WHAT'S NEXT FOR CONTAINERS

Vincent Batts (vbatts)

```
$> finger $(whoami)
Login: vbatts                           Name: Vincent Batts
Directory: /home/vbatts                 Shell: /bin/bash
Such mail.
Plan:
OHMAN
$> id -Gn
devel opencontainers docker appc redhat golang slackware
```

Get Past the Hype

Get Past the Hype

Make Containers Boring

Get Past the Hype

Make Containers Boring

Common Interfaces

Common Interfaces

Standards

- OpenContainers (OCI) runtime and image specification
- Cloud Native Compute Foundation (CNCF)
- Container Network Interface (CNI)

Interface

- Kubernetes Container Runtime Interface (CRI)

What is needed next?

- Unified plumbing
- Discoverable
- UX

“Getting your chi right with systemd”

Then it all makes sense

Let's start a container

Let's start a container

- `systemd-nspawn -M httpd bash`

Let's start a container

- `systemd-nspawn -M httpd bash`
- `machinectl start httpd`

Let's start a container

- `systemd-nspawn -M httpd bash`
- `machinectl start httpd`
- `systemctl start systemd-nspawn@httpd.service`

Let's start a container

- `systemd-nspawn -M httpd bash`
- `machinectl start httpd`
- `systemctl start systemd-nspawn@httpd.service`
- `systemd-run -M httpd /usr/bin/tail -f /dev/null`
  - (has to be an already running container)
  - `machinectl shell httpd /usr/bin/systemctl status run-u69.service`

- `machinectl status httpd`

- `machinectl status httpd`
- `systemctl status systemd-nspawn@httpd.service`

- `machinectl status httpd`
- `systemctl status systemd-nspawn@httpd.service`

Different views/properties for the same service

```
/etc/systemd || /usr/lib/systemd
```

```
/etc/systemd || /usr/lib/systemd
```

- `./nspawn/https.nspawn`

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf
- ./system/lamp-stack.slice

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf
- ./system/lamp-stack.slice
- ./system/machines.target.wants/systemd-nspawn@httpd.service

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf
- ./system/lamp-stack.slice
- ./system/machines.target.wants/systemd-nspawn@httpd.service

- systemd.nspawn(5)

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf
- ./system/lamp-stack.slice
- ./system/machines.target.wants/systemd-nspawn@httpd.service

- systemd.nspawn(5)
- systemd.network(5)

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf
- ./system/lamp-stack.slice
- ./system/machines.target.wants/systemd-nspawn@httpd.service

- systemd.nspawn(5)
- systemd.network(5)
- systemd.resource-control(5)

/etc/systemd || /usr/lib/systemd

- ./nspawn/https.nspawn
- ./system/systemd-nspawn@httpd.service.d/50-Memory.conf
- ./system/lamp-stack.slice
- ./system/machines.target.wants/systemd-nspawn@httpd.service

- systemd.nspawn(5)
- systemd.network(5)
- systemd.resource-control(5)
- systemd.directives(7)

```
/etc/systemd/user || /home/$USER/.config/systemd/user
```

```
/etc/systemd/user || /home/$USER/.config/systemd/user
```

Next?

- Limited user service containers
- bwrap@.service ?
- Different from user namespaced nspawn containers

OTHER POINTS:

Pulling from OpenContainers image

It's in the same vein as portable services

OTHER POINTS:

Really glad to see `--network-zone=<NAME>` in v230

Would be nice if `--port` could expose to localhost

OTHER POINTS:

```
User namespaces. Cool and Crazy, but VFS needs to be right
(without `chown -R `)

overlayfs == copy up

UID/GID shift.
  https://github.com/systemd/systemd/issues/2404
  https://lkml.org/lkml/2016/5/4/411
  Current iteration: https://lkml.org/lkml/2016/5/12/655
Not looking like its issues will get sorted out any time soon.
```

OTHER POINTS:

`cgroup namespace + limited user containers`

currently requires a privileged helper to provide ownership for the pid

Aleksa Sarai (SUSE) is on v10 of patchset on LKML

OTHER POINTS:

`machinectl login <name>` is still blocked by selinux on fedora

https://bugzilla.redhat.com/show_bug.cgi?id=1310464

(the 4 prior BZs were closed as WONTFIX or because EOL)

Questions?

Questions?

Thanks!