

Managing jenkins with multiple components project

Presenter Name Ohad Basan

Presenter Company Www.redhat.com

July 16, 2014

#jenkinsconf

Who am I

- Ohad Basan
- CI engineer at Red Hat
- Member of the RHEV (Red Hat Enterprise Virtualization) CI team
- Full time geek



Agenda

- Testing multiple components w/o loosing your mind AKA 'Factory Flow'
- "Puppet for jenkins jobs", possible?
- Giving developers access to automation infra
- Automate cool tasks in jenkins, foreman and other devops tools via fabric



What is Factory flow?







Problems with testing multiple components project at once

- Project is made out of multiple components that are constantly changing and depend on each other
- If two or more components change before the tests being – you have no idea which component broke the functionality
- We had to isolate the environment of each component while keep the other component stable.



What we had

- Every job builds all the components
 - Time consuming.
 - Load on git.
 - Error prone.
- In order to isolate the environments we created distinct jenkins slaves for each component. Each job compiles its own product.
 - Time Consuming
 - Error prone.
 - Increase load of git.
 - Resource heavy. (many slaves)



Solution - FACTORY

Listener job > Polls for SCM changes.



Create custom rpms > Builds the relevant component



Add missing from stable – Links the compiled component with the rest of the missing stable component and



Publish - publishes the files on a repository that contains all the needed components. (full project)



Automation jobs – Receive the repository link from the publisher, deploy it on the jenkins slaves and run the tests.



How did we solve it? Multijob!

rhevm engine 3.3 git listener > #328





Build #328 (Jun 19, 2014 5:39:49 PM)



No changes.



Started by user obasan



Revision: 60b0f8a243a63154499d42619239477b41652c20

• origin/ovirt-engine-3.3



	Building RPMs	rpms_factory		build #1183	(11 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_ho	tplug_sanity_nfs_rest_factory	<u>build #244</u>	(7.6 sec)
	Running Automation Jobs	rhevm_3.3_automation_storage_flo	ating_disk_sanity_nfs_rest_factory	<u>build #250</u>	(12 sec)
	Running Automation Jobs	update_latest_is_to_3.3_factory		<u>build #646</u>	(5 min 0 sec)
	Running Automation Jobs	update_latest_sf_to_3.3_factory		<u>build #742</u>	(5 min 54 sec)
	Running Automation Jobs	rhevm_3.3_automation_coretools_re	estapi_hosts_nfs_rest_factory	<u>build #250</u>	(9 min 48 sec)
	Running Automation Jobs	rhevm_3.3_automation_coretools_rl	hevh_restapi_hosts_nfs_rest_factory	<u>build #261</u>	(11 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_nfs	s_options_sanity_nfs_rest_factory	<u>build #240</u>	(13 min)
	Running Automation Jobs	rhevm_3.3_automation_infra_one_h	nost_restapi_mixed_nfs_rest_factory	<u>build #388</u>	(14 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_sar	nity_unittest_nfs_rest_factory	<u>build #236</u>	(16 min)
	Running Automation Jobs	rhevm_3.3_automation_network_sa	nity_unittest_iscsi_rest_factory	<u>build #135</u>	(17 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_live	_upgrade_sanity_nfs_rest_factory	<u>build #240</u>	(19 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_au	torecovery_sanity_nfs_rest_factory	<u>build #255</u>	(20 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_sar	nlock_sanity_nfs_rest_factory	<u>build #241</u>	(21 min)
	Running Automation Jobs	rhevm_3.3_automation_coretools_v	irt_vm_migration_basic_nfs_factory	<u>build #386</u>	(25 min)
	Running Automation Jobs	rhevm_3.3_automation_storage_clo	ne_vm_sanity_nfs_rest_factory	<u>build #258</u>	(26 min)
1	Running Automation Jobs	rhevm_3.3_automation_storage_ex	port_import_server_nfs_rest_factory	<u>build #249</u>	(51 min)
	Running Automation Jobs	rhevm_3.3_automation_coretools_t	wo_hosts_restapi_vms_nfs_rest_factory	/ <u>build #243</u>	(1 hr 4 min)





Multijob Jenkins Plugin

Pros

- Really easy to configure
- Nice and comfortable view

Cons

- Huge mess when running the the same job concurrently.
- Not flexible can't dynamically chose which jobs to run.
- Can't pass dynamic parameters from upstream jobs.



Evolution — Build Flow Plugin

Pros

- Very dynamic and flexible
- Highly maintained

Cons

- View is not always comfortable (We customized it to make it more readable)
- Graphs not always working

































Build #1191 (Jul 8, 2014 4:01:21 PM)









1. frontend: Fix NPE in ConsoleProtocol (detail)

Started by an SCM change

Revision: 0849d86fa1de37b47951b88a4e9e54e297c4a9ac

origin/ovirt-engine-3.4



Repo URL: factory-rhevm 3.4 listener-b1191

Package (only custom)	Version
rhevm	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-backend	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-dbscripts	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-lib	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-restapi	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-base	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-allinone	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-ovirt-engine	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-ovirt-engine-common	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-websocket-proxy	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-tools	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-userportal	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-webadmin-portal	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-websocket-proxy	3.4.1-0.1000.12.9c933e9.el6ev.noarch

rhevm 3.4 automation storage sanity unittest nfs rest factory#1027 - SUCCESS

rheym 3.4 automation infra one host restapi mixed nfs rest factory#1268 - SUCCESS

rheym 3.4 automation infra one host restapi mixed nfs cli factory#224 - SUCCESS







No Tags



Rebuild



Build Graph



Previous Build



Next Build



Repo URL: factory-rheym 3.4 listener-b1191	
Package (custom stable)	Version
rhevm	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-backend	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-dbscripts	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-lib	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-restapi	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-base	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-allinone	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-ovirt-engine	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-ovirt-engine-common	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-setup-plugin-websocket-proxy	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-tools	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-userportal	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-webadmin-portal	3.4.1-0.1000.12.9c933e9.el6ev.noarch
rhevm-websocket-proxy	3.4.1-0.1000.12.9c933e9.el6ev.noarch
jasperreports-server-pro	5.5.0-9.el6ev.noarch
mom	0.4.0-1.el6ev.noarch
novnc	0.4-8.1.el6.noarch
otopi	1.2.1-1.el6ev.noarch
otopi-devtools	1.2.1-1.el6ev.noarch
otopi-java	1.2.1-1.el6ev.noarch
otopi-repolib	1.2.1-1.el6ev.noarch
ovirt-host-deploy	1.2.2-1.el6ev.noarch
ovirt-host-deploy-java	1.2.2-1.el6ev.noarch
ovirt-host-deploy-offline	1.2.1-1.el6ev.x86_64
ovirt-host-deploy-repolib	1.2.2-1.el6ev.noarch
ovirt-hosted-engine-ha	1.1.2-5.el6ev.noarch
ovirt-hosted-engine-setup	1.1.3-1.el6ev.noarch
ovirt-node-plugin-vdsm	0.1.1-24.el6ev.noarch
ovirt-node-plugin-vdsm-recipe	0.1.1-24.el6ev.noarch
ovirt-scheduler-proxy	0.1.3-6.el6ev.noarch
python-argparse	1.2.1-2.1.el6.noarch
python-cpopen	1.3-2.el6_5.x86_64

How do we configure so many jobs?

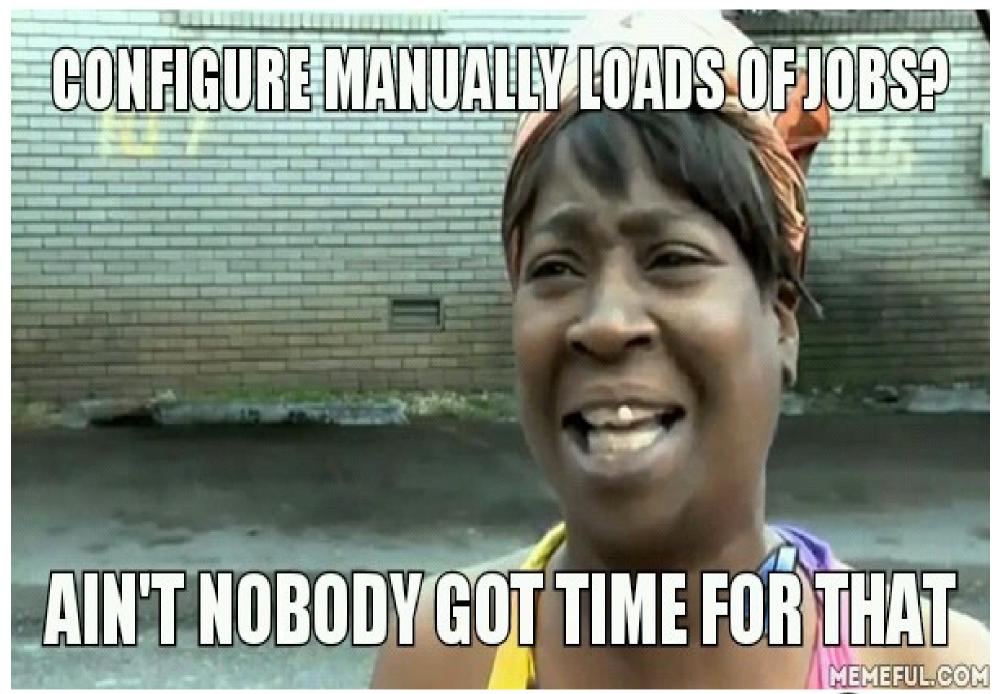
- We test a huge matrix
 - Different operating systems versions
 - Different product components
 - Different product versions
 - Backward compatible restrictions
 - Layered products consumption (jboss, kvm)



Layered product

- Ovirt (Backend) > Application Server (JBOSS)
 > Operating system (Red Hat Linux).
- VDSM (Hypervisor) > Libvirt > KVM (Kernel Virtual Machine) > Operating system
- Red Hat does ALL the components and EVERYTHING is changing all the time.







"Puppet for jenkins jobs", possible? Meet jenkins-job builder

http://ci.openstack.org/jenkins_jobs.html

 This tool uses YAML files that act as template file to generate jenkins jobs with different values and configuration



Example of YAML file

```
Project Definitions
name: clone vm sanity
version
suffix
    - factory vdsm
    - factory
   - factory developer
   - storage
   - clone vm sanity
    - rest
storage type:
   - nfs
jobs
```



```
ob-template:
parameters:
    - default params:
       num hosts: 1
    - string:
        name: CONF FILE
        default: rhevm 3.3 automation {test group} {test name}.conf
         description: "Configuration file to use for ART tests.
    - string:
        name: STORAGE TYPE
        default: '{storage_type}'
    - strina:
        name: HOST HOSTGROUP
        default: P-RHEVM-{version}-RHEL65-HOSTS
    - string:
        name: ENGINE
        default: '{engine}'
node: RHEVM-{version}-ENGINE-RHEL65
scm:
    - jenkins-{version}
    - ci-dev-tools
builders:
    - default builders job runner:
        version: '{version}
wrappers:
    - default wrappers
    - check conf param
properties:
    - storage owners
publishers:
     - default publishers job runner {version}
```



Giving developers access to automation infra

How can a dev test a patch?

- Unit tests
- Lint
- Very basic sanity

But for system tests, they need massive infra to run it, almost impossible for a dev





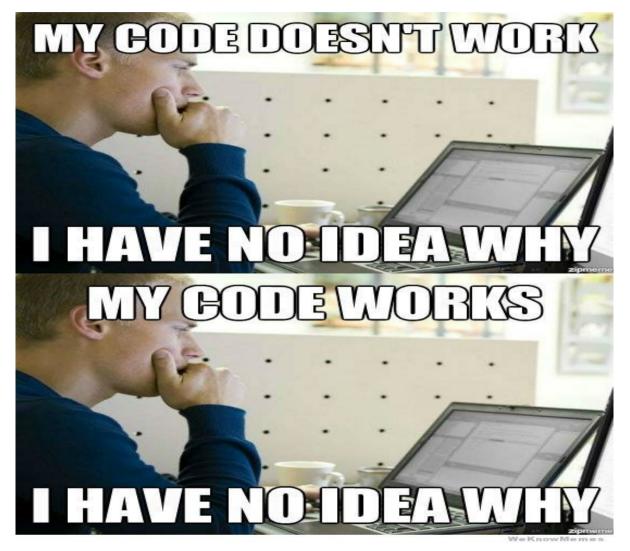


Mission: Patches verification

- Every patch has to be verified before merging
- Some patches are really hard to verify and require a complicated setup (several bare metal machines, storage servers, complicated configuration and preparation)
- The Solution...



Developer job





Developer job

- We give developers a special job to test their patches with
- It will checkout their branch, build the component that they chose and will run an automation job that they can choose.
- Build flow to rescue.. (again!)
- Using the factory flow infrastructure.



Build Flow engine_3.5_developer

This build requires parameters:

Build

This balla requires parameters.				
JOBS	rhevm_3.5_automation_infra_one_host_restapi_mixed_nfs_rest_factory rhevm_3.5_automation_coretools_webadmin_factory rhevm_3.5_automation_coretools_userportal_factory rhevm_3.5_automation_storage_sanity_unittest_nfs_rest_factory rhevm_3.5_automation_network_sanity_iscsi_rest_factory Jobs to run			
RHEVM_PATCHES				
	Comma separated of refspecs to apply to the engine code before packageing. For exameple: refs/changes/28/18328/39,refs/changes/85/19185/2			
CHECKOUT_FIRST_PATCH €				
	If set, will checkout the first patch passed so it includes all the commits it depends on, and cherry-pick the rest over that			
COMPILE_GWT				
	Enable to compile also gwt for the rpms			
PLAIN_JS				
	If set, will compile plain javascripts for the ui			
EMAIL				
	If set, will send an email with the results to that email			
UPSTREAM_BRANCH	master			
	If you don't know if you should change this value, you shouldn't			
UPSTREAM_REPO				
	If you don't know if you should change this value, you shouldn't			
DOWNSTREAM_BRANCH	master-rhevm			
	If you don't know if you should change this value, you shouldn't			
DOWNSTREAM_REPO				
	If you don't know if you should change this value, you shouldn't			
BASE_REPO	latest			
BASE_REPO_URL				

Managing the system

- Keeping track on problems can get complicated
- Problems can originate in infrastructure problems, automation problems, or product problems.
- We are using python fabric task to maintain and keep track on our environment.



What is fabric?

- Fabric is a python library and cli for easy deployments through ssh
- http://www.fabfile.org/
- We use it as a wrapper for performing all sorts of commands on our ci environment.
- We use Foreman api to use Fabric for foreman tasks.



Foreman Tasks

do.foreman.info.get_properties	Gets the given properties from the hosts. From the host properties or
do.foreman.info.get_subnet	
do.foreman.info.hostgroups_summary	Show a summary of the current used and unused profiles and machines
do.foreman.setup_host	Sets up and adds a host to foreman.
do.foreman.setup_host.add_to_foreman	Adds the host to foreman usig the API.
<pre>do.foreman.setup_host.deploy_existing_host</pre>	Sets up and adds a host to foreman.
do.foreman.setup_host.fix_hostname	Correctly set up the hostname
do.foreman.setup_host.install_pkgs	Installs the required packages to add a host to foreman
do.foreman.setup_host.install_repos	Installs all the needed repos to add the host into foreman

JENKINS TASKS

do.jenkins.generate_job_report	Generates a daily job report for the given view (by default 3.2, 3.3 and
do.jenkins.show_queue	
do.jenkins.update_view	This tasks take all the jobs from the given view and updated their bugs
do.jenkins.job.add_bz	Adds a bugzilla to the job description
do.jenkins.job.add trac	Adds a trac to the job description
do.jenkins.job.add work	Adds a work entry to the job description
do.jenkins.job.del bz	Deletes a bugzilla entry from the given job
do.jenkins.job.del_trac	Deletes a trac entry from the given job
do.jenkins.job.del_work	Deletes a user's or all the work entries from the job
do.jenkins.job.get_job_logs	Fetches all the related vdsm and engine logs for the given build.
do.jenkins.job.get_last_build	Gets the latest jbuild for the given job and state.
do.jenkins.job.update	Updates the description bugs from the given job
<pre>do.jenkins.nagios.check_build_load</pre>	cheacks if the build load on all the slaves is acceptable
do.jenkins.nagios.check_executors	Checks if the executors statuses are acceptable
do.jenkins.nagios.check_queue	Check if the build queue is in an acceptable state
do.jenkins.slave.set_offline	Takes the slave offline
do.jenkins.slave.set_online	Brings the slave online
do.jenkins.slave.show	H H
do.jenkins.slave.show_offline	Shows a list of the offline slaves



```
[obasan@localhost fabric ci]$ fab do.jenkins.slave.show offline
Not enabling jira related tasks, install jira module to get full functionality
       State
               Reason
name
                                                <span class=error style='display:inline-block'>Time out for
jenkins-automation-rpm-vm27
                                permanently
jenkins-upgrade-jobs-vm03
                                temporarily
                                                725311488
jenkins-upgrade-jobs-vm12
                                temporarily
                                                No reason
jenkins-upgrade-jobs-vm14
                                temporarily
                                                1045454848
jenkins-automation-rpm-vm48
                                temporarily
                                                problem
taking offline
                                                <span class=error style='display:inline-block'>Time out for
jenkins-upgrade-jobs-vm08
                                permanently
jenkins-automation-rpm-vm46
                                temporarily
                                                No reason
jenkins-upgrade-jobs-vm01
                                                <span class=error style='display:inline-block'>Time out for
                                permanently
jenkins-upgrade-jobs-vm04
                                temporarily
                                                312115200
[obasan@localhost fabric ci]$
```



one can perform a command on a foreman query with fabric.

```
obasan@localhost fabric ci]$ sudo fab on.foreman.search:"hostgroup = RHEVM-3.5-ENGINE-RHEL65" -- yum update]
[sudo] password for obasan:
Not enabling jira related tasks, install jira module to get full functionality
Query used:
        "hostgroup = RHEVM-3.5-ENGINE-RHEL65"
Got 9 hosts:
        jenkins-automation-rpm-vm01.eng.lab.tlv.redhat.com
        enkins-automation-rpm-vm05.eng.lab.tlv.redhat.com
        enkins-automation-rpm-vm09.eng.lab.tlv.redhat.com
        enkins-automation-rpm-vm12.eng.lab.tlv.redhat.com
        enkins-automation-rpm-vm20.eng.lab.tlv.redhat.com
        enkins-automation-rpm-vm21.eng.lab.tlv.redhat.com
        enkins-automation-rpm-vm22.eng.lab.tlv.redhat.com
        jenkins-automation-rpm-vm27.eng.lab.tlv.redhat.com
        jenkins-automation-rpm-vm35.eng.lab.tlv.redhat.com
[s what you expected? y|n [y]
```



What is fabric?

- Fabric is a python library and cli for easy deployments through ssh
- http://www.fabfile.org/
- We use it as a wrapper for performing all sorts of commands on our ci environment.
- We use Foreman api to use Fabric for foreman tasks.



Fabric view update tasks

- We keep track on automation problems with JIRA ticket tracking system. For bugs we use bugzilla.
- We have fabric task that keeps the jobs description updated for tracking the bugs' statuses. (using jira and bugzilla api)
- All you have to do is write BZ:XX in the description and a cronjob with fabric will update the status





rheym 3.5 automation coretools userportal factory



<u>JIRA:RHEVM-1523</u> In Progress(movciari@redhat.com) - Please fix UI tests for 3.5



rheym 3.5 automation coretools virt vm migration basic nfs factory



BZ:1111069 NEW(Tal Nisan) - NumberFormatException after re-electing spm

[obasan@localhost fabric_ci]\$ fab -d do.jenkins.job.update
Not enabling jira related tasks, install jira module to get full functionality
Displaying detailed information for task 'do.jenkins.job.update':

 Updates the description bugs from the given job

 :param job: Name of the job to update
 :param jenkins_url: Url to the jenkins server (JENKINS_URL in the config by default)
 :param jenkins_user: JENKINS_USER in the config by default
 :param jenkins_passwd: JENKINS_PASSWORD in the config by default

Arguments:



The Future

- Provisioner provisioning slaves dynamically at the beginning of each test and then tearing them down at the end. Provisioning from foreman hostgroup using an openstack setup
- Nested virtualization Run the hypervisor tests on a vm using nested kvm.
- Adding more flow and components to the factory flow
- Adopting these solutions to other Red Hat products



Links

- Fabric http://www.fabfile.org/
- Build flow https://wiki.jenkinsci.org/display/JENKINS/Build+Flow+Plugin
- Multi job https://wiki.jenkinsci.org/display/JENKINS/Multijob+Plugin
- Job builder http://ci.openstack.org/jenkins-jobbuilder/
- Foreman api https://pypi.python.org/pypi/pythonforeman



Open Source

- Red Hat Enterprise Virtualization is using code from the upstream project "oVirt" - an open source virtual data center manager.
- http://www.ovirt.org
- Public gerrit is on gerrit.ovirt.org
- FEEL FREE TO CONTRIBUTE CODE!
- My E-mail: obasan@redhat.com







Thank You To Our Sponsors

Platinum





Gold









Silver

