



CONTAINERCON



THE LINUX FOUNDATION

OPEN SOURCE SUMMIT

NORTH AMERICA

# Container Registries, No Longer Just for Containers!

#osummit @vbatts@fosstodon.org



# Howdy



```
$> whoami  
vbatts  
$> id -Gn  
devel azure flatcar kinvolk redhat oci appc  
docker golang slackware ruby
```

# The situation

- Cloud native!
- `image: docker.io/nginx:stable`



# Hidden in the reference

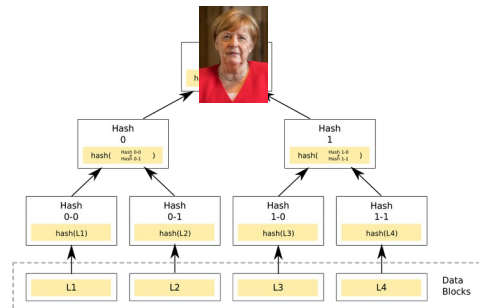
- ♥ Better yet..
- ♥ `image:`  
`docker.io/nginx@sha256:b1a2c7bcc61be621eae24851a976179bfb72591e43c1fb340f7497ff72128ff`
- ♥ The API for moving that image
- ♥ The format of that packaged image
- ♥ A crux of this whole deal

# Some Background ...

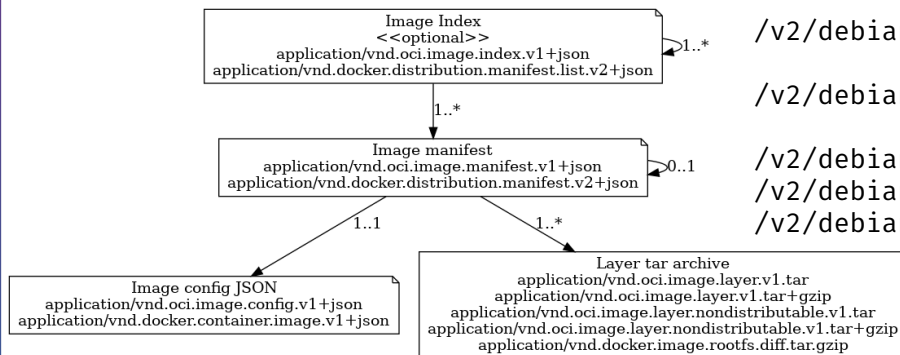
- ♥ Evolution of containers, as we know them, was *fast*
- ♥ Standards/Specs were needed
- ♥ Opening the governance,
  - ♣ and releasing with consensus took time
- ♥ First it was runtime
  - ♣ Then image format
    - ♣ And *then* it was distribution
- ♥ The OCI mediaType adoption is *still* happening today

# Container images?

- ♥ Merkle trees
- ♥ Container image pieces
- ♥ Content addressable registry



Source for [Merkle](#) and [Merkel](#)



/v2/debian/manifests/testing

/v2/debian/manifests/sha256:bc75dee24f5f0c4da9998b7214ae5f02c4

/v2/debian/blobs/sha256:e26bc5d581e3d496366ab7f5c68d21f494979a

/v2/debian/blobs/sha256:5cf1ea968ffbb3fe3e74e126b87e54ebec4b34

/v2/debian/blobs/sha256:37aea29754268aad209de0e749656a0b9e069c

# Deployments are more than just container

- ♥ Evolutions of the ecosystem
  - ♣ Deployment configs (helm charts)
  - ♣ Signatures
  - ♣ SBOMs
  - ♣ New types of byte code

# Needs of the Ecosystem

- ♥ Storing related state and configs
- ♥ Storing related metadata
- ♥ Not reinventing the wheel
- ♥ Connecting the dots



# It's not so straight forward

- ♥ Ship has sailed on some design choices
- ♥ A new mediaType would find us in the same migration
- ♥ The “solution”?: continue with the existing container image [for now] ..



# Hands on

- ♥ Various tools to experiment with
  - ♣ crane and oras
  - ♣ regclient and skopeo too
- ♥
- ♥
- ♥ Fun, though most of the following ideally is transparent to the user!

# The way it's been ...

```
vbatts@possibly:~/talk/scratch$ docker build -t ossna.azurecr.io/scratch .
Sending build context to Docker daemon 3.072kB
Step 1/2 : FROM scratch
---->
Step 2/2 : ADD hi.txt / (Preview)
----> 1b508df0ef11
Successfully built 1b508df0ef11
Successfully tagged ossna.azurecr.io/scratch:latest
vbatts@possibly:~/talk/scratch$ docker push ossna.azurecr.io/scratch
Using default tag: latest
The push refers to repository [ossna.azurecr.io/scratch]
077024883d1b: Pushed
latest: digest: sha256:a5b9c2190ff74f728374502126d260035868e1623d2f44120725b9fcf7530d95 size: 524
vbatts@possibly:~/talk/scratch$
```

```
vbatts@possibly:~/talk/scratch$ crane manifest ossna.azurecr.io/scratch | jq '.layers[].mediaType'
"application/vnd.docker.image.rootfs.diff.tar.gzip"
```

# A bit more direct

```
vbatts@possibly:~/talk/scratch$ oras push ossna.azurecr.io/hi-txt:latest hi.txt:application/vnd.me.hi
Exists 315a437df83e hi.txt
Pushed [registry] ossna.azurecr.io/hi-txt:latest
Digest: sha256:ce51f54fdac067a16c33fcfa2cb0ac9d836bb1f03191f9cbfdb0dcf55116afb8
vbatts@possibly:~/talk/scratch$ crane manifest ossna.azurecr.io/hi-txt:latest | jq '.layers[].mediaType'
"application/vnd.me.hi"
```

# A bit more direct

```
vbatts@possibly:~$ crane --insecure manifest 0:1234/hi-txt:latest | jq
{
  "schemaVersion": 2,
  "mediaType": "application/vnd.oci.image.manifest.v1+json",
  "config": {
    "mediaType": "application/vnd.unknown.config.v1+json",
    "digest": "sha256:44136fa355b3678a1146ad16f7e8649e94fb4fc21fe77e8310c060f61caaff8a",
    "size": 28
  },
  "layers": [
    {
      "mediaType": "application/vnd.me.hi",
      "digest": "sha256:315a437df83ef78cd147d4f270eef2d490dec4c4e12ea84886d3d45c57d3dd71",
      "size": 4,
      "annotations": {
        "org.opencontainers.image.title": "hi.txt"
      }
    }
  ],
  "annotations": {
    "org.opencontainers.image.created": "2023-05-10T15:18:02Z"
  }
}
```

# A bit more direct

```
vbatts@possibly:~$ crane push ./talk/debian-testing ossna.azurecr.io/debian:testing
2023/05/10 09:37:51 existing blob: sha256:5cf1ea968ffbb3fe3e74e126b87e54ebee4b34f3ffcf7b2ea6bbec9669161693
2023/05/10 09:37:51 existing blob: sha256:37aea29754268aad209de0e749656a0b9e069ce7d3c82ee5aa54eb78b5009e50
2023/05/10 09:37:51 existing blob: sha256:e26bc5d581e3d496366ab7f5c68d21f494979a28934dd5b9335c6feacc2bb8fa
2023/05/10 09:37:51 ossna.azurecr.io/debian:testing: digest: sha256:bc75dee24f5f0c4da9998b7214ae5f02c4c5d55e03d9c021379cf8332dd43f37 size: 736
ossna.azurecr.io/debian@sha256:bc75dee24f5f0c4da9998b7214ae5f02c4c5d55e03d9c021379cf8332dd43f37
```

```
"subject": {
  "mediaType": "application/vnd.oci.image.manifest.v1+json",
  "size": 736,
  "digest": "sha256:bc75dee24f5f0c4da9998b7214ae5f02c4c5d55e03d9c021379cf8332dd43f37"
},
```

```
vbatts@possibly:~/talk/scratch$ crane edit manifest ossna.azurecr.io/hi-txt:latest
2023/05/10 09:45:59 ossna.azurecr.io/hi-txt:latest: digest: sha256:ef775c92ad78d3f277eafeda81d354aa9a4f8bf4bdcd874eb362b5291a2451ae size: 669
ossna.azurecr.io/hi-txt:latest
```

# Examples of arbitrary content?

- ♥ Deployment configs (helm charts)
- ♥ Signatures (sigstore and notation)
- ♥ SBOMs
- ♥ types of byte code/packages
  - ♣ WASM (fermyon spin, and more)
  - ♣ BPF (bumblebee)
- ♥ And so many more I'm certainly missing

# referrers

[github.com/opencontainers/distribution-spec/blob/main/spec.md](https://github.com/opencontainers/distribution-spec/blob/main/spec.md)

- > Referrers List: a list of manifests with a subject relationship to a specified digest. The referrers list is generated with a query to a registry.

```
"subject": {  
  "mediaType": "application/vnd.oci.image.manifest.v1+json",  
  "size": 736,  
  "digest": "sha256:bc75dee24f5f0c4da9998b7214ae5f02c4c5d55e03d9c021379cf8332dd43f37"  
},
```

```
vbatts@possibly:~$ crane --insecure edit manifest 0:1234/hi-txt:latest  
2023/05/10 08:56:08 updating fallback tag sha256-bc75dee24f5f0c4da9998b7214ae5f02c4c5d55e03d9c021379cf83  
32dd43f37 with new referrer  
2023/05/10 08:56:08 0:1234/hi-txt:sha256-bc75dee24f5f0c4da9998b7214ae5f02c4c5d55e03d9c021379cf8332dd43f3  
7: digest: sha256:5ff7b6d56b188ad880b1c01c88d0c5546224991f36841a87ee0f7e70d3e72edf size: 296  
2023/05/10 08:56:08 0:1234/hi-txt:latest: digest: sha256:e7572f1ae499993c39fc75c7d443db3d6168fcb1996cec7  
bd25f3c2fe7357b53 size: 807  
0:1234/hi-txt:latest
```



# But why?

- ♥ K.I.S.S.
- ♥ Use the content addressable store
  - ♣ And enable extensibility
- ♥ So much debate and deliberation
  - ♣ “with much power comes much responsibility”

# So many links

- ▼ [github.com/oras-project/oras](https://github.com/oras-project/oras)
- ▼ [github.com/google/go-containerregistry/tree/main/cmd/crane](https://github.com/google/go-containerregistry/tree/main/cmd/crane)
- ▼ [github.com/regclient/regclient](https://github.com/regclient/regclient)
- ▼ [github.com/containerd/skopeo](https://github.com/containerd/skopeo)
- ▼ [github.com/sigstore/cosign](https://github.com/sigstore/cosign)
- ▼ [github.com/notaryproject/notation](https://github.com/notaryproject/notation)
- ▼ [helm.sh/docs/topics/registries](https://helm.sh/docs/topics/registries)
- ▼ [github.com/fermyon/spin](https://github.com/fermyon/spin)
- ▼ [github.com/solo-io/bumblebee](https://github.com/solo-io/bumblebee)
- ▼ [techcommunity.microsoft.com](https://techcommunity.microsoft.com) - Azure Container Registry support for OCI v1.1
- ▼ [chainguard.dev](https://chainguard.dev) - Intro to OCI Reference Types
- ▼ [explore.ggcr.dev](https://explore.ggcr.dev) - Super handy web view of container innards
- ▼ [github.com/emporous](https://github.com/emporous) - object connector using OCI
- ▼ [github.com/package-url/purl-spec](https://github.com/package-url/purl-spec) - OCI type
- ▼ [github.com/distribution/distribution](https://github.com/distribution/distribution) - PR 3834 for references support

# Thank you!

## Vincent Batts

Azure Linux

vbatts

- ♣ Reach out to learn more or get involved



CONTAINERCON



OPEN SOURCE SUMMIT

NORTH AMERICA

THE LINUX FOUNDATION

