

# Continuous Integration of an Operating System in Kubernetes

Stef Walter  
Red Hat



# *Continuous Integration*

# *Continuous Integration*


Assemble everything together like in production,  
and then drive it like a user.



# *Continuous Integration*

Do that integration for every single "change".

# September 2017







RED HAT ENTERPRISE LINUX SERVER

- Falcon
- System
- Logs
- Storage
- Networking
- Containers
- oVirt Machines
- Accounts
- Services
- Applications
- Diagnostic Reports
- Kernel Dump
- SELinux
- Software Updates
- Subscriptions
- Terminal

Dashboard Cluster Image Registry

Hardware TAROX Basic 7000BD  
 Asset Tag 1437471  
 Machine ID a20c82e128524937ad8...  
 Operating System Employee SKU  
 Secure Shell Keys [Show fingerprints](#)  
 Host Name Falcon (falcon.thewalter...)  
 Domain [Join Domain](#)  
 System Time 2017-10-17 13:47 ⓘ  
 Power Options  ▾  
 Performance Profile none ⓘ  
 Store Performance Data

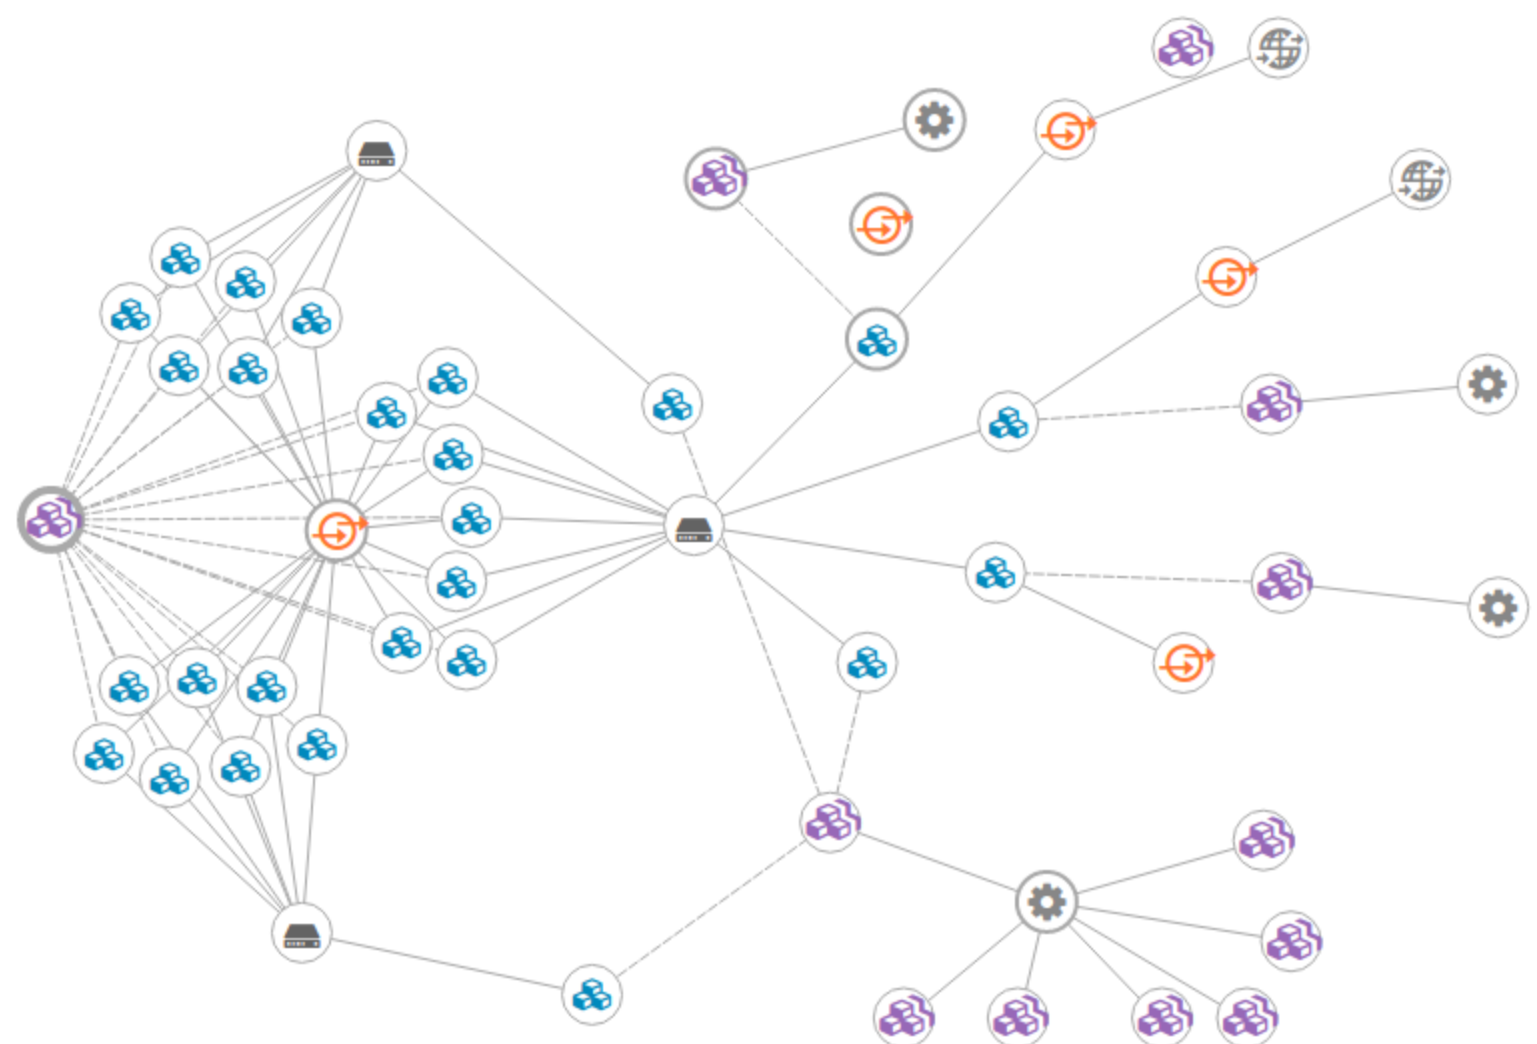




RED HAT ENTERPRISE LINUX SERVER

- Overview
- Nodes
- Container
- Topology**
- Details
- Volumes
- Images
- Projects

All Projects



ReplicationController

**Name** cockpit-tasks  
**Namespace** verify  
**Created** Sep 20, 2017 10:42:49 AM  
**Replicas** 20

Selector

infra cockpit-tasks

Pod Template

**Restart policy** Always  
**DNS policy** ClusterFirst  
**Service acco...** cockpituous

Containers

**Name** cockpit-tasks  
**Image** docker.io/cockpit/tests  
**Ports** none  
**Env vars** TEST\_JOBS=8  
 TEST\_PUBLISH=sink

**Name** cockpit-images  
**Image** docker.io/cockpit/images  
**Ports** 80/TCP

System

Hardware TAROX Basic 7000BD

Services

Asset Tag 1437471

Containers

BIOS American Megatrends Inc. 2001 (06/16/2014)

Logs

Operating System Fedora 23 (Workstation Edition)

Storage

Host Name Falcon (falcon.thewalter.lan)

Networking

Domain [Join Domain](#)

Elements

Console

Sources

Network

Timeline

Profiles

Resources



&lt;top frame&gt;

 Preserve log

&gt; zz = cockpit.spawn



System

Services

Containers

Logs

Storage

Networking

Hardware TAROX Basic 7000BD

Asset Tag 1437471

BIOS American Megatrends Inc. 2001 (06/16/2014)

Operating System Fedora 23 (Workstation Edition)

Host Name Falcon (falcon.thewalter.lan)

Domain 

Elements

Console

Sources

Network

Timeline

Profiles

Resources



&lt;top frame&gt;

 Preserve log

## 90+ APIs: File, Command, REST, DBus, Socket

abrt	AppStream	apt-get	atomic	Candlepin /candlepin/
chpasswd	CloudForms	cryptsetup	curl	dbus-daemon
device-mapper	docker-storage-setup	docker	e2fsprogs	etcd
/etc/kdump.conf	/etc/passwd	firewalld	FreeIPA	GnuTLS
GSSAPI	hostnamed	ipa-client	ipa-client	iproute
iptables	iscsi-tools	journalctl	kdump	klist
krb5	Kubernetes /api/	lastlog	libvirt	loginctl
lvm	mdadm	NetworkManager	NetworkManager-te...	oddjob
Openshift /oapi/	Openshift OAuth2	openssl	ostree	oVirt /api/
PackageKit	passwd	PCP	PolicyKit	/proc/meminfo...
/proc/mounts	/proc/net/dev...	procps-ng	/proc/stat	pwquality
qemu	realmd	rm ...	rpmostreed	rpm
selinux-policy-target...	selinux-utils	setenforce	Setroubleshootd	shadow-utils
shutdown	sosreport	ssh-agent	sshd	ssh-keygen
ssh	sssd	storaged	subscription-manager	sudo
/sys/fs/cgroup	/sys/kernel	/sys/power	systemd	timedated
Tuned	udev	UDisks2	/usr/bin/kubectl	/usr/bin/timedatectl
/usr/bin/virt-install	/var/log/wtmp	/var/run/utmp	virsh	who/w
xfspgms	yum			

## 15+ Linuxes and Products

---

CentOS (7.x, Atomic)

---

Fedora (25, 26, 27, Atomic)

---

Ubuntu (stable, 16.04)

---

RHEL (7.x, 7.4, Extras, Atomic)

---

Debian (stable, testing)

---

Openshift

---

RHEV Hypervisor

## 5+ maintained branches

---

master

---

rhel-7.4

---

rhel-7.3.x

---

i386

---

s390x

## 3+ browsers

---

Google Chrome

---

Internet Explorer


---

Firefox






$$92 \times 15 \times 5 \times 3 \times 50$$



Can i borrow  
a cup of  
robots?



Containerize your bots, yo



Scaling up and down



Communicating between bots



**Containerize your bots, yo!**

## Cockpit Task Bots

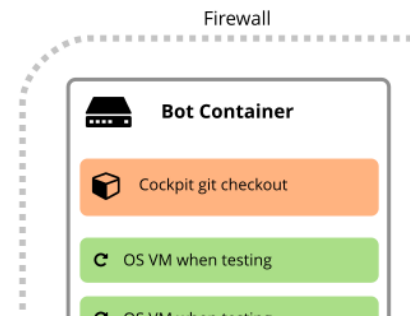
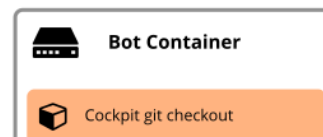
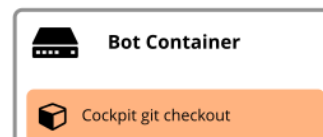
Task system is distributed with GitHub as the single point of failure



Every pull request is booted 100's of times in Atomic, Debian, Fedora, RHEL, before it hits master.

Containers that start 1,000 - 10,000 test VMs a day

These containers can run anywhere



- Self contained
- Schlep around all the junk you need
- Easy to update atomically
- Rolling and asynchronous updates
- Canary testing of bot environment

Testing **containers**  
in **Virtual Machines**  
in **container** bots  
running on **virtual machines**  
on **turtles!**

**Yo Dawg...**  
v



**Scaling up and down**



```
$ kubectl scale --replicas=30 cockpit-tasks
```

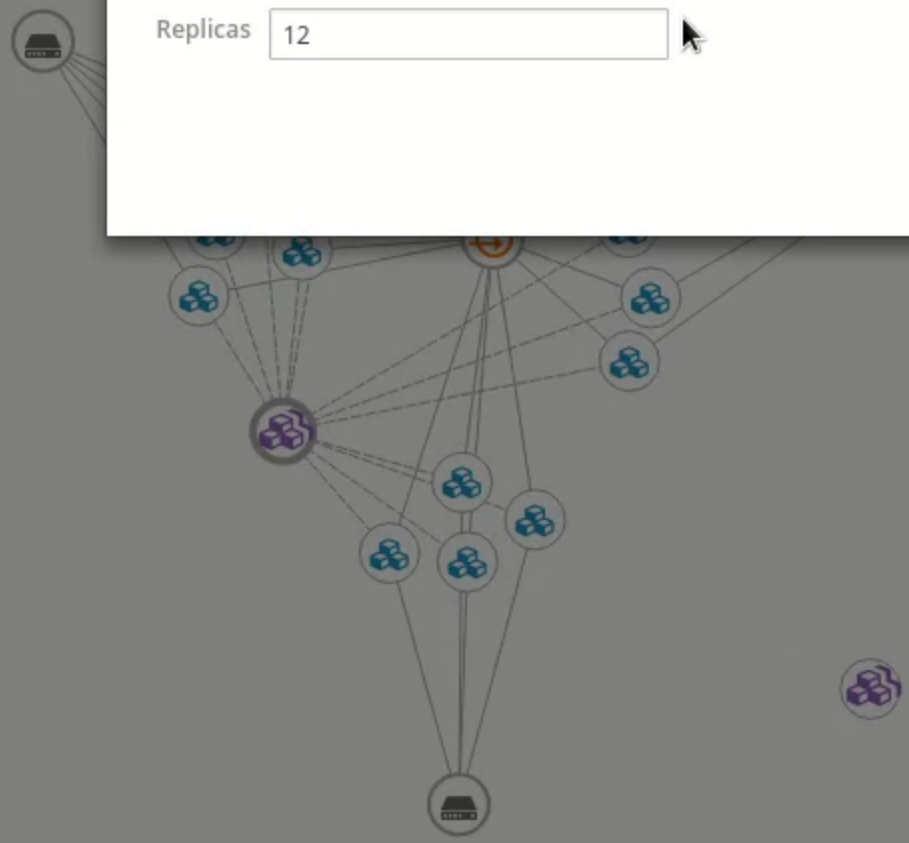
- Overview
- Nodes
- Containers
- Topology
- Details
- Volumes
- Images
- Projects

Project: verify

**Adjust Replication Controller cockpit-tasks**

Replicas

Cancel Change



### ReplicationController

<b>Name</b>	cockpit-tasks
<b>Namespace</b>	verify
<b>Created</b>	Sep 20, 2017 10:42:49 AM
<b>Replicas</b>	12

### Selector

<b>infra</b>	cockpit-tasks
--------------	---------------

### Pod Template

<b>Restart policy</b>	Always
<b>DNS policy</b>	ClusterFirst
<b>Service acco...</b>	cockpituous

### Containers


<b>Name</b>	cockpit-tasks
<b>Image</b>	docker.io/cockpit/tests
<b>Ports</b>	none
<b>Env vars</b>	TEST_JOBS=8 TEST_PUBLISH=sink

<b>Name</b>	cockpit-images
<b>Image</b>	docker.io/cockpit/images



**Sharing state between pods**

# Bots share state



*Bot gossip*

## Headless Services

```
{
  "kind": "Service",
  "metadata": {
    "name": "cockpit-images"
  },
  "spec": {
    "clusterIP": "None",
    "selector": { "infra": "cockpit-tasks" }
  }
}
```

```
$ host cockpit-images.verify.svc.cluster.local
cockpit-images.verify.svc.cluster.local has address 10.128.1.158
cockpit-images.verify.svc.cluster.local has address 10.128.1.159
cockpit-images.verify.svc.cluster.local has address 10.128.1.160

...
$ dig +short cockpit-images.verify.svc.cluster.local | while read addr; do
curl -fsS http://$addr/rhel-7-ccd75fa36e17ffb1fa8d0da29fe2dc1afed28cd2662
curl: (22) The requested URL returned error: 404 Not Found
curl: (22) The requested URL returned error: 404 Not Found
curl: (22) The requested URL returned error: 404 Not Found
curl: (22) The requested URL returned error: 404 Not Found
#####
...

```

**Going off road**





/dev/kvm in a Container



OCI Hook: oci-kvm-hook




Sharing state via host mounts



High density via kernel caching

**!** `/dev/kvm` in a container





VM lifecycle under control of a container

/dev/kvm is already privilege separated

```
$ sudo docker run -ti busybox ls -l /dev/kvm
ls: /dev/kvm: No such file or directory
$ sudo docker run -ti --device=/dev/kvm busybox ls -l /dev/kvm
crw-rw-rw-  1 root    36          10, 232 Oct 19 08:20 /dev/kvm
```

```
$ kubectl run -ti test --image=busybox --restart=Never -- ls -l /dev/kvm  
ls: /dev/kvm: No such file or directory
```



**add support for host devices #5607**

[github.com/kubernetes/kubernetes/issues/5607](https://github.com/kubernetes/kubernetes/issues/5607)

**!** `oci-kvm-hook`



**OPEN** CONTAINER  
INITIATIVE

[github.com/opencontainers/runtime-spec](https://github.com/opencontainers/runtime-spec)

```
$ sudo yum install -y oci-kvm-hook
$ kubectl run test --image=fedora --restart=OnFailure -- ls -l /dev/kvm
$ kubectl logs test
[root@cockpit-9 ~]# kubectl run -ti test --image=busybox --restart=Never
crw-rw-rw-    1 root    root      10, 232 Oct 18 10:41 /dev/kvm
```

[github.com/stefwalter/oci-kvm-hook](https://github.com/stefwalter/oci-kvm-hook)

**!** **Sharing state via host mounts**

```
"containers": [{
  "name": "cockpit-tasks",
  "image": "docker.io/cockpit/tests",
  "securityContext": { "runAsUser": 1111 },
  "volumeMounts": [{
    "mountPath": "/secrets",
    "readOnly": true
  },
  {
    "mountPath": "/cache",
    "readOnly": false
  },
  {
    "mountPath": "/tmp",
    "readOnly": false
  }
}
```

```
"volumes": [  
  {  
    "name": "secrets",  
    "secret": { "secretName": "cockpit-tests-secrets" }  
  },  
  {  
    "name": "cache",  
    "hostPath": { "path": "/var/cache/cockpit-tests" }  
  },  
  {  
    "name": "tmp",  
    "emptyDir": { }  
  }  
]
```

**!** High density via kernel caching



VMs: 50 - 70

Load: 50 to 60

IO: ~0 MB/s read 30-50 MB/s write

```
top - 07:35:41 up 11 days, 16:23, 3 users, load average: 39,71, 19,51, 11,65
Tasks: 948 total, 12 running, 925 sleeping, 0 stopped, 11 zombie
%Cpu(s): 82,4 us, 8,8 sy, 0,0 ni, 4,6 id, 4,1 wa, 0,0 hi, 0,1 si, 0,0 st
KiB Mem : 13161798+total, 24729728 free, 29423104 used, 77465152 buff/cache
KiB Swap: 8388604 total, 8388604 free, 0 used. 95578128 avail Mem
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
34327	1111	20	0	2045008	840812	16156	R	129,4	0,6	0:29.64	qemu-syst+
22549	1111	20	0	2108520	1,026g	16160	R	111,2	0,8	1:05.54	qemu-syst+
23714	1111	20	0	2059348	1,027g	16160	S	105,3	0,8	1:02.00	qemu-syst+
34772	1111	20	0	1979464	459912	16160	R	101,3	0,3	0:20.26	qemu-syst+
23346	1111	20	0	2571584	1,031g	16160	R	100,7	0,8	1:04.11	qemu-syst+
26147	1111	20	0	2111580	1,053g	16184	S	100,0	0,8	0:48.22	qemu-syst+
22431	1111	20	0	2063452	1,035g	16160	S	99,3	0,8	1:07.68	qemu-syst+
31953	1111	20	0	2056272	1,056g	16156	S	99,3	0,8	0:37.35	qemu-syst+
34155	1111	20	0	2029648	935168	16148	S	99,3	0,7	0:29.36	qemu-syst+
30742	1111	20	0	2068556	1,032g	16156	S	97,7	0,8	0:39.72	qemu-syst+
35031	1111	20	0	4728	952	556	R	97,0	0,0	0:15.58	gzip
34424	1111	20	0	1997900	616740	16160	R	96,7	0,5	0:22.92	qemu-syst+
22486	1111	20	0	2161752	1,030g	16172	R	96,0	0,8	1:05.95	qemu-syst+
34778	1111	20	0	1992784	458540	16160	S	95,4	0,3	0:20.29	qemu-syst+
34925	1111	20	0	2010188	411684	16152	S	95,4	0,3	0:15.26	qemu-syst+
34857	1111	20	0	1987660	431524	16156	S	92,4	0,3	0:17.93	qemu-syst+
22462	1111	20	0	2084968	557976	16164	S	92,1	0,4	0:20.51	qemu-syst+



Overview All Projects 33 Pods All running

Nodes 3 Nodes All healthy

Containers Topology

Details

Volumes

Images

Projects

Services

Name	IP	Ports	Replicas	Config
docker-registry				
registry-console				
router	172.30.204.61	80 443 1936	1	default
cockpit-images	None	80 443	54	verify

Containers	IP
redhat.com	19
cockpit-8.e2e.bos.redhat.com	19
cockpit-9.e2e.bos.redhat.com	22

### Connection Settings

Cluster: cockpit-9-e2e-bos-redhat-com:8443

Address: https://cockpit-9.e2e.bos.redhat.com:8443

Skip Certificate Verification

Requires Authentication

User: cockpit/cockpit-9-e2e-bos-redhat-com:8443/cwqq

Token: CQMhfE-MsgquqdnG0SDGYAxYBJl8J\_n9FJaKQ0FbnqE

Cancel Connect

# Questions?

[cockpit-project.org](https://cockpit-project.org)  
#cockpit on FreeNode

Credits:

