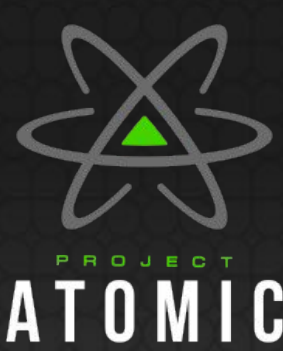

Fedora + Kube

Deployment Options

Jason Brooks
Project Atomic

Yu Qi (Jerry) Zhang
Project Atomic

Why Kubernetes

- Containerize all the things
- Production apps need to span multiple hosts
- Kubernetes is great for this
  - Run your container(s) on a cluster
  - Communicate via services
- Lots of communities have adopted it:
  - Project Atomic
  - Fedora Atomic WG
  - Red Hat (via OpenShift)
  - Kubevirt
  - Many others in the container-loving community

Kube or Origin?

- Kubernetes
  - First: leads origin by ~1 release
  - Friends: very popular
- OpenShift Origin
  - Features: build tools, slick UI, multi-tenancy
  - Friends: fits well w/ Fedora, lots of RHers
- Fedora can do both
  - Contributors are the key

fedora™

Part 1: The Bits

Upstream Binaries

- Kubernetes project provides them for every release, and there are lots of releases
  - three active streams at a time, plus alpha/beta
  - multiple archs
  - many installation methods expect to call on a specific version
- OpenShift Origin also provides binaries
- These are outside of our control/responsibility
  - We sometimes need/want to patch things

Fedora RPMs

- We distribute what we've built
  - we can patch things if we need to
- For Kubernetes, we haven’t been tracking all the releases closely
  - We need more contributors / automation
  - We’re bumping up against version / release issues
  - Modularity might help
- Origin moves at a more Fedora-friendly pace

Part 2: Packaging

None

- copy to /usr/bin/local and go
- write a systemd unit file, if you like

Hyperkube Image

- From upstream: debian+hyperkube binary
- Used by kubeadm, various other installers
- We don't have a Fedora-based hyperkube image, yet, but we can/will
  - We'll need to patch scripts/tools somewhat to use a Fedora image
- OpenShift Origin has "origin" image, similar deal, CentOS-based
- How do we feel about non-Fedora base images, anyway?

Fedora RPMs

- installable on regular fedora
  - baked into the fedora atomic image
  - we're planning to remove
- installable on fedora atomic via package layering
  - rpm-ostree install kubernetes
- managed by systemd

- We also have OpenShift Origin rpms in Fedora

FROM fedora:26

- Per-component images based on Fedora RPMs
- Run as
  - system container
  - from systemd unit files
  - from kubelet manifest
- Drop-in replacement for baked-in RPMs
- We need better versioning / tagging
- Modularity might help us here

- No official Fedora-based OpenShift Origin container

fedora

Part 3: Deployment

From Scratch

- Lots of people want to get their hands dirty & avoid scripts
- Can use rpm, system containers, plain binaries
- Kubernetes the Hard Way
  - My atomic-friendly fork: https://github.com/jasonbrooks/kubernetes-the-hard-way
- Project Atomic Getting Started Guide
  - http://www.projectatomic.io/docs/gettingstarted/
  - Needs work
- Not sure about OpenShift Origin

Minikube / Minishift

- Single-node cluster in a VM for quickstart
- Cross-platform friendliness
- Not generally Fedora-based
  - Boot2docker VM
  - Minishift can use CentOS VM
  - optional '--vm-driver=none' uses host's docker
- https://github.com/kubernetes/minikube
- https://github.com/minishift/minishift

fedora

Kubeadm

- single or multi-node
  - kubeadm init
  - considered “beta”
  - not HA
- Mix of installed rpms and container images
  - Upstream ships CentOS rpms for former, Debian-based hyperkube image for latter
- Fedora packages available
  - Package layering, still uses upstream hyperkube
- Kubeadm system container on the way

fedora

oc cluster up

- single or multi-node*
  - * oc cluster join, but I haven't gotten it working myself
  - not HA?
- Mix of installed rpms and container images
  - origin-clients is in Fedora, CentOS-based container image
- Probably a good candidate for system containerization

contrib/ansible

- https://github.com/kubernetes/contrib/tree/master/ansible
- Supports multiple distros, for Fedora, uses rpms or assumes kube components baked in
  - Can use w/ containers, system or no
- Support for openstack, aws via vagrant
- Needs more contributors
  - HA etcd, no HA master
  - Talk of moving under kube-deploy

kubespray

- https://github.com/kubernetes-incubator/kubespray
- kubernetes incubator project
- Ansible-based
- AWS, GCE, Azure, OpenStack or Baremetal
- High-availability
- supports RHEL/CentOS, not Fedora, but I got Fedora working w/ small tweaks
- hyperkube images from CoreOS, but image repo can be modded easily
- setenforce 0

openshift-ansible

- https://github.com/openshift/openshift-ansible
- Supports Fedora/CentOS/RHEL, uses rpms or containers
- HA etcd, HA master
- Very active
  - https://github.com/openshift/openshift-ansible-contrib
    - Vagrant, GCP, AWS, VMware, Azure, OpenStack and oVirt
- http://www.projectatomic.io/blog/2016/12/part1-install-origin-on-f25-atomic-host/

fedora

openshift-ansible

- Can use system containers for origin!
  - openshift_use_openvswitch_system_container=True
  - openshift_use_node_system_container=True
  - openshift_use_master_system_container=True
  - openshift_use_etcd_system_container=True
  - system_images_registry="docker.io"
- Can also use system containerized docker
  - openshift_docker_use_system_container=True
  - Uses /etc/docker/container-docker.json for config
  - Service is "container-engine"

fedora

SELinux

- too much setenforce 0 out there
  - spc_t for unconfining specific containers
  - chcon -Rt container_share_t
- network plugins, addons need modification
- help upstreaming this stuff is needed

Contributing

- https://pagure.io/atomic/kubernetes-sig
- https://bodhi.fedoraproject.org/updates/?packages=kubernetes
- https://www.openshift.org/#contribute
- https://fedoraproject.org/wiki/Container:Guidelines
- https://github.com/projectatomic/atomic-host-docs

fedora

Questions?

Contact:
jbrooks@redhat.com
jerzhang@redhat.com