Kubernetes, Docker volume mounts, and autofs

Environment details

- Machine_Type: Virtual
- OS: Oracle Enterprise Linux 7.x
- Software: Docker 1.12.6, Kubernetes 1.7.1

Scenario: Can’t Login via ssh public key

Unable to login to docker host using public key authentication
Able to login to the host using my password
Once at the console, I observed an error similar to:
Could not chdir to home directory /home/myuser: Too many levels of symbolic links
-bash: /home/myuser/.bash_profile: Too many levels of symbolic links

Hmm wtf …

Troubleshooting Steps

A fellow admin suggested I check for docker mapped volumes that point to /home

Here’s the command I used to query for that:

```
sudo docker ps --filter volume=/opt --format "Name:\n\t{{.Names}}\nID:\n\t{{.ID}}\nMounts:\n\t{{.Mounts}}\n"
```

Boom, looks like the kubernetes weaver container is using that mapping:

```
Name:
k8s_weave_weave-net-ljzn9_kube-system_740c10c5-d6b8-11e7-838f-005056b5384e_0
ID:
dc95801e4442
Mounts:
/opt/kubernetes,/lib/modules,/run/xtables.lo,
/var/lib/kubele,/var/lib/weave,/etc,/var/lib/dbus,/var/lib/kubele,/opt
```

Ok, so why would a docker volume mapped to /home induce such a problem?

Turns out that in some cases, binding autofs-mounted paths to docker containers can cause problems on the docker host.

This is due to the way in which kubernetes performs the volume mapping, which utilizes docker volume binds under the hood.

And, depending on how you map a volume to a docker container, you might conflict with autofs volume mounting.

For insight into a similar issue, see:

- Issue with AutoFS mounts and docker 1.11.2: https://github.com/moby/moby/issues/24303

According to the above issue description, the problem we're seeing might be fixed by adjusting the bind propagation for the volume mount in question,
see:
https://docs.docker.com/engine/admin/volumes/bind-mounts/#choosing-the–v-or-mount-flag

However, there's no way to control that setting via a kubernetes manifest, not at present at least, since HostPath bind propagation is currently a proposed feature in

kubernetes,
see:
https://github.com/kubernetes/community/blob/master/contributors/design-proposals/node/propagation.md

So the best course of action is to simply change hostPath setting in the weave-kube manifest, e.g.

- Change:
  hostPath:
  path: /home
- To:
  hostPath:
  path: /opt/kubernetes/bind-mounts/weave-kube/home

You can then simply redeploy the offending container (sudo docker stop ecfa204283d3 && sudo docker rm ecfa204283d3 && kubectl apply -f net.yaml)

Note: You'll have to perform similar changes to the weave manifest according to whatever other autofs mounts its hostPath(s) might conflict with.

Ensure you review your autofs settings!