

PuzzleFS

Bringing the PuzzleFS pieces together

Ariel Miculaș-Trif - Cisco

github.com/project-machine/puzzlefs

Design goals

Solve the most pertinent OCI v1 problems

- Reduced duplication
- Reproducible image build
- Direct mounting support
- Data integrity
- Memory safety guarantees
- Same implementation in userspace and kernel

PuzzleFS

- PuzzleFS is an immutable filesystem which shares design goals with the [OClv2 proposal brainstorm](#)
- It uses Content defined chunking (CDC) to split a filesystem into variable-sized chunks
- There is a PuzzleFS image rootfs which contains the necessary metadata in order to reassemble the filesystem when mounting the image

PuzzleFS integration goals

- build PuzzleFS images easily (similar to `docker build`)
- run LXC containers with the rootfs backed by a PuzzleFS image
- copy PuzzleFS images with existing tools (e.g. `skopeo`)
- integrate PuzzleFS into the OCI Image standard

Oci disk layout

```
/home/amiculas/.local/share/puzzlefs/ubuntu/
├── blobs
│   └── sha256
│       ├── 31e907dcc94a592a57796786399eb004dcbba714389fa615f5efa05a91316356
│       ├── 35a88802559dd2077e584394471ddaa1a2c5bfd16893b829ea57619301eb3908
│       ├── 3c645031de2917ade93ec54b118d5d3e45de72ef580b8f419a8cdc41e01d042c
│       ├── 6f6ec53d36a9504f01e3636cf68e0e03761a3b6947a95ba430ae553ee3aaf4d9
│       ├── 77d57fd89366f7d16615794a5b53e124d742404e20f035c22032233f1826bd6a
│       ├── 7af9ba4f0a47d9bc8b1ffa492c6b0276476f1889cf4e699fba2236924e5932ed
│       ├── 9c704ecd0c694c4cbdd85e589ac8d1fc3fd8f890b7f3731769a5b169eb495809
│       ├── aa772c98400ef833586d1d517d3e8de670f7e712bf581ce6053165081773259d
│       ├── b1e9cef3f2977f8bdd19eb9ae04f83b315f80fe4f5c5651fedf41482c12432f7
│       ├── bccd10f490ab0f3fba61b193d1b80af91b17ca9bdca9768a16ed05ce16552fcb
│       ├── c920ba4cfca05503764b785c16b76d43c83a6df8d1ab107e7e6610000d94315c
│       ├── ca2b0f26964cf2e80ba3e084d5983dab293fdb87485dc6445f3f7bbfc89d7459
│       ├── d35dfc2fe3ef66bcc085ca00d3152b482e6cafb23cdda1864154caf3b19094ba
│       ├── dafa2b0c44d2cfb0be6721f079092ddf15dc8bc537fb07fe7c3264c15cb2e8e6
│       └── edbfe74c41f8a3501ce542e137cf28ea04dd03e6df8c9d66519b6ad761c2598a
├── index.json
└── oci-layout
```

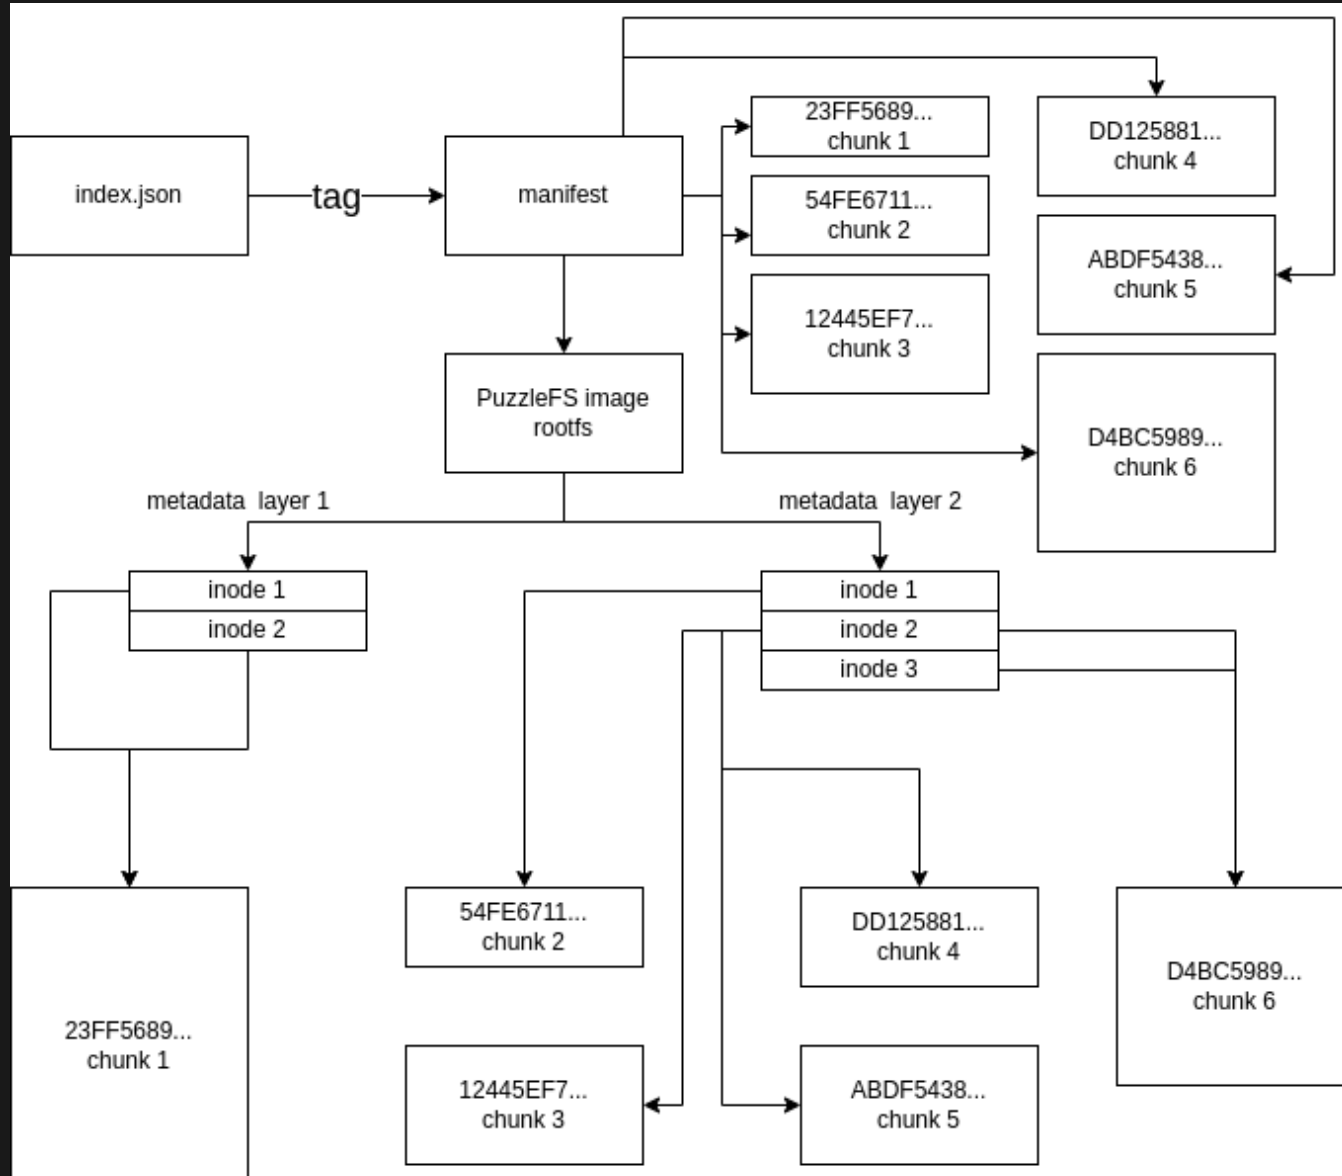
Oci specification: Image Manifest

```
{
  "schemaVersion": 2,
  "mediaType": "application/vnd.oci.image.manifest.v1+json",
  "config": {
    "mediaType": "application/vnd.oci.image.config.v1+json",
    "size": 2296,
    "digest": "sha256:b1e9cef3f2977f8bdd19eb9ae04f83b315f80fe4f5c5651fedf41482c12432f7"
  },
  "layers": [
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
      "size": 29749828,
      "digest": "sha256:dafa2b0c44d2cfb0be6721f079092ddf15dc8bc537fb07fe7c3264c15cb2e8e6"
    }
  ]
}
```

PuzzleFS Oci format

```
"layers": [  
  {  
    "annotations": {  
      "io.puzzlefsoci.puzzlefs.puzzlefs_verity_root_hash":  
        "343c67e13be60bfdc390feed9520d4da59456b2e2671212a2f2b9c8f66d51fb7"  
    },  
    "digest": "sha256:b66472a6cdef868cfecb865f1bfd869cc07e046f723151311d89bdfc682c49dc",  
    "mediaType": "application/vnd.puzzlefs.image.rootfs.v1",  
    "size": 652824  
  },  
  {  
    "digest": "sha256:ade5db0eb90c177afb4e2e617bfff7102d6c2089bd2e7ee9b36b590cebac6ff",  
    "mediaType": "application/vnd.puzzlefs.image.filedata.v1",  
    "size": 78546  
  },  
  {  
    "digest": "sha256:0bc46147d163b17d1f1b8c56422da0cc4993886388526e842ed7442e78c7e22c",  
    "mediaType": "application/vnd.puzzlefs.image.filedata.v1",  
    "size": 109342  
  },  
]
```

PuzzleFS Format



Building a PuzzleFS image

Get the ubuntu:latest image

```
mkdir /tmp/ubuntu-image  
skopeo copy docker://ubuntu:latest oci:/tmp/ubuntu-image:latest  
umoci unpack --rootless --image /tmp/ubuntu-image:latest /tmp/ubuntu
```

Convert it to a PuzzleFS image

```
cargo install puzzlefs  
puzzlefs build /tmp/ubuntu/rootfs /tmp/pfs_ubuntu:eg
```

Creating the LXC container

```
sudo lxc-create --name mycontainer --template oci -- --url oci:/tmp/pfs_ubuntu:eg \  
--no-cache --mount-helper pfs-ovl-mount-helper
```

```
sudo lxc-start --name mycontainer --foreground /bin/bash
```

lxc config for mycontainer:

```
lxc.rootfs.path = dir:/var/lib/lxc/mycontainer/rootfs  
lxc.hook.pre-mount = pfs-ovl-mount-helper mount \  
    ${LXC_ROOTFS_PATH}/../oci:${LXC_NAME} ${LXC_ROOTFS_PATH}
```

Overlay mount helper

```
COMMAND=$1
if [ "$COMMAND" = "mount" ]; then
    OCI_DIR=$2
    MOUNTPOINT=$3
    PFS_MOUNTPOINT=$(mktemp -d)
    UPPERDIR="/tmp/persist"
    mkdir -p $UPPERDIR
    WORKDIR=$(mktemp -d)
    puzzlefs mount $OCI_DIR $PFS_MOUNTPOINT

    echo "mounted PFS on $PFS_MOUNTPOINT"
    mount -t overlay overlay -o lowerdir=$PFS_MOUNTPOINT,upperdir=$UPPERDIR,\
workdir=$WORKDIR $MOUNTPOINT
fi
```

Next steps

Integration with `stacker`

```
hello-stacker:  
  from:  
    type: docker  
    url: docker://zothub.io/tools/busybox:stable  
  run: |  
    mkdir -p /hello-stacker-app  
    echo 'echo "Hello Stacker!"' > /hello-stacker-app/hello.sh  
    chmod +x /hello-stacker-app/hello.sh  
  entrypoint: /hello-stacker-app/hello.sh
```