

Building OCI Images Without Privilege

tycho@tycho.ws, tycander@cisco.com
github.com/tych0









FFFFFFFF

FFFFFFFF

FFFFFFF

FFFUU


UUUU

UUUU

UUUU

UUUU

UUUU-





jonboulle spec: add conversion references db4d6de on Jun 26, 2017

8 contributors

Open Container Initiative

Image Format Specification

This specification defines an OCI Image, consisting of a [manifest](#), an [image index](#) (optional), a set of [filesystem layers](#), and a [configuration](#).

The goal of this specification is to enable the creation of interoperable tools for building, transporting, and preparing a container image to run.

Table of Contents

- [Introduction](#)
- [Notational Conventions](#)

\$ not #


```
yum -y install ...  
./configure ...  
sed -i ...
```

What exists?

shiftfs

- In kernel solution to uid-map files based on namespace map
- Author uses it for building container images
- Other interesting applications
- <https://lwn.net/ml/linux-fsdevel/1529098514.4048.41.camel@HansenPartnership.com/>

“rootless” containers

- umoci has rootless support without user namespaces
- Buildah has (recent) support for user namespaces

	Works on mainline kernels today	Can run arbitrary commands	Works without privilege
ShiftFS	X	✓	✓
 umoci	✓	X	✓
 buildah	✓	✓	✓

stacker

How do I use it?

first:

from:

type: docker

url: docker://centos:latest

import:

- config.json

- install.sh

run: |

```
mkdir -p /etc/myapp
```

```
cp /stacker/config.json /etc/myapp/
```

```
/stacker/install.sh
```

How do I use it?

first:

from:

type: docker

url: docker://centos:latest

import:

- config.json

- install.sh

run: |

mkdir -p /etc/myapp

cp /stacker/config.json /etc/myapp/

/stacker/install.sh

How do I use it?

```
first:
```

```
  from:
```

```
    type: docker # or tar, oci, etc.
```

```
    url: docker://centos:latest
```

```
import:
```

```
  - config.json
```

```
  - install.sh
```

```
run: |
```

```
  mkdir -p /etc/myapp
```

```
  cp /stacker/config.json /etc/myapp/
```

```
  /stacker/install.sh
```

How do I use it?

```
first:
  from:
    type: docker
    url: docker://centos:latest
import:
  - config.json
  - install.sh
run: |
  mkdir -p /etc/myapp
  cp /stacker/config.json /etc/myapp/
  /stacker/install.sh
```

How do I use it?

```
first:
  from:
    type: docker
    url: docker://centos:latest
  import:
    - config.json
    - install.sh
  run: |
    mkdir -p /etc/myapp
    cp /stacker/config.json /etc/myapp/
    /stacker/install.sh
```

How does it work?

- liblxc
- go-lxc
- umoci
 - Newly merged API!
- skopeo
 - No API :(
- btrfs
 - Multiple images built from the same source are only extracted once

What does the run environment look like?

- User namespaces
- Host network namespace
- Bind mounted `/etc/resolv.conf`
- `/proc/sys` and `/proc/sysrq-trigger` readonly (proc:mixed in LXC)
- Reasonable devices in `/dev` (`lxc.autodev = 1`)
- Bind mounted `/sys` from host
- `/stacker` directory mounted r/o for import:s
- Reasonable default `$PATH`
- Mostly looks like a reasonable system, yum, apt, etc. work fine

odds & ends



--shell-fail



```
$ stacker inspect --oci-dir oci
```

```
a
```

```
    layer 0: sha256:256b176b... (75 MB)
```

```
    layer 1: sha256:276a625d... (156 kB)
```

```
Annotations:
```

```
  ws.tycho.stacker.stacker_yaml: ...
```

```
Image config:
```

```
{  
  "created": "2018-08-06T16:33:04.379695767-06:00",  
  "os": "linux",  
  "config": {  
    "Env": [  
      "PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin"  
    ], ...  
  }  
}
```


Thanks / Merci

We are hiring! Linux, containers, secure boot, etc.

tycho@tycho.ws, tycander@cisco.com

<http://github.com/tych0>

