SCaLE

Build Containers:

The Hard Way, the Easy Way and the Other Ways :-)

the DockerSlim guy

KYLE QUEST

- Created **DockerSlim** (aka SlimToolkit/minToolkit)
- **Good Dockerfiles** project
- Founder/CEO, **Autonomous Layer** (aka AutonomousPlane)
- Founder/CTO, **Slim.AI**
- **https://linktr.ee/kcqon**

Same Size, Different Load Time! What’s Going On?

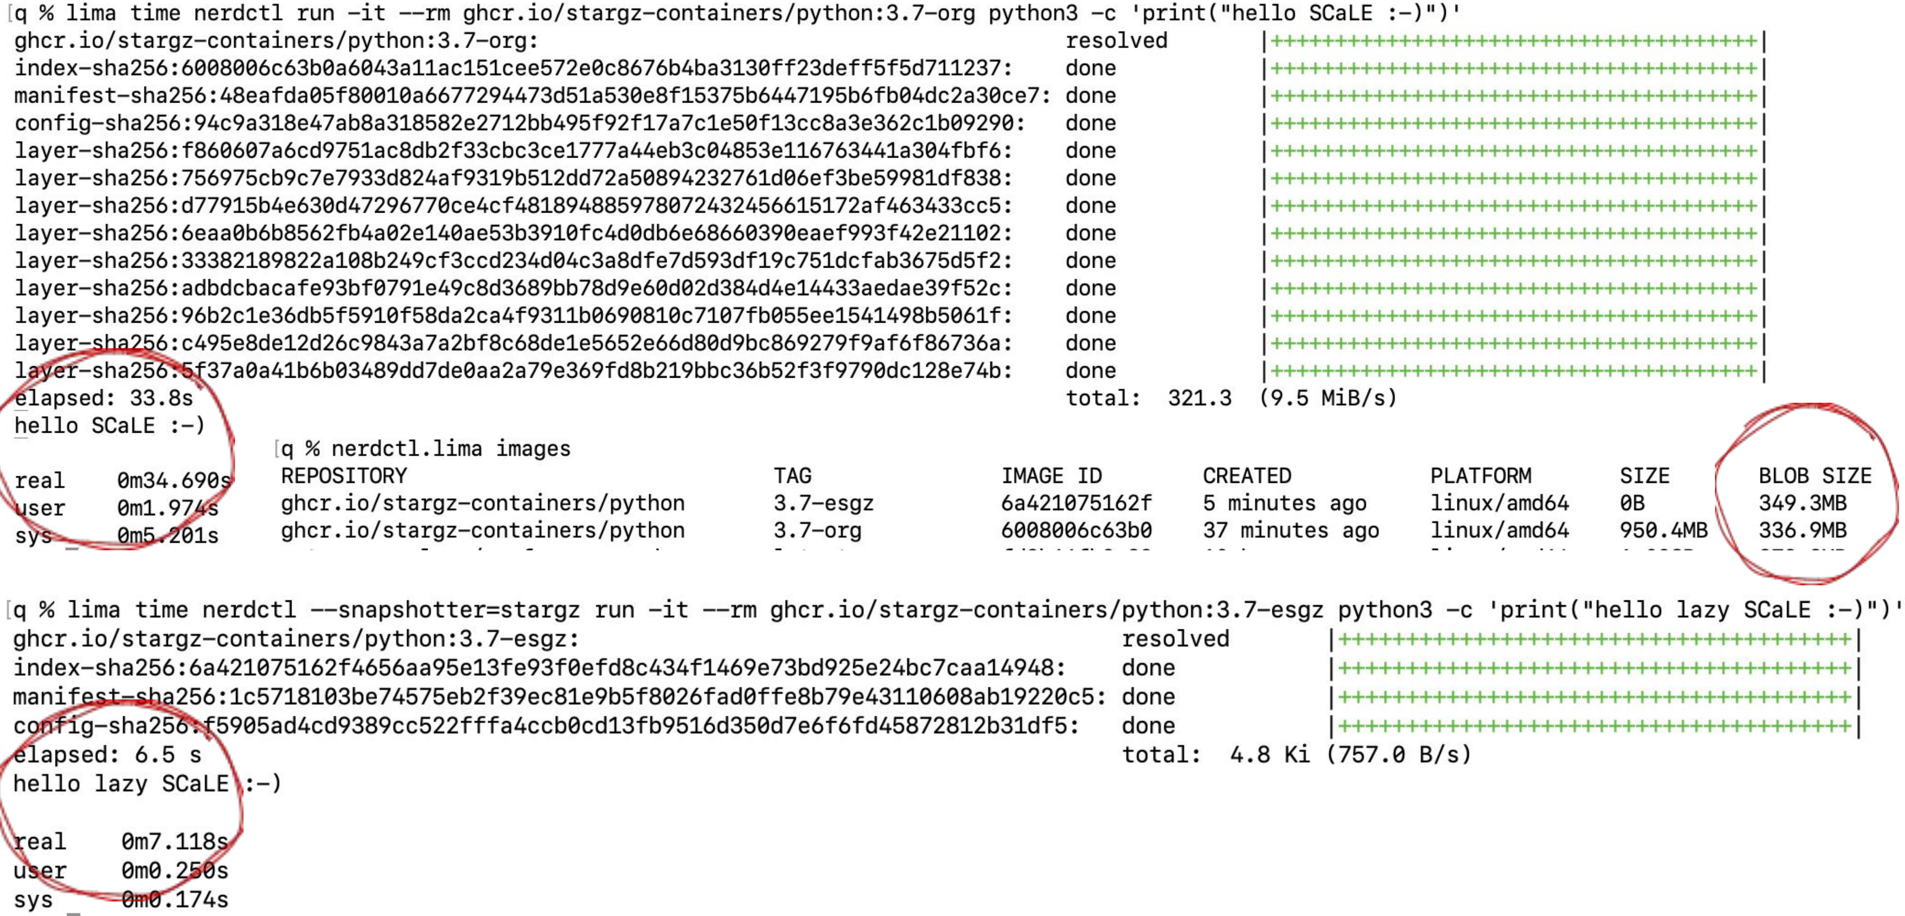

```
q % lima time nerdctl run -it --rm ghcr.io/stargz-containers/python:3.7-org python3 -c 'print("hello SCaLE :-)")'
ghcr.io/stargz-containers/python:3.7-org:                                        resolved       |++++++++++++++++++++++++++++++++++++++|
index-sha256:6008006c63b0a6043a11ac151cee572e0c8676b4ba3130ff23deff5f5d711237:    done           |++++++++++++++++++++++++++++++++++++++|
manifest-sha256:48eafda05f80010a6677294473d51a530e8f15375b6447195b6fb04dc2a30ce7: done           |++++++++++++++++++++++++++++++++++++++|
config-sha256:94c9a318e47ab8a318582e2712bb495f92f17a7c1e50f13cc8a3e362c1b09290:   done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:f860607a6cd9751ac8db2f33cbc3ce1777a44eb3c04853e116763441a304fbf6:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:756975cb9c7e7933d824af9319b512dd72a50894232761d06ef3be59981df838:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:d77915b4e630d47296770ce4cf48189488597807243245661517 2af463433cc5:   done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:6eaa0b6b8562fb4a02e140ae53b3910fc4d0db6e68660390eaef993f42e21102:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:33382189822a108b249cf3ccd234d04c3a8dfe7d593df19c751dcfab3675d5f2:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:adbdcbacafe93bf0791e49c8d3689bb78d9e60d02d384d4e14433aedae39f52c:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:96b2c1e36db5f5910f58da2ca4f9311b0690810c7107fb055ee1541498b5061f:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:c495e8de12d26c9843a7a2bf8c68de1e5652e66d80d9bc869279f9af6f86736a:    done           |++++++++++++++++++++++++++++++++++++++|
layer-sha256:5f37a0a41b6b03489dd7de0aa2a79e369fd8b219bbc36b52f3f9790dc128e74b:    done           |++++++++++++++++++++++++++++++++++++++|
elapsed: 33.8s                                                                    total:  321.3  (9.5 MiB/s)
hello SCaLE :-)

real    0m34.690s
user    0m1.974s
sys     0m5.201s
```

```
q % nerdctl.lima images
REPOSITORY                          TAG          IMAGE ID        CREATED           PLATFORM       SIZE       BLOB SIZE
ghcr.io/stargz-containers/python    3.7-esgz     6a421075162f    5 minutes ago     linux/amd64    0B         349.3MB
ghcr.io/stargz-containers/python    3.7-org      6008006c63b0    37 minutes ago    linux/amd64    950.4MB    336.9MB
```

```
q % lima time nerdctl --snapshotter=stargz run -it --rm ghcr.io/stargz-containers/python:3.7-esgz python3 -c 'print("hello lazy SCaLE :-)")'
ghcr.io/stargz-containers/python:3.7-esgz:                                        resolved       |++++++++++++++++++++++++++++++++++++++|
index-sha256:6a421075162f4656aa95e13fe93f0efd8c434f1469e73bd925e24bc7caa14948:    done           |++++++++++++++++++++++++++++++++++++++|
manifest-sha256:1c5718103be74575eb2f39ec81e9b5f8026fad0ffe8b79e43110608ab19220c5: done           |++++++++++++++++++++++++++++++++++++++|
config-sha256:f5905ad4cd9389cc522fffa4ccb0cd13fb9516d350d7e6f6fd45872812b31df5:   done           |++++++++++++++++++++++++++++++++++++++|
elapsed: 6.5 s                                                                    total:  4.8 Ki (757.0 B/s)
hello lazy SCaLE :-)

real    0m7.118s
user    0m0.250s
sys     0m0.174s
```

Let’s Start the Journey! Let’s Create a Container!

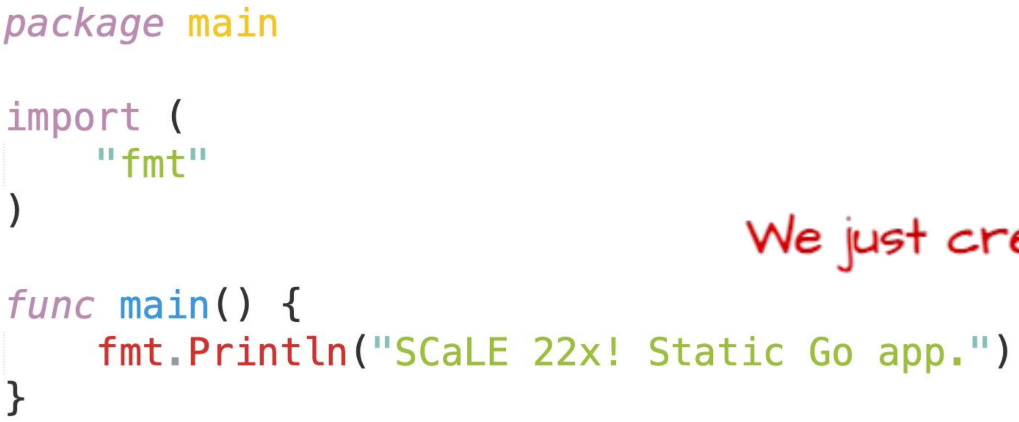

```go
package main

import (
    "fmt"
)

func main() {
    fmt.Println("SCaLE 22x! Static Go app.")
}
```

We just created a container the easy way :-)

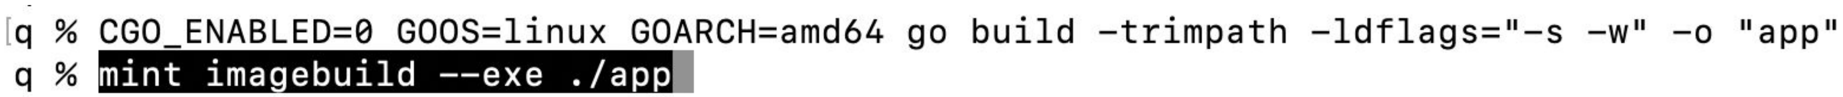

```
[q % CGO_ENABLED=0 GOOS=linux GOARCH=amd64 go build -trimpath -ldflags="-s -w" -o "app"
 q % mint imagebuild --exe ./app
```

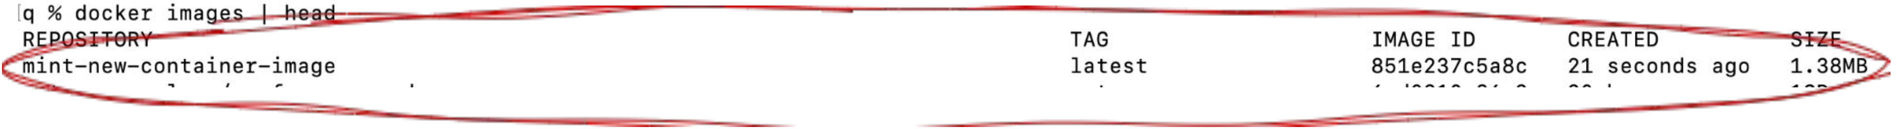

```
[q % docker images | head
REPOSITORY                                            TAG             IMAGE ID       CREATED          SIZE
mint-new-container-image                              latest          851e237c5a8c   21 seconds ago   1.38MB
```

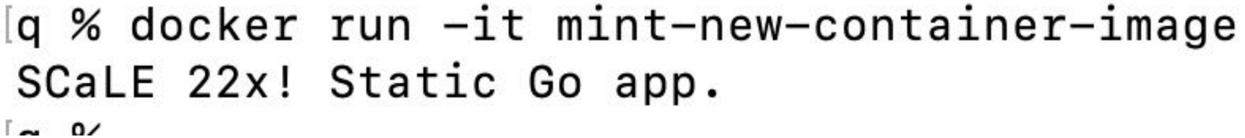

```
[q % docker run -it mint-new-container-image
SCaLE 22x! Static Go app.
```

What’s in the Container? OCI Image Spec TLDR

```
q % docker save -o conference-node-app.slim_image_data.tar conference-node-app.slim:latest
q % tar -tf conference-node-app.slim_image_data.tar | tree --fromfile
.
├── blobs
│   └── sha256
│       ├── 3d04aa34e1a18e3eef316df4888fe2bc3b3f24b4b294d59ea36ca62c4ebfd728
│       ├── 4402a79d98c4362b8fec4ef30e693fe47f6f87c8ca4cc5bb28ebfd5b14239978
│       ├── 496163609c48cc0d6abe558dd424d5fd87ec32d6874f3238e55634d9a81de553
│       └── 6703bc44da4e0c2ff0579a5a0b518b98f2bde74648f9a655ebe7e75998b1f09f
├── index.json
├── manifest.json
├── oci-layout
└── repositories
```

```
q % cat index.json | jq
{
  "schemaVersion": 2,
  "mediaType": "application/vnd.oci.image.index.v1+json",
  "manifests": [
    {
      "mediaType": "application/vnd.oci.image.manifest.v1+json",
      "digest": "sha256:4402a79d98c4362b8fec4ef30e693fe47f6f87c8ca4cc5bb28ebfd5b14239978",
      "size": 402,
      "annotations": {
        "io.containerd.image.name": "docker.io/library/conference-node-app.slim:latest",
        "org.opencontainers.image.ref.name": "latest"
      }
    }
  ]
}
```

```
q % cat blobs/sha256/4402a79d98c4362b8fec4ef30e693fe47f6f87c8ca4cc5bb28ebfd5b14239978 | jq
{
  "schemaVersion": 2,
  "mediaType": "application/vnd.oci.image.manifest.v1+json",
  "config": {
    "mediaType": "application/vnd.oci.image.config.v1+json",
    "digest": "sha256:6703bc44da4e0c2ff0579a5a0b518b98f2bde74648f9a655ebe7e75998b1f09f",
    "size": 2457
  },
  "layers": [
    {
      "mediaType": "application/vnd.oci.image.layer.v1.tar",
      "digest": "sha256:496163609c48cc0d6abe558dd424d5fd87ec32d6874f3238e55634d9a81de553",
      "size": 94071808
    }
  ]
}
```

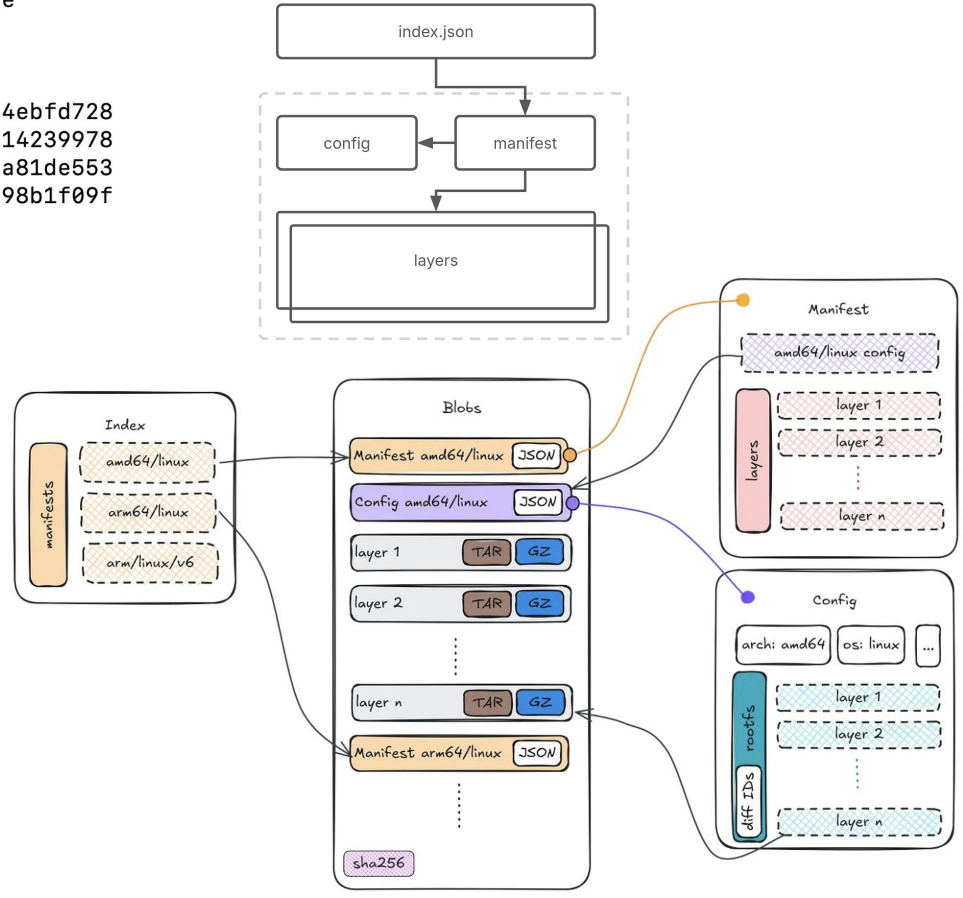

Build Containers By Hand… How Hard Is It?

- Create the **layer archive** (tar + gzip)
  - Generate SHA256 for it
  - Calculate layer archive file size
  - Rename layer with SHA256
  - Generate uncompressed SHA256 for layer archive to use as DiffID in config
- Create the (runtime param) **config file** (which parts do we really need?) ,
  - Generate SHA256 for it
  - Calculate config file size
  - Rename config file with SHA256
- Create the **image manifest**
  - Use config file SHA256 and size
  - Use layer archive SHA256 and size
  - Generate SHA256 for image manifest
  - Calculate image manifest file size
  - rename image manifest file with SHA256
- Create **index.json**
  - Use image manifest file SHA256 and size
- Create the static **oci-layout** file

Build Containers By Hand… Can We Cheat?

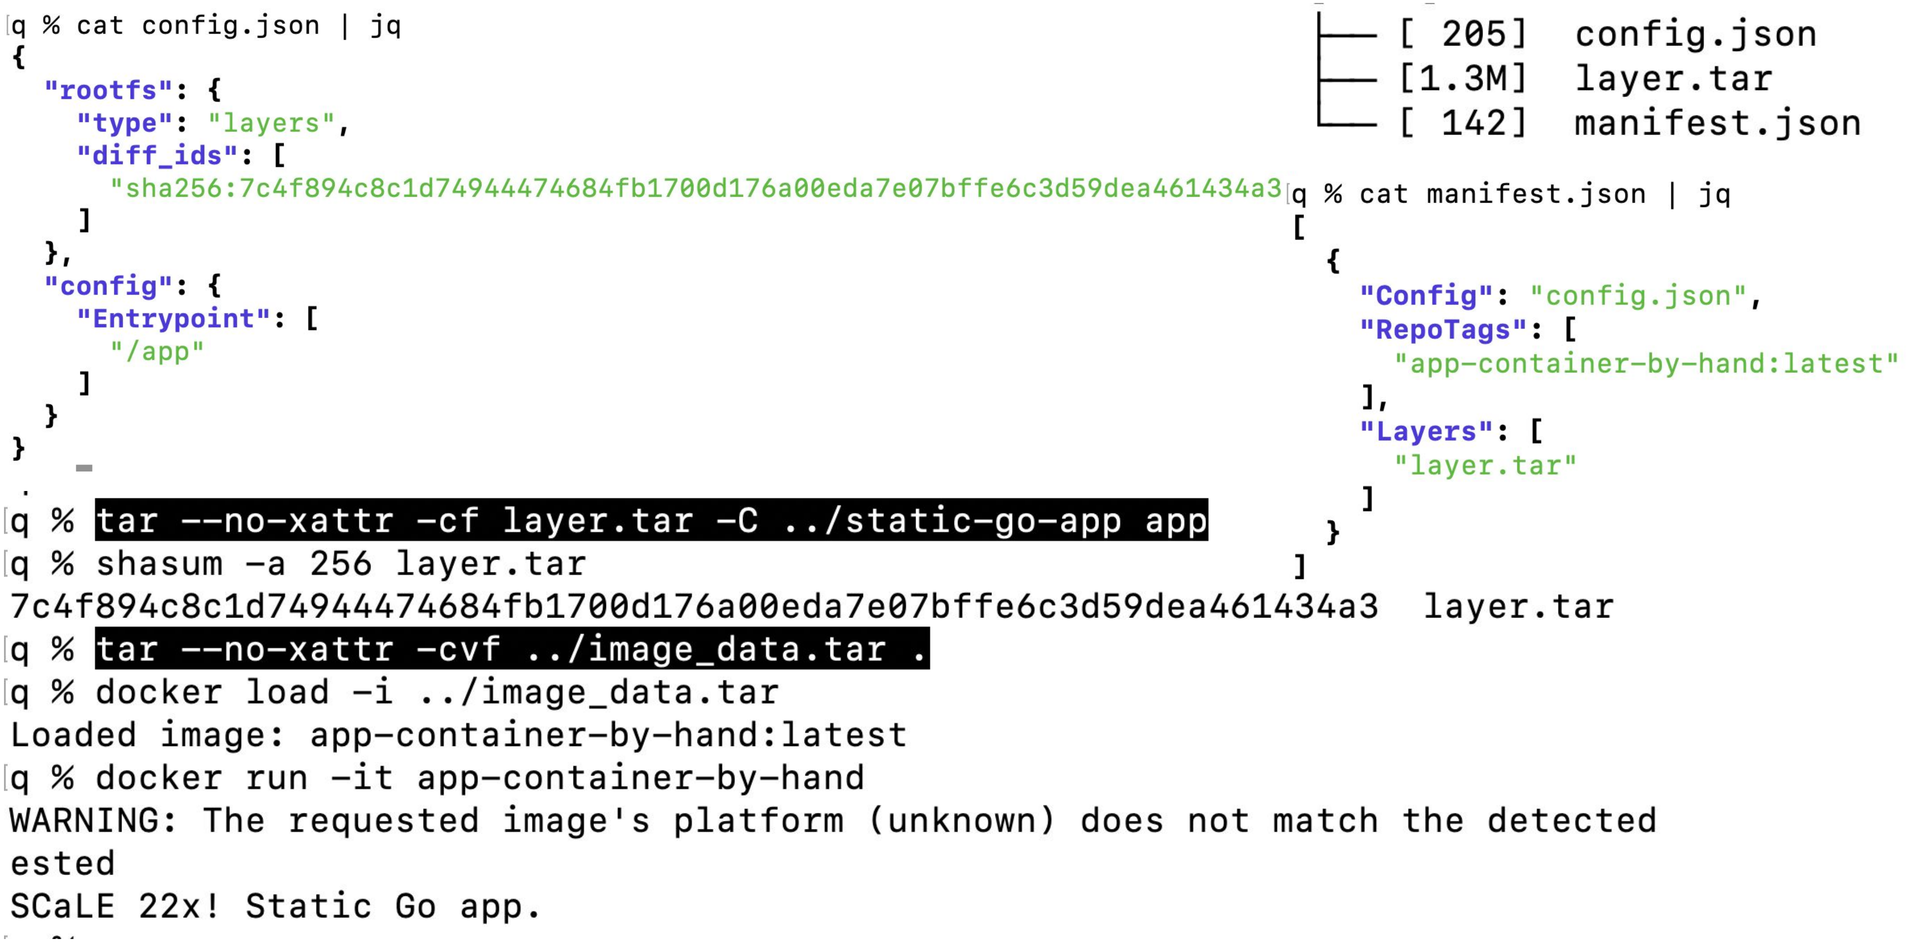

```
q % cat config.json | jq
{
  "rootfs": {
    "type": "layers",
    "diff_ids": [
      "sha256:7c4f894c8c1d74944474684fb1700d176a00eda7e07bffe6c3d59dea461434a3
    ]
  },
  "config": {
    "Entrypoint": [
      "/app"
    ]
  }
}
```

```
├── [ 205]  config.json
├── [1.3M]  layer.tar
└── [ 142]  manifest.json
```

```
q % cat manifest.json | jq
[
  {
    "Config": "config.json",
    "RepoTags": [
      "app-container-by-hand:latest"
    ],
    "Layers": [
      "layer.tar"
    ]
  }
]
```

```
q % tar --no-xattr -cf layer.tar -C ../static-go-app app
q % shasum -a 256 layer.tar
7c4f894c8c1d74944474684fb1700d176a00eda7e07bffe6c3d59dea461434a3  layer.tar
q % tar --no-xattr -cvf ../image_data.tar .
q % docker load -i ../image_data.tar
Loaded image: app-container-by-hand:latest
q % docker run -it app-container-by-hand
WARNING: The requested image's platform (unknown) does not match the detected
ested
SCaLE 22x! Static Go app.
```

Enough Doing it the Hard Way… Let’s Use Tools!

We’ll cover:

- **Buildpacks**
- **NixPacks**
- **Nixery**
- **Kontain.me**
- **MinToolkit** / Mint (aka DockerSlim) - already covered some of it

But we’ll skip:

- **BuildKit** - You are using it when you are using Docker (custom frontends deserve their own talk :-))
- **Nix** - has several ways to build containers (Nix is its own world, HB2E)
- **ko** - Golang only (language specific)
- **JIB** - Java only (language specific)
- **Kaniko** - to do container builds in Kubernetes
- **Podman/Buildah** - Same as Docker, but a bit buggy / less compatible :-)

Building Images with Buildpacks

Buildpacks.io

- https://buildpacks.io
- **Builds application source and generates a container**
- **PaaS origins (Heroku), needs a Procfile or special config**

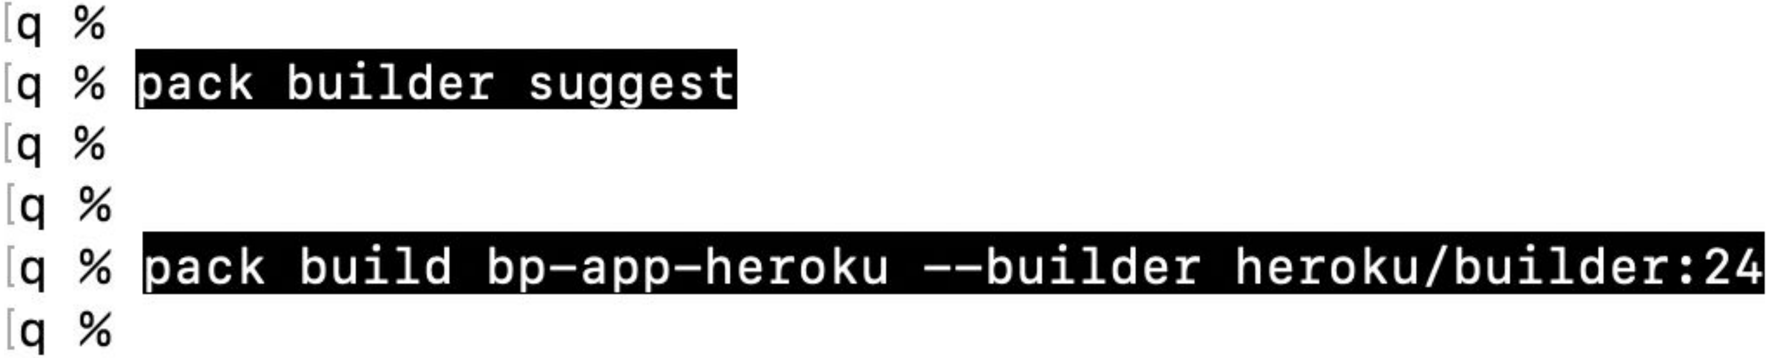

```
[q %
[q % pack builder suggest
[q %
[q %
[q % pack build bp-app-heroku --builder heroku/builder:24
[q %
```

Heroku Buildpacks

buildpacks

paketo
buildpacks

Building Images with NixPacks

- https://nixpacks.com (made by Railway)
- **App source + Nix packages + Docker = Image**
- **Buildpacks that don’t suck :-)**
- **Also generates a Dockerfile (if you want, but it has Nix)**

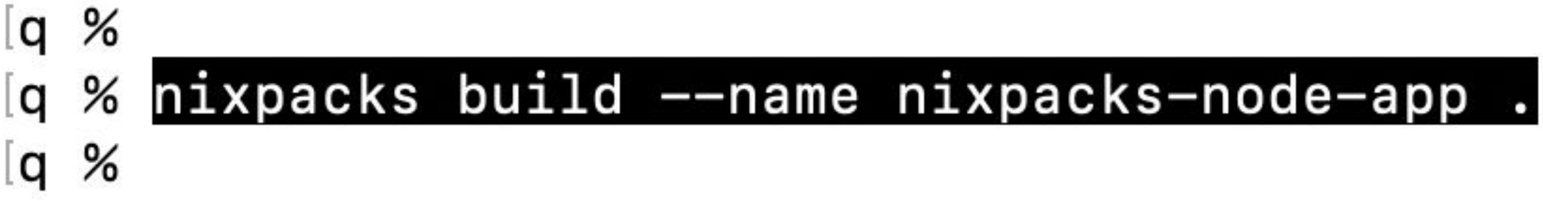

```
[q %
[q % nixpacks build --name nixpacks-node-app .
[q %
```

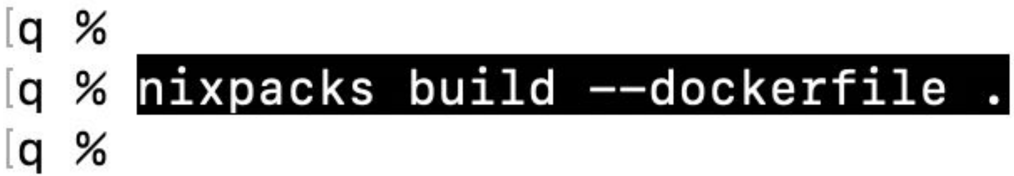

```
[q %
[q % nixpacks build --dockerfile .
[q %
```

Building Images with Nixery

- https://nixery.dev
- **Ad-hoc container image registry that provides packages from the Nix package manager**
- Find packages here: https://search.nixos.org/packages

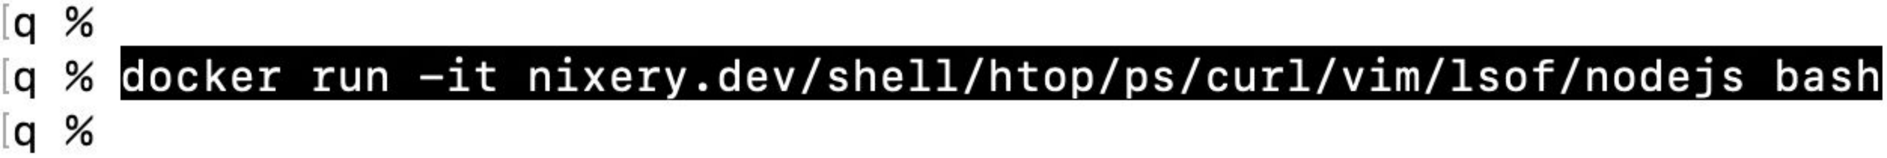

```
[q %
[q % docker run -it nixery.dev/shell/htop/ps/curl/vim/lsof/nodejs bash
[q %
```

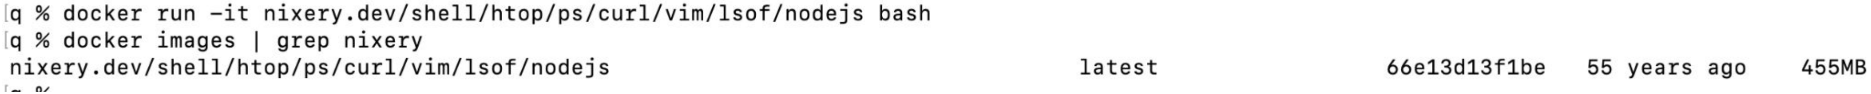

```
[q % docker run -it nixery.dev/shell/htop/ps/curl/vim/lsof/nodejs bash
[q % docker images | grep nixery
nixery.dev/shell/htop/ps/curl/vim/lsof/nodejs                            latest                66e13d13f1be   55 years ago   455MB
```

Building Images with Kontain.me

- https://github.com/imjasonh/kontain.me/tree/main/cmd/apko
- **Like Nixery, but with Wolfi libc-based apk packages**
- **Used to use Alpine musl-based apk package**

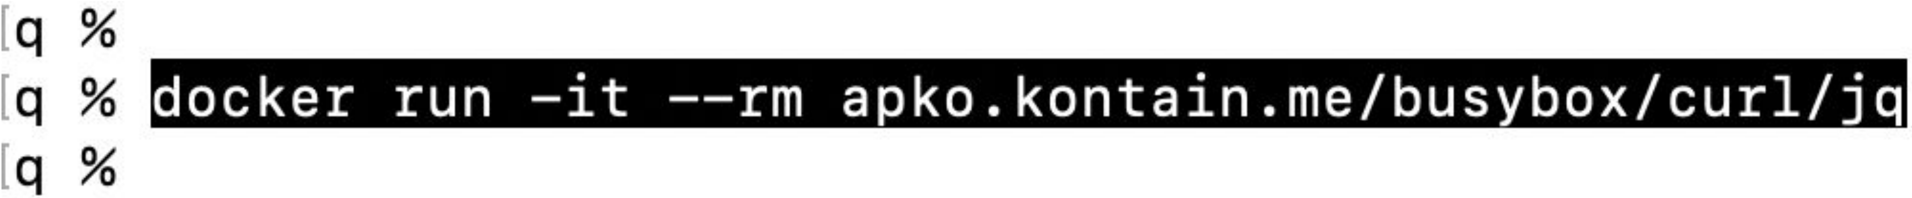

```
[q %
[q % docker run -it --rm apko.kontain.me/busybox/curl/jq
[q %
```

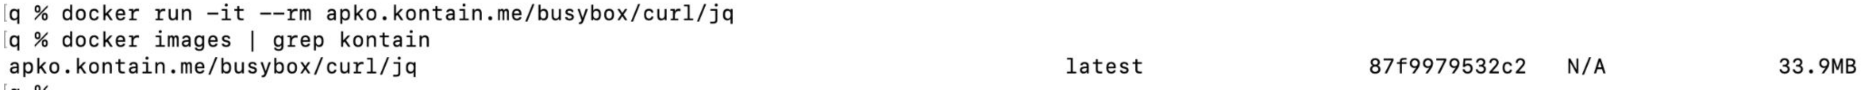

```
[q % docker run -it --rm apko.kontain.me/busybox/curl/jq
[q % docker images | grep kontain
apko.kontain.me/busybox/curl/jq                                latest            87f9979532c2   N/A          33.9MB
[q %
```

How Can We Build Containers That Load Fast?

eStargz

aws soci

Overlaybd

NYDUS

Lazy Loaded Containers with eStargz

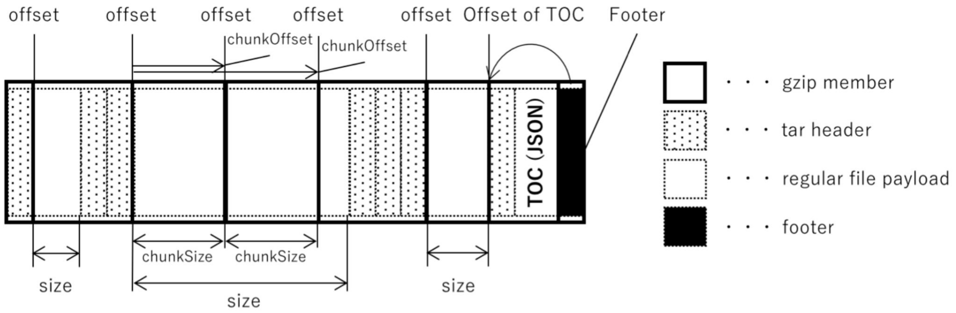

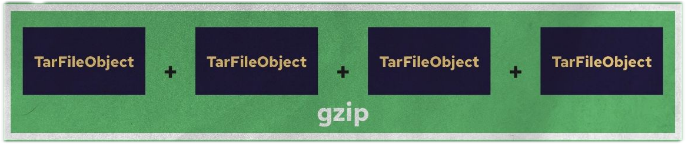

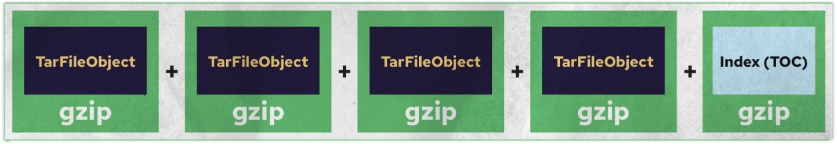

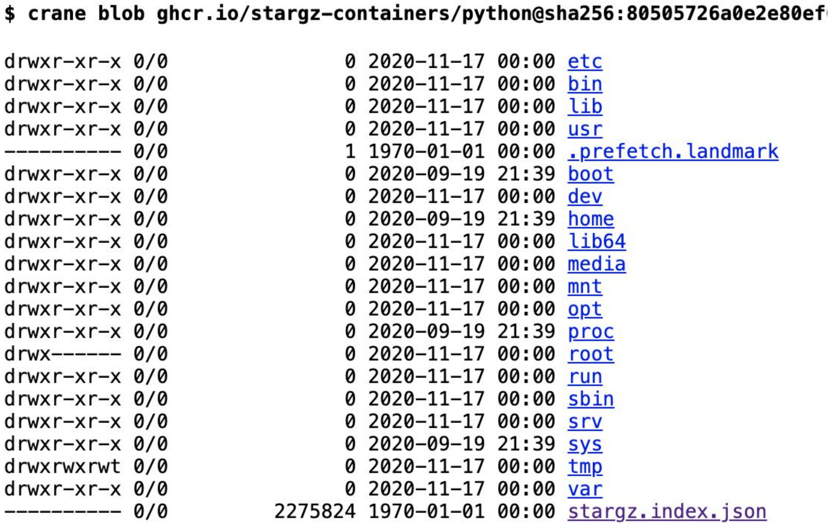

```
$ crane blob ghcr.io/stargz-containers/python@sha256:80505726a0e2e80ef

drwxr-xr-x 0/0                 0 2020-11-17 00:00 etc
drwxr-xr-x 0/0                 0 2020-11-17 00:00 bin
drwxr-xr-x 0/0                 0 2020-11-17 00:00 lib
drwxr-xr-x 0/0                 0 2020-11-17 00:00 usr
---------- 0/0                 1 1970-01-01 00:00 .prefetch.landmark
drwxr-xr-x 0/0                 0 2020-09-19 21:39 boot
drwxr-xr-x 0/0                 0 2020-11-17 00:00 dev
drwxr-xr-x 0/0                 0 2020-09-19 21:39 home
drwxr-xr-x 0/0                 0 2020-11-17 00:00 lib64
drwxr-xr-x 0/0                 0 2020-11-17 00:00 media
drwxr-xr-x 0/0                 0 2020-11-17 00:00 mnt
drwxr-xr-x 0/0                 0 2020-11-17 00:00 opt
drwxr-xr-x 0/0                 0 2020-09-19 21:39 proc
drwx------ 0/0                 0 2020-11-17 00:00 root
drwxr-xr-x 0/0                 0 2020-11-17 00:00 run
drwxr-xr-x 0/0                 0 2020-11-17 00:00 sbin
drwxr-xr-x 0/0                 0 2020-11-17 00:00 srv
drwxr-xr-x 0/0                 0 2020-09-19 21:39 sys
drwxrwxrwt 0/0                 0 2020-11-17 00:00 tmp
drwxr-xr-x 0/0                 0 2020-11-17 00:00 var
---------- 0/0           2275824 1970-01-01 00:00 stargz.index.json
```

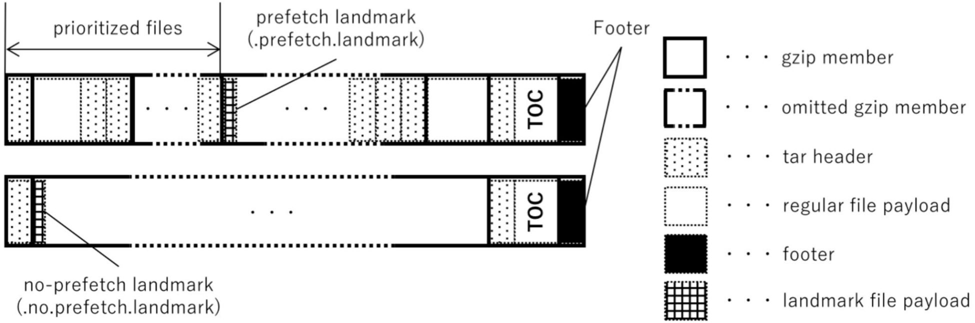

Building Lazy Loaded eStargz Containers

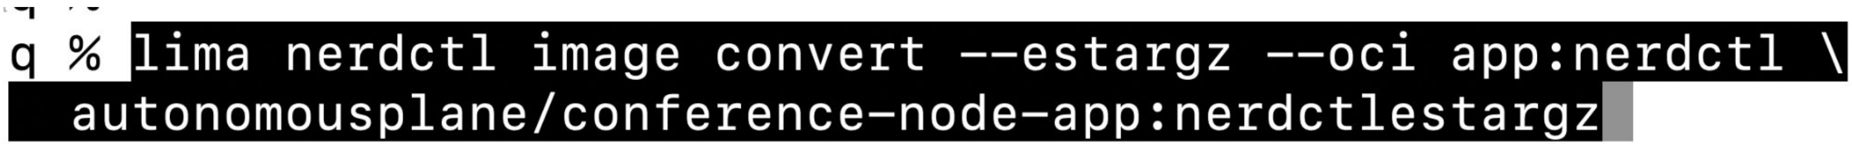

```
q % lima nerdctl image convert --estargz --oci app:nerdctl \
  autonomousplane/conference-node-app:nerdctlestargz
```

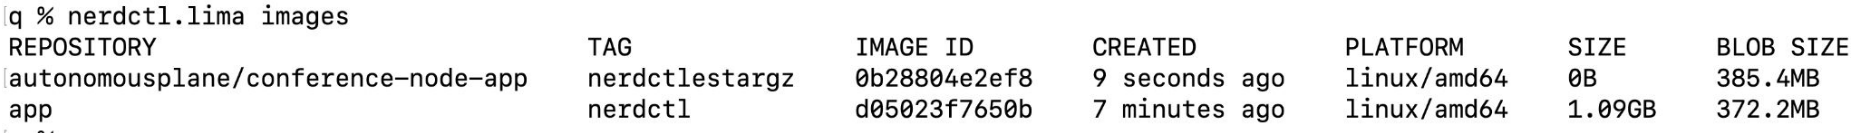

```
q % nerdctl.lima images
REPOSITORY                             TAG                IMAGE ID        CREATED           PLATFORM       SIZE       BLOB SIZE
autonomousplane/conference-node-app    nerdctlestargz     0b28804e2ef8    9 seconds ago     linux/amd64    0B         385.4MB
app                                    nerdctl            d05023f7650b    7 minutes ago     linux/amd64    1.09GB     372.2MB
```

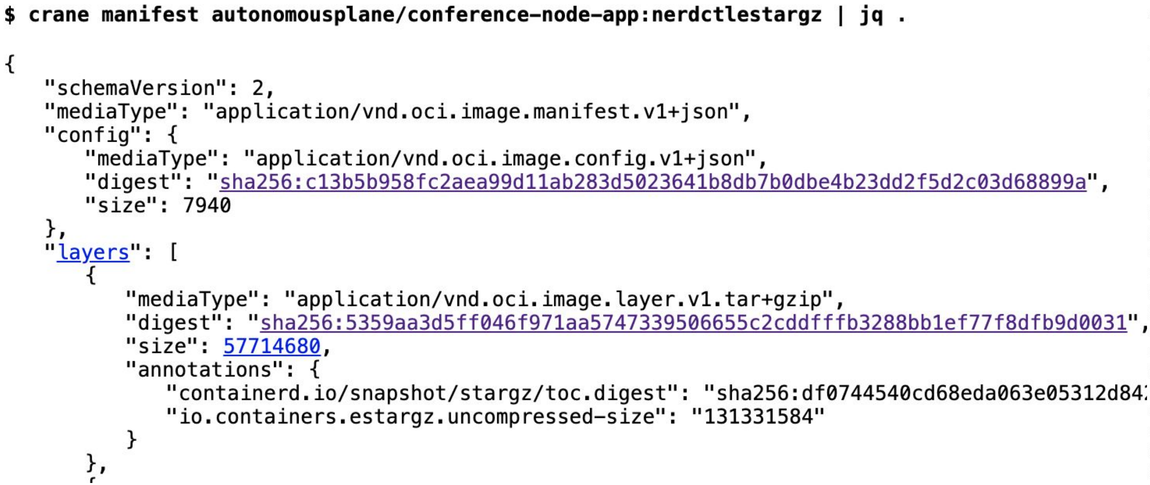

```
$ crane manifest autonomousplane/conference-node-app:nerdctlestargz | jq .

{
   "schemaVersion": 2,
   "mediaType": "application/vnd.oci.image.manifest.v1+json",
   "config": {
      "mediaType": "application/vnd.oci.image.config.v1+json",
      "digest": "sha256:c13b5b958fc2aea99d11ab283d5023641b8db7b0dbe4b23dd2f5d2c03d68899a",
      "size": 7940
   },
   "layers": [
      {
         "mediaType": "application/vnd.oci.image.layer.v1.tar+gzip",
         "digest": "sha256:5359aa3d5ff046f971aa5747339506655c2cddfffb3288bb1ef77f8dfb9d0031",
         "size": 57714680,
         "annotations": {
            "containerd.io/snapshot/stargz/toc.digest": "sha256:df0744540cd68eda063e05312d84
            "io.containers.estargz.uncompressed-size": "131331584"
         }
      },
```

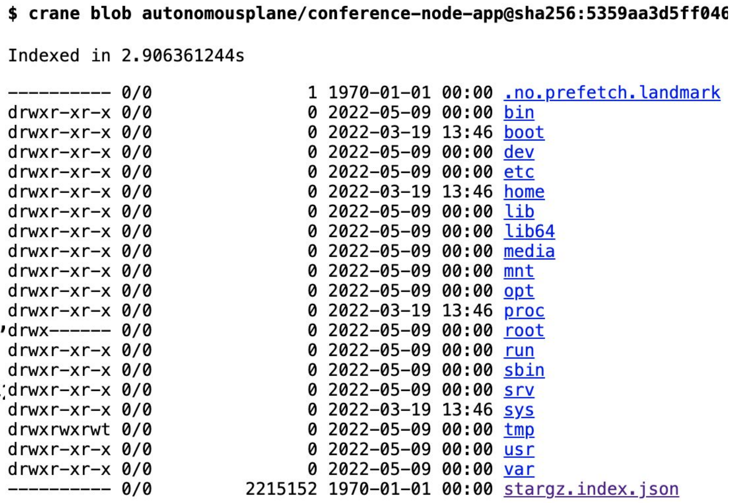

```
$ crane blob autonomousplane/conference-node-app@sha256:5359aa3d5ff046

Indexed in 2.906361244s

---------- 0/0             1 1970-01-01 00:00 .no.prefetch.landmark
drwxr-xr-x 0/0             0 2022-05-09 00:00 bin
drwxr-xr-x 0/0             0 2022-03-19 13:46 boot
drwxr-xr-x 0/0             0 2022-05-09 00:00 dev
drwxr-xr-x 0/0             0 2022-05-09 00:00 etc
drwxr-xr-x 0/0             0 2022-03-19 13:46 home
drwxr-xr-x 0/0             0 2022-05-09 00:00 lib
drwxr-xr-x 0/0             0 2022-05-09 00:00 lib64
drwxr-xr-x 0/0             0 2022-05-09 00:00 media
drwxr-xr-x 0/0             0 2022-05-09 00:00 mnt
drwxr-xr-x 0/0             0 2022-05-09 00:00 opt
drwxr-xr-x 0/0             0 2022-03-19 13:46 proc
drwx------ 0/0             0 2022-05-09 00:00 root
drwxr-xr-x 0/0             0 2022-05-09 00:00 run
drwxr-xr-x 0/0             0 2022-05-09 00:00 sbin
drwxr-xr-x 0/0             0 2022-05-09 00:00 srv
drwxr-xr-x 0/0             0 2022-03-19 13:46 sys
drwxrwxrwt 0/0             0 2022-05-09 00:00 tmp
drwxr-xr-x 0/0             0 2022-05-09 00:00 usr
drwxr-xr-x 0/0             0 2022-05-09 00:00 var
---------- 0/0       2215152 1970-01-01 00:00 stargz.index.json
```

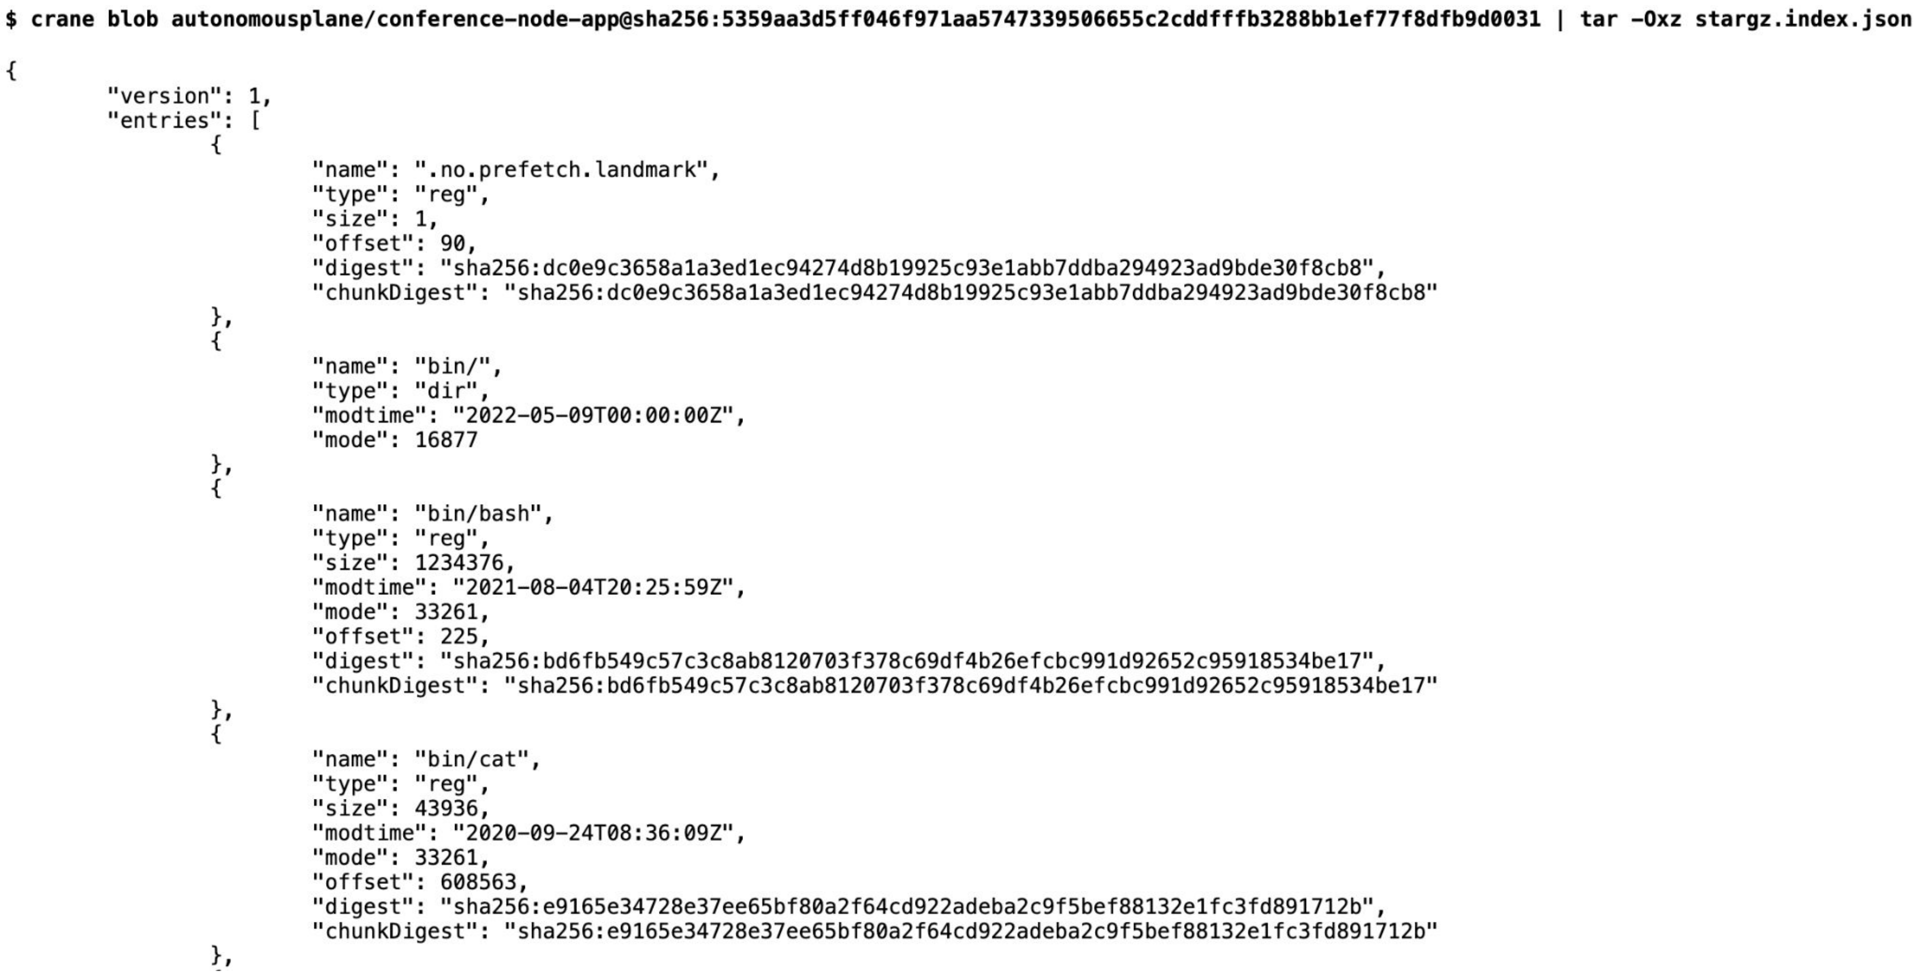

```
$ crane blob autonomousplane/conference-node-app@sha256:5359aa3d5ff046f971aa5747339506655c2cddfffb3288bb1ef77f8dfb9d0031 | tar -Oxz stargz.index.json

{
        "version": 1,
        "entries": [
                {
                        "name": ".no.prefetch.landmark",
                        "type": "reg",
                        "size": 1,
                        "offset": 90,
                        "digest": "sha256:dc0e9c3658a1a3ed1ec94274d8b19925c93e1abb7ddba294923ad9bde30f8cb8",
                        "chunkDigest": "sha256:dc0e9c3658a1a3ed1ec94274d8b19925c93e1abb7ddba294923ad9bde30f8cb8"
                },
                {
                        "name": "bin/",
                        "type": "dir",
                        "modtime": "2022-05-09T00:00:00Z",
                        "mode": 16877
                },
                {
                        "name": "bin/bash",
                        "type": "reg",
                        "size": 1234376,
                        "modtime": "2021-08-04T20:25:59Z",
                        "mode": 33261,
                        "offset": 225,
                        "digest": "sha256:bd6fb549c57c3c8ab8120703f378c69df4b26efcbc991d92652c95918534be17",
                        "chunkDigest": "sha256:bd6fb549c57c3c8ab8120703f378c69df4b26efcbc991d92652c95918534be17"
                },
                {
                        "name": "bin/cat",
                        "type": "reg",
                        "size": 43936,
                        "modtime": "2020-09-24T08:36:09Z",
                        "mode": 33261,
                        "offset": 608563,
                        "digest": "sha256:e9165e34728e37ee65bf80a2f64cd922adeba2c9f5bef88132e1fc3fd891712b",
                        "chunkDigest": "sha256:e9165e34728e37ee65bf80a2f64cd922adeba2c9f5bef88132e1fc3fd891712b"
                },
```

Thank You!

Useful References

- **https://github.com/mintoolkit/mint - Current DockerSlim repo (old is for support only)**
- https://buildpacks.io
- https://nixpacks.com
- https://nixery.dev
- https://lima-vm.io - Nerdctl CLI and Lima VM for Docker/Containerd
- **https://oci.dag.dev** - Explore container image registry data
- **https://ima.ge.cx** - Explore container images online