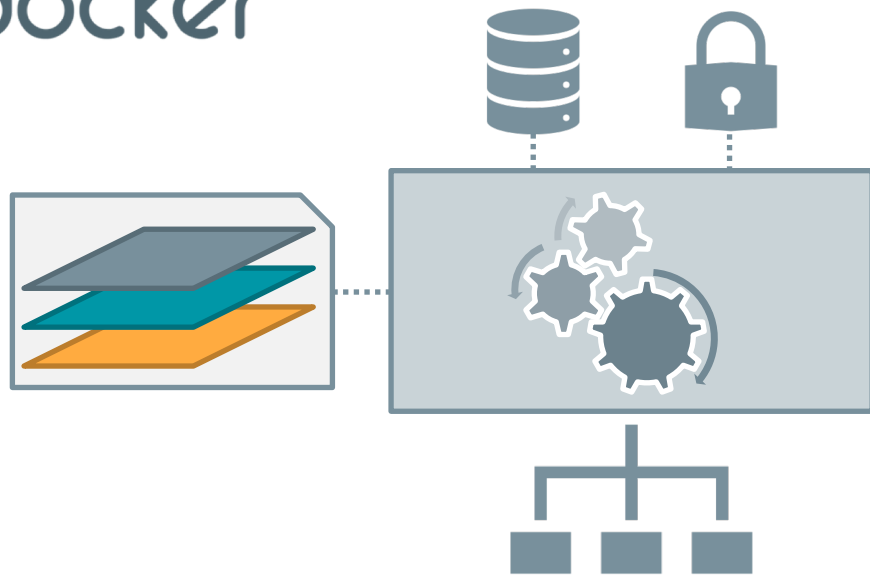


What is Docker?

Container Image Builder



Container Runtime



Docker Images

→ ~ docker history node

IMAGE	CREATED	CREATED BY	SIZE
MENT			
d8c33ae35f44	41 hours ago	/bin/sh -c #(nop) CMD ["node"]	0B
<missing>	41 hours ago	/bin/sh -c #(nop) ENTRYPOINT ["docker-entry...	0B
<missing>	41 hours ago	/bin/sh -c #(nop) COPY file:238737301d473041...	116B
<missing>	41 hours ago	/bin/sh -c set -ex && for key in 6A010...	5.47MB
<missing>	41 hours ago	/bin/sh -c #(nop) ENV YARN_VERSION=1.17.3	0B
<missing>	41 hours ago	/bin/sh -c ARCH= && dpkgArch="\$(dpkg --print...	66.6MB
<missing>	41 hours ago	/bin/sh -c #(nop) ENV NODE_VERSION=12.10.0	0B
<missing>	41 hours ago	/bin/sh -c groupadd --gid 1000 node && use...	333kB
<missing>	45 hours ago	/bin/sh -c set -ex; apt-get update; apt-ge...	562MB
<missing>	45 hours ago	/bin/sh -c apt-get update && apt-get install...	142MB
<missing>	45 hours ago	/bin/sh -c set -ex; if ! command -v gpg > /...	7.81MB
<missing>	45 hours ago	/bin/sh -c apt-get update && apt-get install...	23.2MB
<missing>	46 hours ago	/bin/sh -c #(nop) CMD ["bash"]	0B
<missing>	46 hours ago	/bin/sh -c #(nop) ADD file:9788b61de35351489...	101MB

Container Images

- Manifest / Metadata

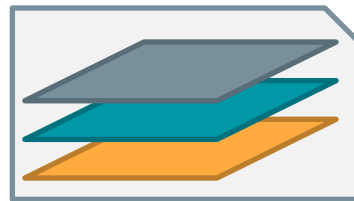
ENV | WORKDIR | USER | CMD

- Default configuration for creating containers
- Content hashes of layers to ensure integrity

Layer1: 7d97e98f8af71
Layer2: e703abc8f639e

- Layers

- File system packed with *tar*
- Multiple layers → root file system for containers



Container Image Format Evolution

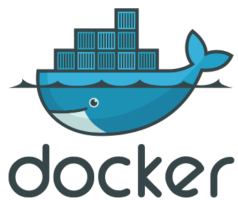


Image Spec v1

Registry v2

Image Spec v1.2

2013

2014

2015

2016

2017



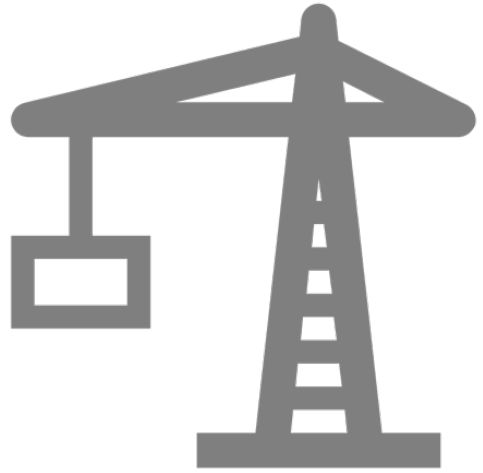
OPEN CONTAINER
INITIATIVE

OCI Image Spec v1

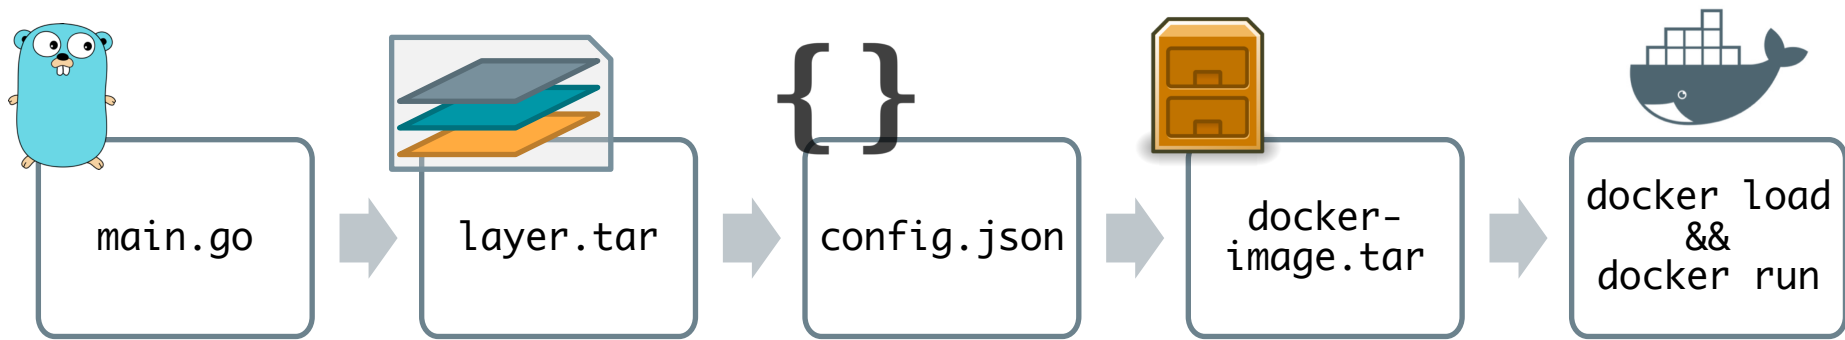
Open Container Initiative



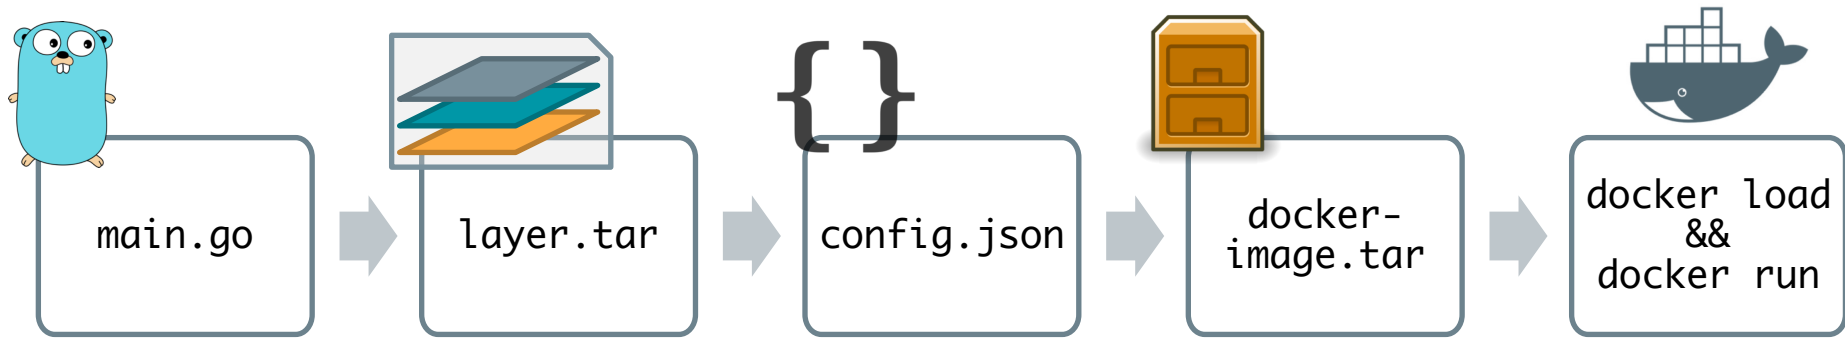
DEMO: Build a Container Image from Scratch




DEMO: Build a Container Image from Scratch



DEMO: Build a Container Image from Scratch



- No elevated privileges required
- No Dockerfile
- No *docker build*



What's wrong with building images via Docker?



→ **Security**

→ **Scalability**

→ **Flexibility**

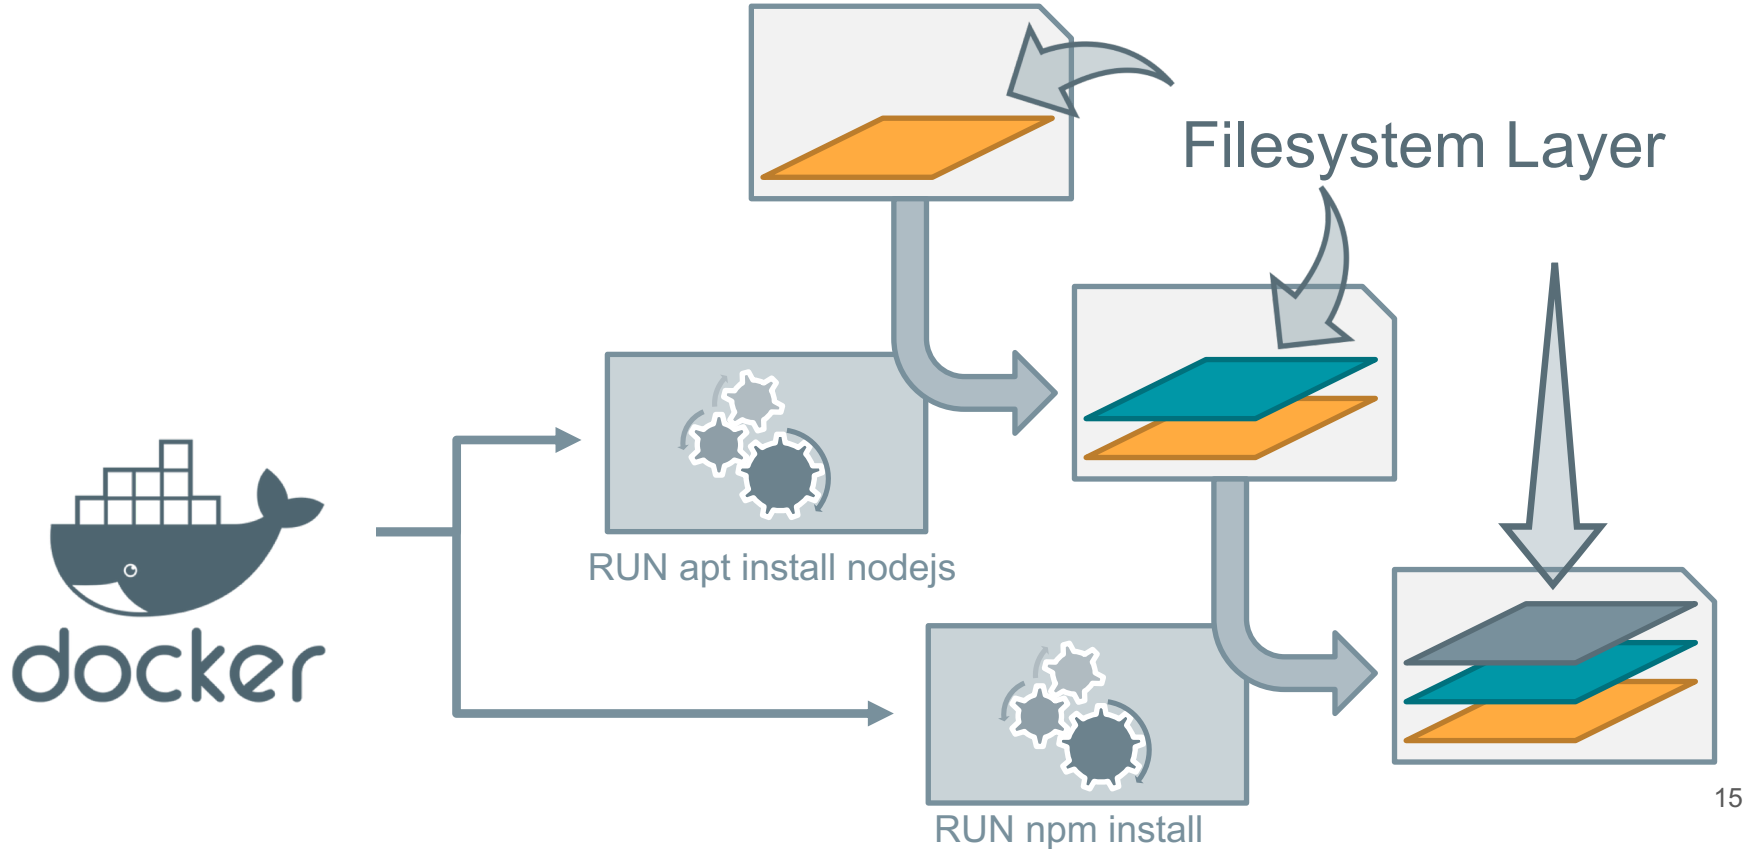


→ **Security**

→ **Scalability**

→ **Flexibility**

How Docker builds container images



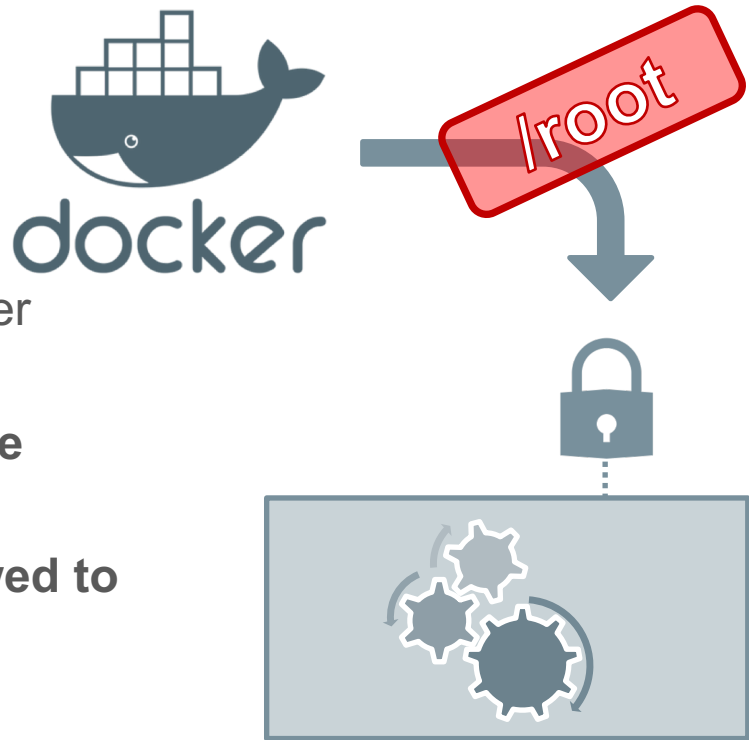
How Docker builds container images

- docker build uses Docker containers
- Docker containers require isolation
- Docker requires elevated privileges
- Build pipelines / developers can access Docker

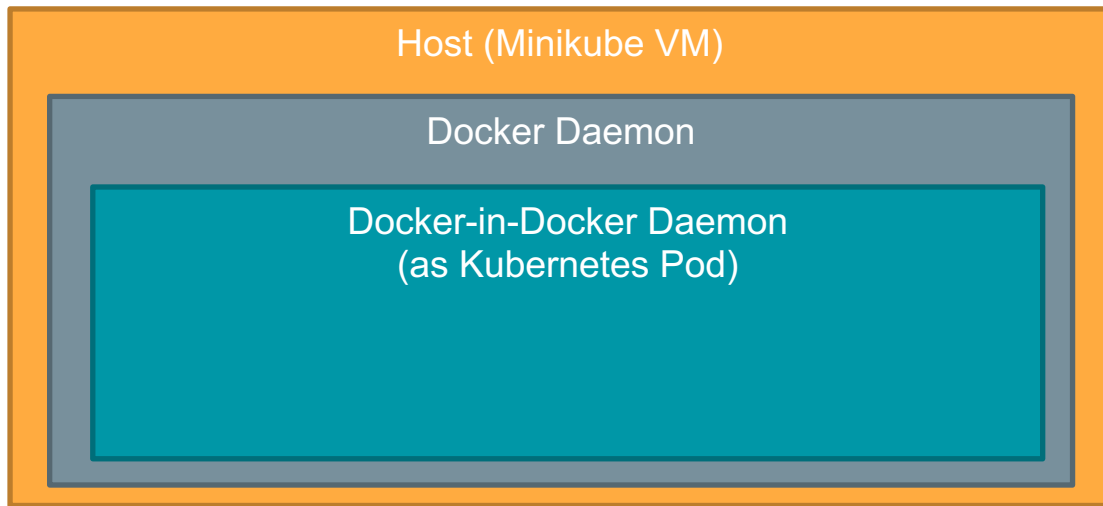
➔ Security nightmare on shared infrastructure

“First of all, **only trusted users should be allowed to control your Docker daemon.**”

<https://docs.docker.com/engine/security/security/#docker-daemon-attack-surface>

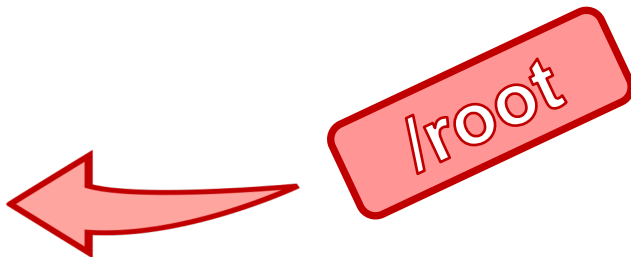


DEMO: Host Access via privileged container

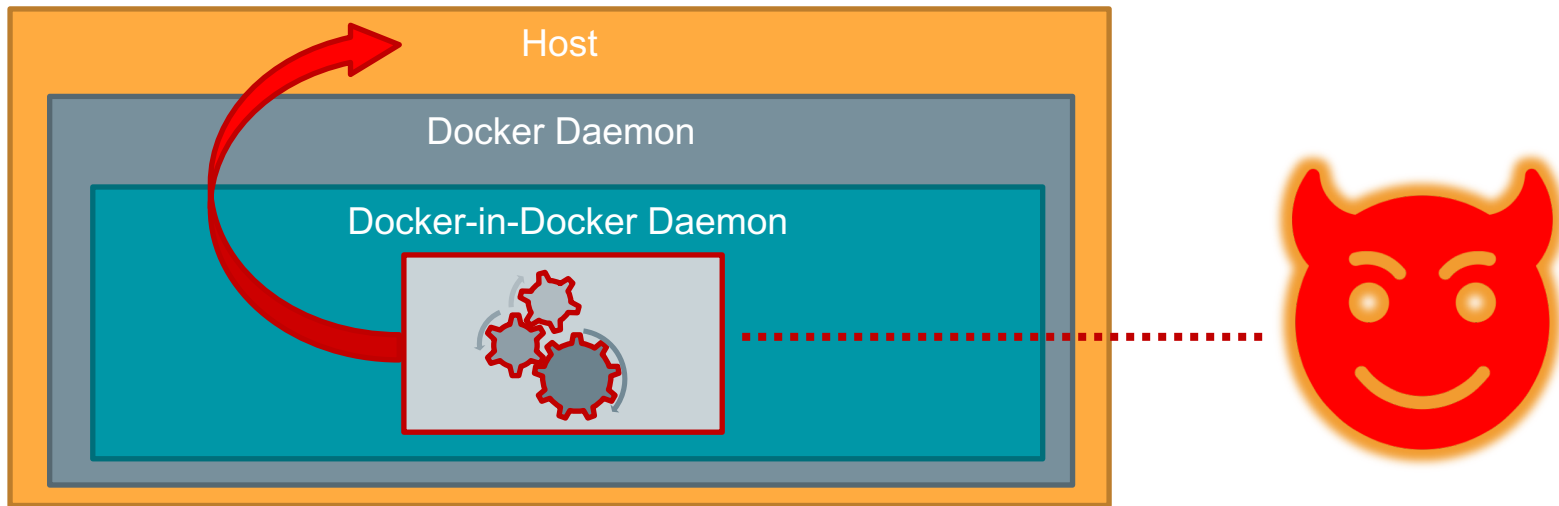


Docker in Docker Kubernetes Pod Spec

```
apiVersion: v1
kind: Pod
metadata:
  name: dind
spec:
  hostname: dind-pod
  containers:
  - name: dind
    image: docker:dind
    securityContext:
      privileged: True
    ports:
      - containerPort: 2375
```



DEMO: Host Access via privileged container



Security Risks?

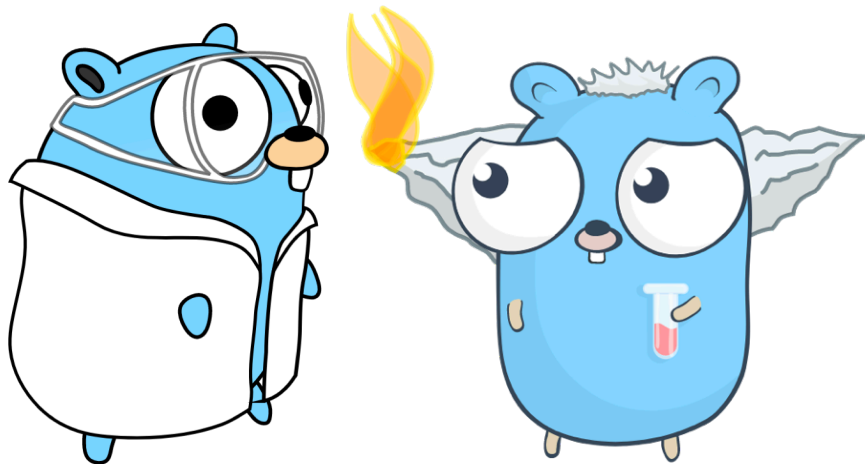
- Privileged Docker-in-Docker → **full host access**
 - Mounting or exposing Docker socket → **full host access**
 - Base image runs as container *root* → larger vulnerability surface
- Easy to break and lose container isolation



Remark: Hermetic Builds and Reproducibility

→ Hermetic: sandboxed build process

→ Reproducible builds result in verifiable artifacts





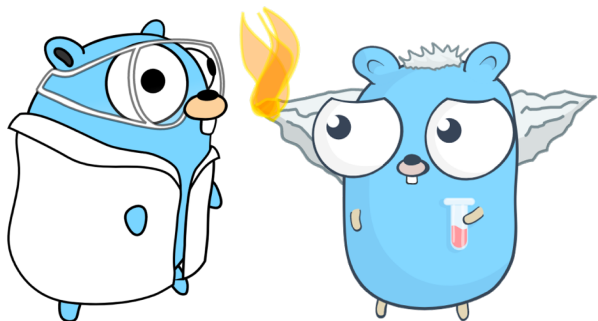
→ Security

→ **Scalability**

→ Flexibility

Caching

- Allows scaling up CI/CD pipelines
- Reuse base layers across different branches and builds
- Reproducible builds improve caching



Build Pipeline



kubernetes



Jenkins

Github

- source code repository



docker build

- docker-in-docker
- privileged pod



Docker Registry

- pull cache
- push images



Scalability Issues

- One Docker daemon does not scale for parallel builds
- No distributed caching support



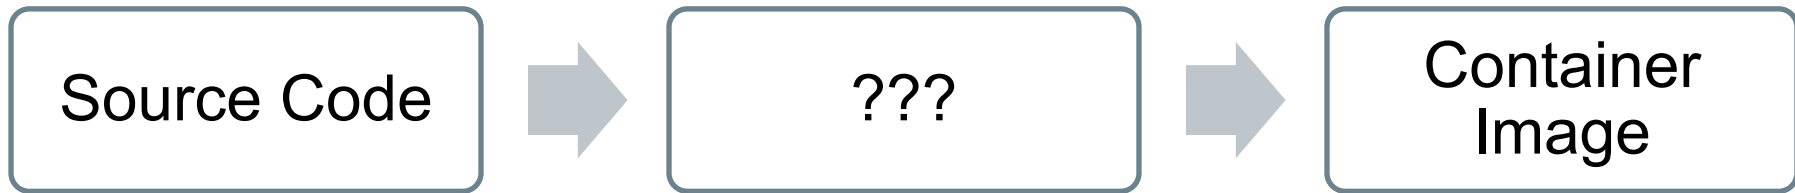
→ Security

→ Scalability

→ **Flexibility**

Flexibility

- How restricted is the build process and image definition?
 - Can developers use any tools and languages they want?
 - How well does it integrate into an existing development pipeline?



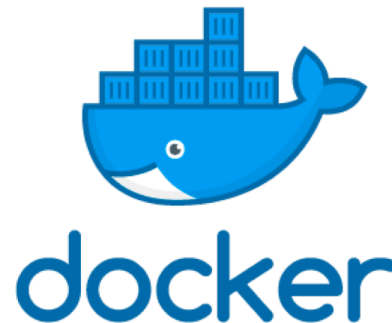
Dockerfile based Tools

- Extract base layer(s)
- Run a command in sub container or directly
- Snapshot Filesystem

✓ Generic

– Problem: Supported Dockerfile Features?

- USER – run commands as specific user
- Multistage builds
- Root vs. non-root



```
FROM ubuntu:18.04
```



```
RUN apt-get update
```

```
RUN apt-get install -y nginx
```

```
COPY nginx.conf /etc/nginx/
```

```
EXPOSE 8080
```

```
ENTRYPOINT ['/usr/bin/nginx']
```

Tool	Primary Maintainer	Security	Scalability	Flexibility
BuildKit	Docker			Dockerfile

- Focus on scalability, performance, extensibility
- Experimental support in newer Docker versions and with *docker buildx*
- Optional rootless mode

moby / buildkit
Watch 60

Code
Issues 96
Pull requests 8
Projects 0
Insights

concurrent, cache-efficient, and Dockerfile-agnostic builder toolkit <https://github.com/moby/moby/issues/3>

buildkit
dockerfile
docker
oci-image
oci
containers
builder



1,793 commits
2 branches
17 releases
61 contributors

README.md

```

# ./example | buildctl build
INFO[0000] tracing logs to /tmp/buildctl325403184
[+] Building 16.6s (2/19)
=> docker-image://docker.io/library/golang:1.8-alpine
=> => resolve docker.io/library/golang:1.8-alpine
=> => sha256:f8dc884e6d7b8b3cd2899efdc89d1a6949e9fe18df08be482bfff5c2ae9a5 1.572kB /
=> => sha256:b259295eadd4bee43b57ca16ab840851277ac5fb15e77b34e2fa09efce5adb77 487B
=> => sha256:5e73dfc64116fcf9fada3a2f372783ed0088812a107e26bf017dc9b2b96e7a38 126B
=> => sha256:8b97decbbcecab3c45982ba6fe6e2a44487fecafe5007c0af215bb76c39254 1.35kB
=> => sha256:6555f92e77d3062fc8fc5a7d5f8b589f0bf0e38f2624397ac502b300dc6 350.6kB /
=> => sha256:6f821164d5b7ec94868795c1fb8dc6f51e04f97a6cf3a487868f2f 1.98MB /
=> => sha256:e7baf3b1a3a56a214658c5cb102c54e2b91e95245b6f1b6ca5 4.144kB /

```

Tool	Primary Maintainer	Security	Scalability	Flexibility
Buildah	Red Hat			Dockerfile

- Secure and flexible builds of OCI images
- Intended as a Docker replacement together with Podman

containers / buildah


Watch 67 Unstar 1,530 Fork

Code Issues 47 Pull requests 13 Projects 0 Wiki Insights

A tool that facilitates building OCI images



1,353 commits 3 branches 29 releases 58 contributors Apache-2.0

README.md



buildah

Buildah - a tool that facilitates building **Open Container Initiative (OCI)** container images

Tool	Primary Maintainer	Security	Scalability	Flexibility
Kaniko	Google			Dockerfile

- Designed for Kubernetes
- Compatible with
 - AppArmor / SELinux
 - gVisor
- Focus on security and performance
- Reproducible builds

GoogleContainerTools / kaniko

Watch 104 Unstar 3,542 Fork

<> Code Issues 95 Pull requests 19 Projects 2 Wiki Insights


Build Container Images In Kubernetes



652 commits 4 branches 10 releases 57 contributors Apache-2.0

README.md

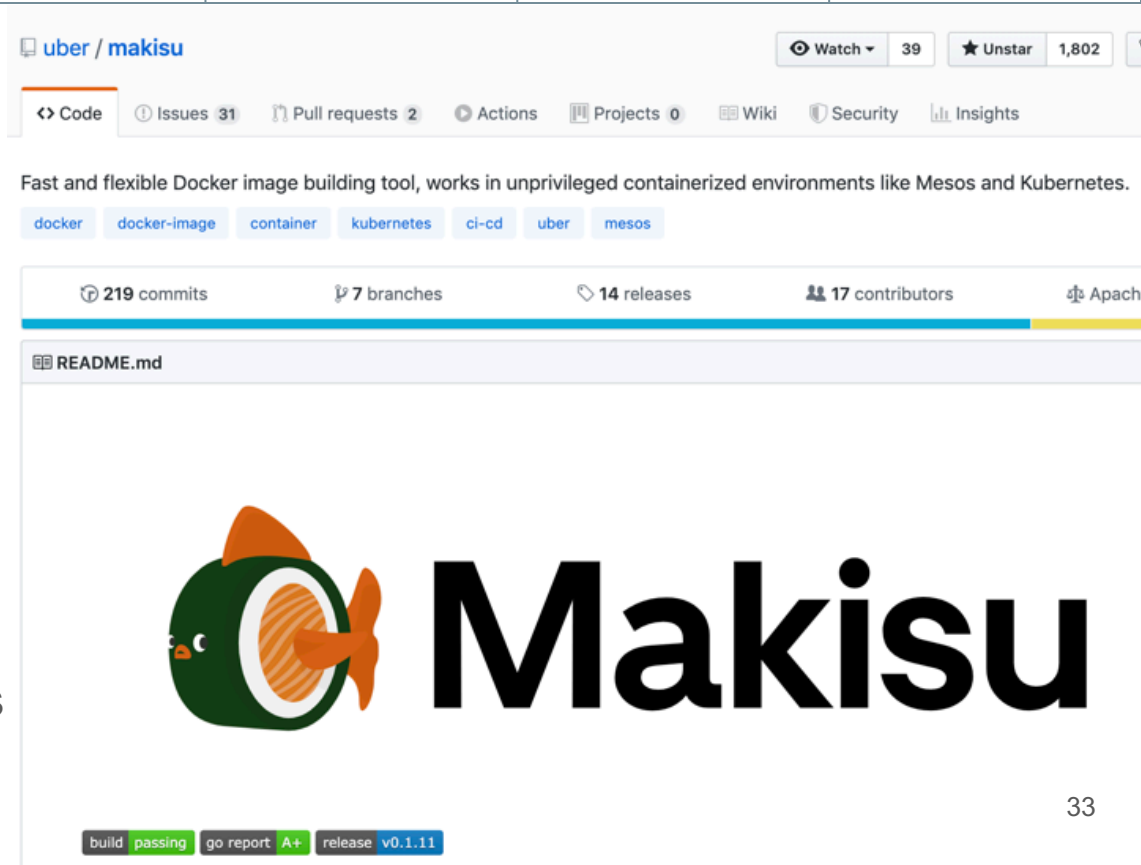
kaniko - Build Images In Kubernetes

build passing



Tool	Primary Maintainer	Security	Scalability	Flexibility
Makisu	Uber			Dockerfile

- Focus on security and performance
- Dockerfile support with opinionated modifications
- Distributed caching of layers





The screenshot shows the GitHub repository page for 'uber / makisu'. At the top, the repository name is displayed with options to 'Watch' (39), 'Unstar' (1,802), and a search icon. Below this, navigation tabs include 'Code', 'Issues' (31), 'Pull requests' (2), 'Actions', 'Projects' (0), 'Wiki', 'Security', and 'Insights'. A description states: 'Fast and flexible Docker image building tool, works in unprivileged containerized environments like Mesos and Kubernetes.' Below the description are tags for 'docker', 'docker-image', 'container', 'kubernetes', 'ci-cd', 'uber', and 'mesos'. A progress bar shows repository statistics: 219 commits, 7 branches, 14 releases, 17 contributors, and Apache 2.0 license. The 'README.md' file is open, showing the 'Makisu' logo (a green fish with an orange mouth) and the word 'Makisu' in large black text. At the bottom, a status bar indicates 'build passing', 'go report A+', and 'release v0.1.11'.

Tailored image construction

- Tailored for a distinct language and build-system
- The actual **build is not performed in a (child) container**
- The build result is often combined with a base image
 - e.g. Node.js runtime + node_modules + application

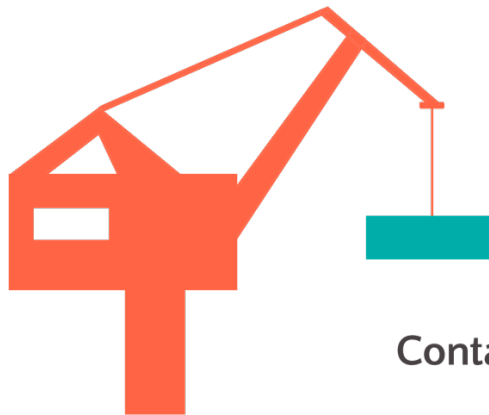
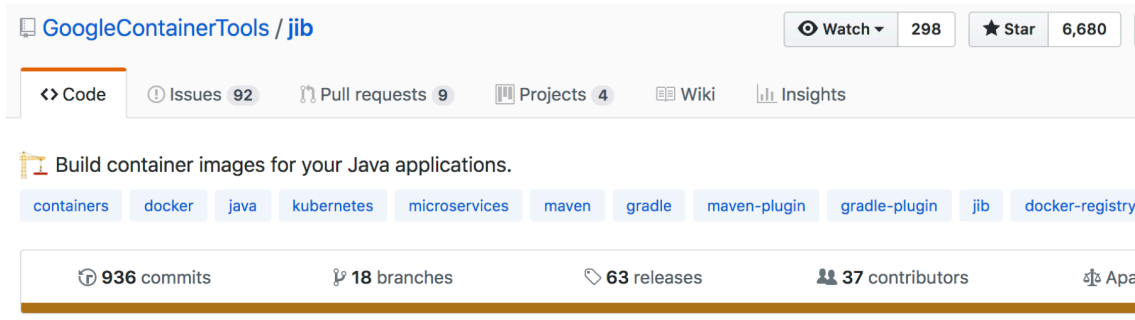


Tool	Primary Maintainer	Security	Scalability	Flexibility
Jib	Google			Java only

- Maven / Gradle plugin
- Distroless Java base image



Builds are

- ✓ Minimal
- ✓ Reproducible
- ✓ Fast (caching)



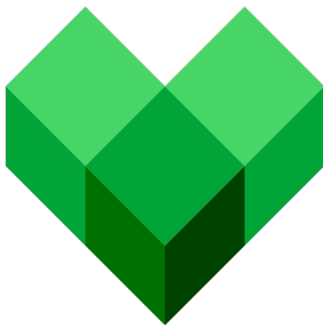
Jib

Containerize your Java application.

Tool	Primary Maintainer	Security	Scalability	Flexibility
Bazel	Google			Starlark rules


- Supports Python, Node.js, Java, C/C++, Go, Rust, ...


Builds are




- ✓ Minimal
- ✓ Reproducible
- ✓ Fast (caching)
- Complex rules written in Starlark

 [bazelbuild / rules_docker](#)

 Code

 Issues 52

 Pull requests 6

 Projects 0

 Insights

Rules for building and handling Docker images with Bazel

[bazel](#)

[docker](#)


[docker-image](#)

[bazel-rules](#)

[cloud](#)

[google](#)

 656 commits

 36 branches

 11 releases

 README.md

Bazel Container Image Rules


Travis CI


Bazel CI


build passing


build passing

Tool	Primary Maintainer	Security	Scalability	Flexibility
OpenShift Source-to-Image	Red Hat			Common stacks

 openshift / source-to-image


 Watch ▾ 248

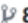
 Star 1,275


 Fork 394


[Code](#)
[Issues 12](#)
[Pull requests 2](#)
[Projects 0](#)
[Insights](#)


A tool for building/building artifacts from source and injecting into docker images


 1,340 commits

 8 branches

 31 releases

 61 contributors

 Apache-2.0

 README.md

Source-To-Image (S2I)

Overview



go report
A+


godoc
reference




build
passing

license
Apache-2.0

Source-to-Image (S2I) is a toolkit and workflow for building reproducible Docker images from source code. S2I produces ready-to-run images by injecting source code into a Docker container and letting the container prepare that source code for execution. By creating self-assembling **builder images**, you can version and control your build environments exactly

Tool	Primary Maintainer	Security	Scalability	Flexibility
Cloud Native Buildpacks	Heroku / Pivotal / CNCF			Common stacks


[buildpack / pack](#)

 Watch ▾
 13
  Star
 377
  Fork
 24

<> Code

🔔 Issues 35

🔗 Pull requests 4

📁 Projects 0

📖 Wiki

📊 Insights

Local CLI for building apps using Cloud Native Buildpacks <https://buildpacks.io>

🔖 398 commits

🌿 20 branches

📦 10 releases

👤 17 contributors

📄 Apache-2.0



















📖 README.md

pack - Buildpack CLI build passing

pack makes it easy for

- Application developers to use [Cloud Native Buildpacks](#) to convert code into runnable images
- Buildpack authors to develop and package buildpacks for distribution

Ready to embark on your adventure with **pack** but not sure where to start? Try out our tutorial, [An App's Brief Journey from Source to Image](#).

Tool	Primary Maintainer	Security	Scalability	Flexibility
docker build	Docker			Dockerfile
docker buildx (BuildKit)	Docker			Dockerfile
Buildah	Red Hat			Dockerfile
Kaniko	Google			Dockerfile
Makisu	Uber			Dockerfile
Jib	Google			Java only
Bazel	Google			Starlark rules
OpenShift Source-to-Image	Red Hat			Common stacks
Cloud Native Buildpacks	Heroku / Pivotal / CNCF			Common stacks

What should
I use now?



Use case: Small Team

- No strict security requirements for team isolation
- Teams have full access to CI infrastructure

→ **Docker**

Still a valid choice

→ **Buildah**

Flexible, only parts of Dockerfile syntax supported securely

→ **BuildKit**

Are you feeling adventurous? Potential transition path with *docker buildx*



Use case: Multiple teams, Provided K8s infrastructure

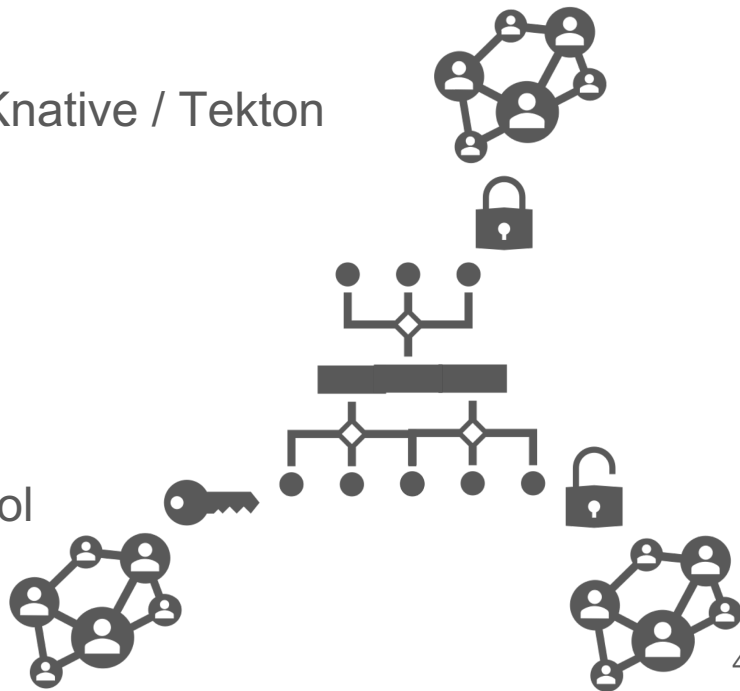
- Cannot modify K8s infrastructure, no privileged containers, no container nesting
- Teams are isolated, e.g. on namespace level

➔ **Kaniko**, e.g. combined with Skaffold or Knative / Tekton

- Shared volume caching
(e.g. on Google Cloud Platform)

➔ Take a look at: **Makisu**

- Fine grained distributed cache control



Use case: No Dockerfile required

e.g. modern dev stack, Container Native Team

→ Bazel

→ Jib

→ Cloud Native Buildpacks



~~apt-get install python-dev~~

No classical ops pattern!

Docker-less Infrastructure?



Re-evaluate your container build process!



[/martinhoefling](#)



[/Pharb](#)



[/TNG](#)

TNG  **TECHNOLOGY
CONSULTING**