

OpenStack

Operational details of a large python project

http://www.pixelbeat.org/talks/openstack_python/

Pádraig Brady – Red Hat

Agenda

- OpenStack Overview
- Public / Private cloud description
- Development Process
- Related python projects

About Me

Open Source user for 12 years

Python user for 10 years

Gnu Coreutils co-maintainer (4 of us) for 4 years

Red Hatter for 9 months

OpenStacker for 7 months

Fedora OpenStack release manager for 4 months

What is OpenStack

- IaaS platform (like AWS)
- Open Source
 - unlike VMware, AWS, Compute Engine etc.
- Leverages lots of existing Linux technologies
- Written mainly in $2.6 \leq \text{python} < 3.0$
- 2 years old
- 300K lines of python

The shift to cloud

- Division of labour
- Economies of scale
- Automation of provisioning
- Increased resiliency
 - Decreased MTTR
 - Due to scale and automation

Public or Private


- Need for private
 - Leverage existing systems
 - Specific requirements
 - Large scale economies
 - Sensitive logic or data
- Amazon acknowledged in March 2012 the need for private cloud by partnering with Eucalyptus
- Traditionally separate implementations
- OpenStack aims to support both with consistent implementation and APIs

Who uses OpenStack?

- private clouds
 - mercado libre (ebay of south america)
 - Wikimedia
 - scientific community
 - NASA
 - HEPIX
 - CERN (15K servers, 300K VMs by 2015)
 - ISI.edu
- Public clouds
 - Rackspace
 - HP

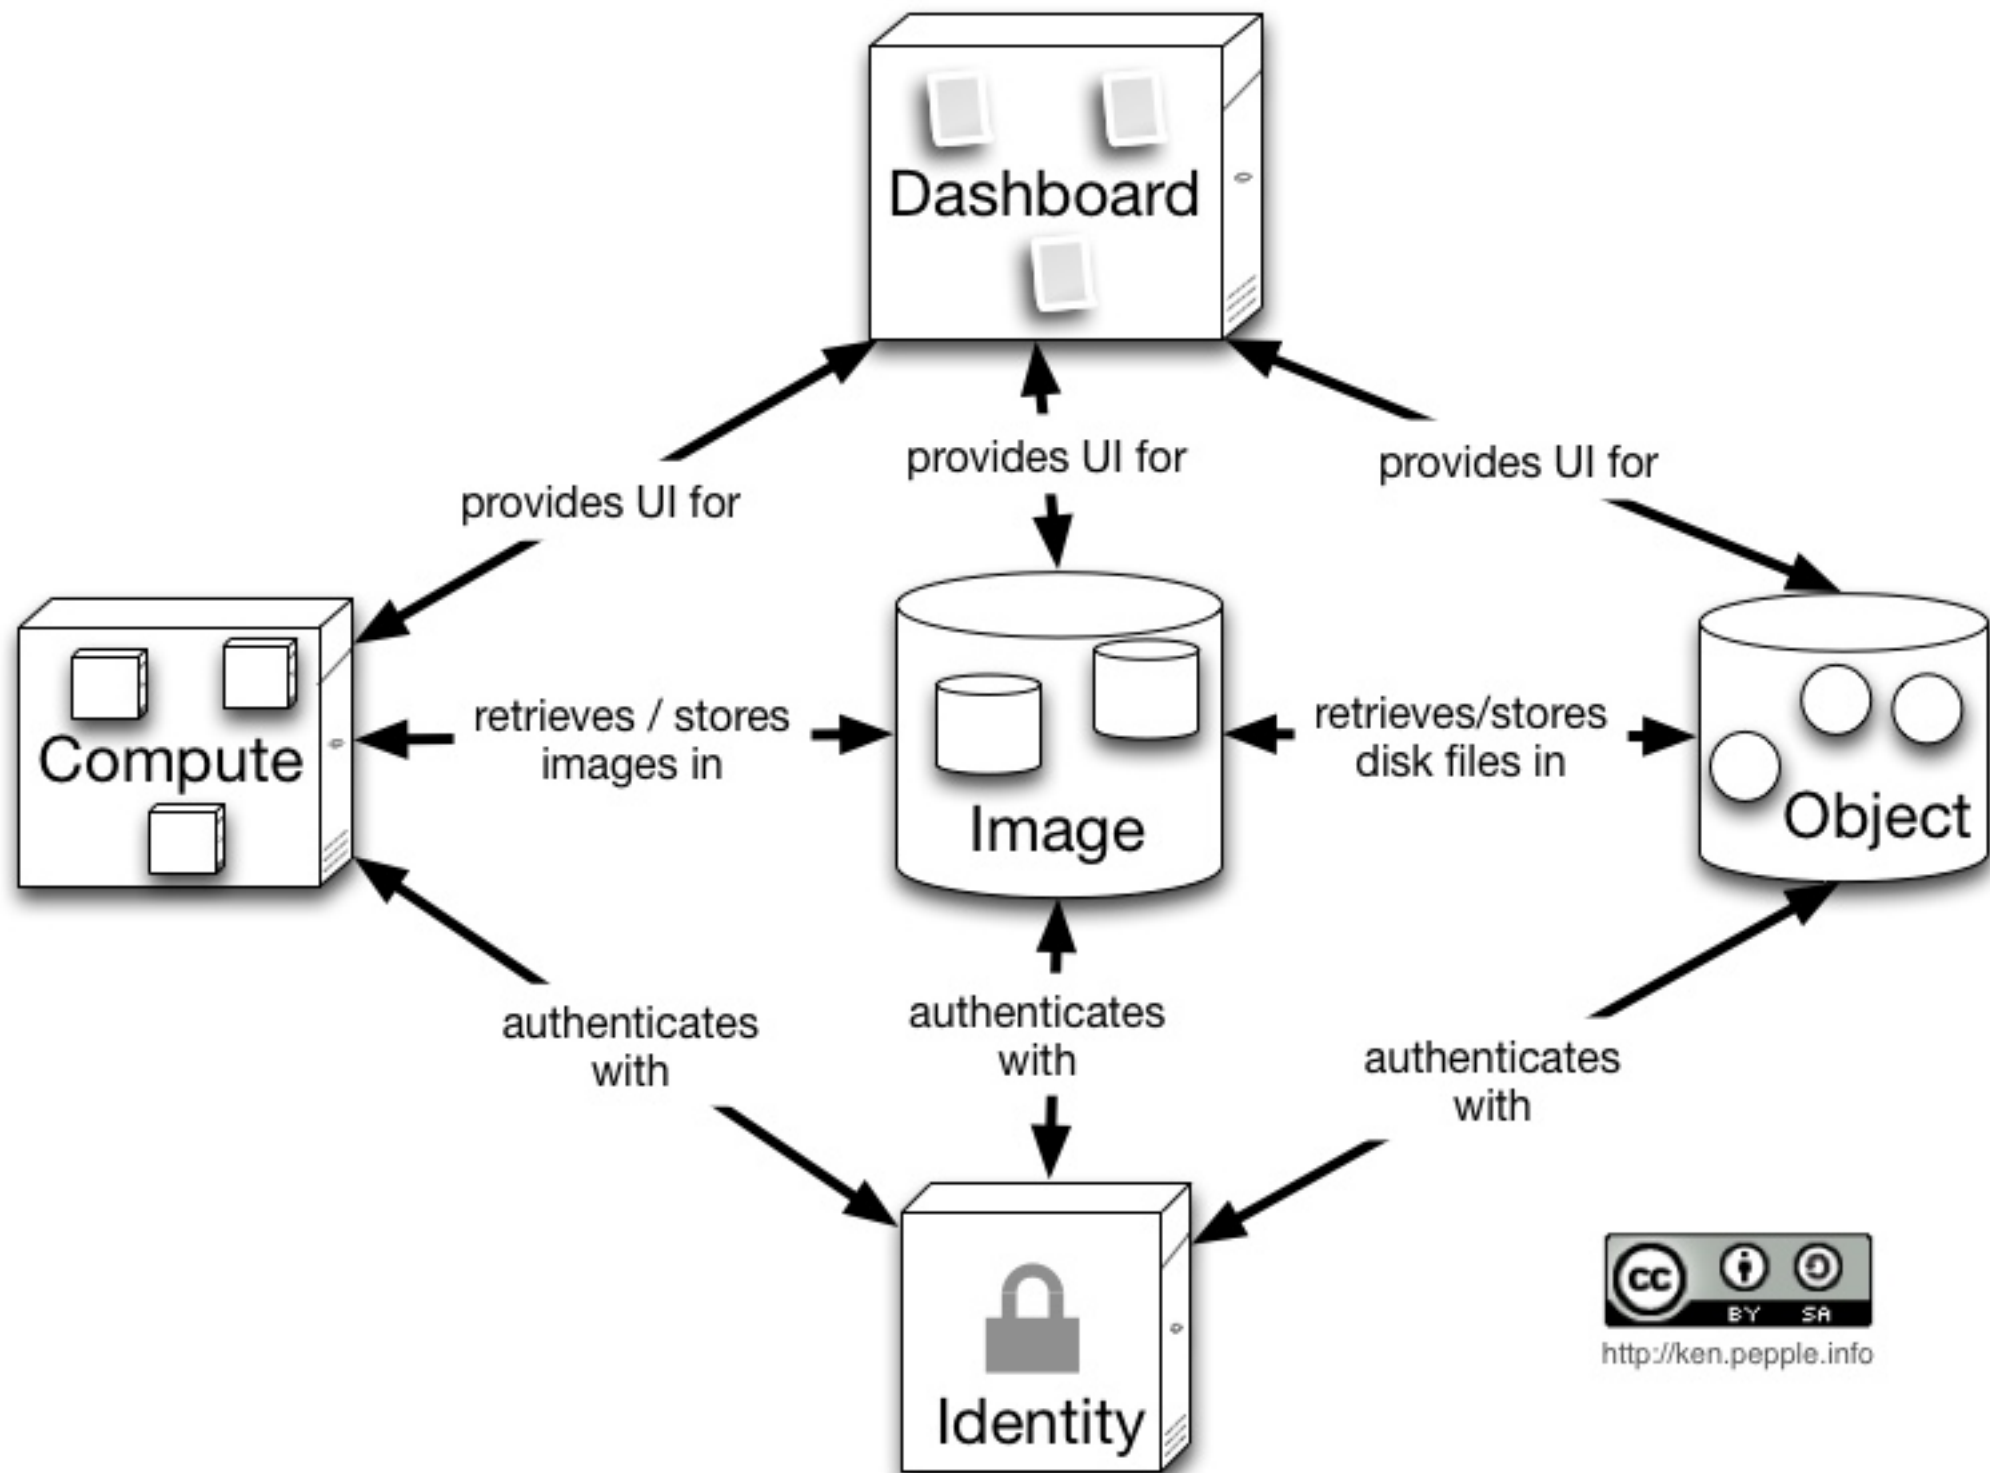
Linux Kernel Similarities





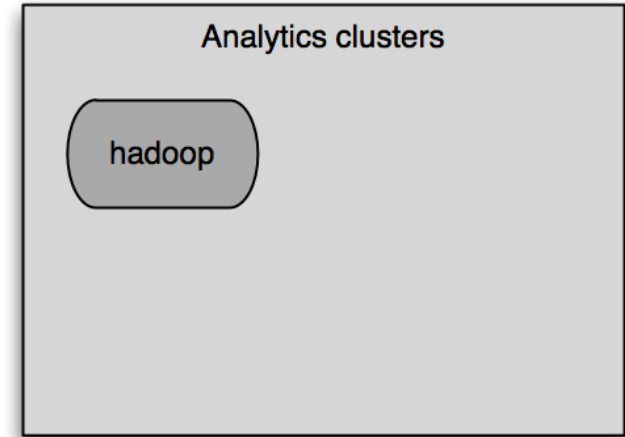
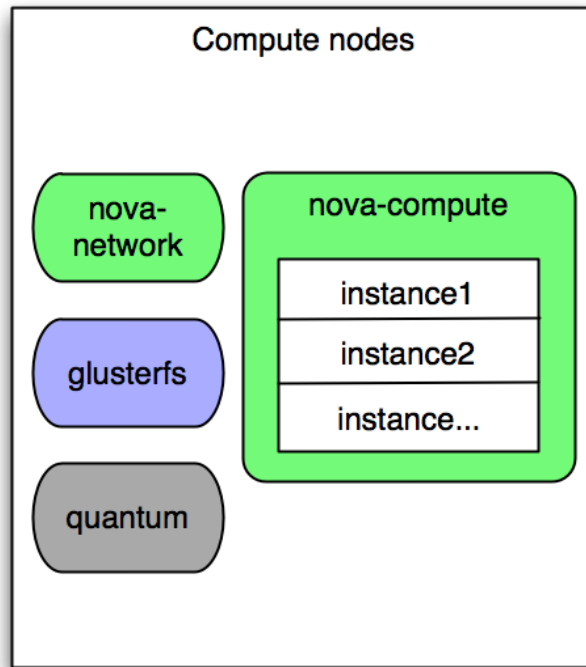
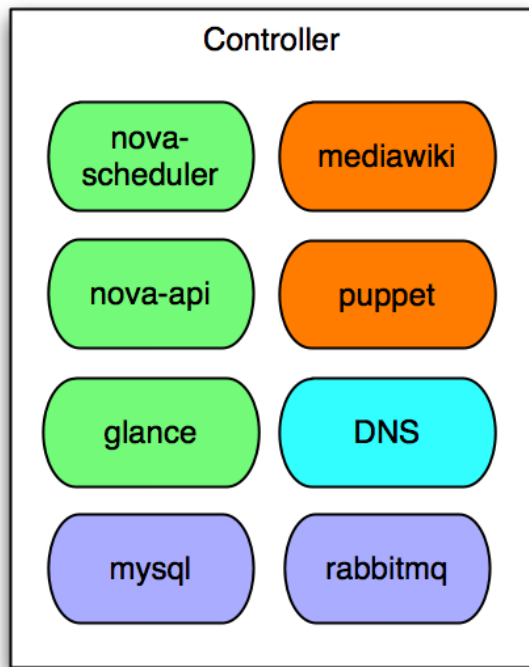
- Open Source
- Released early
- OpenStack Foundation is being formed akin to the Linux Foundation
- One could consider the abstractions quite similar too:



Core Projects Overview

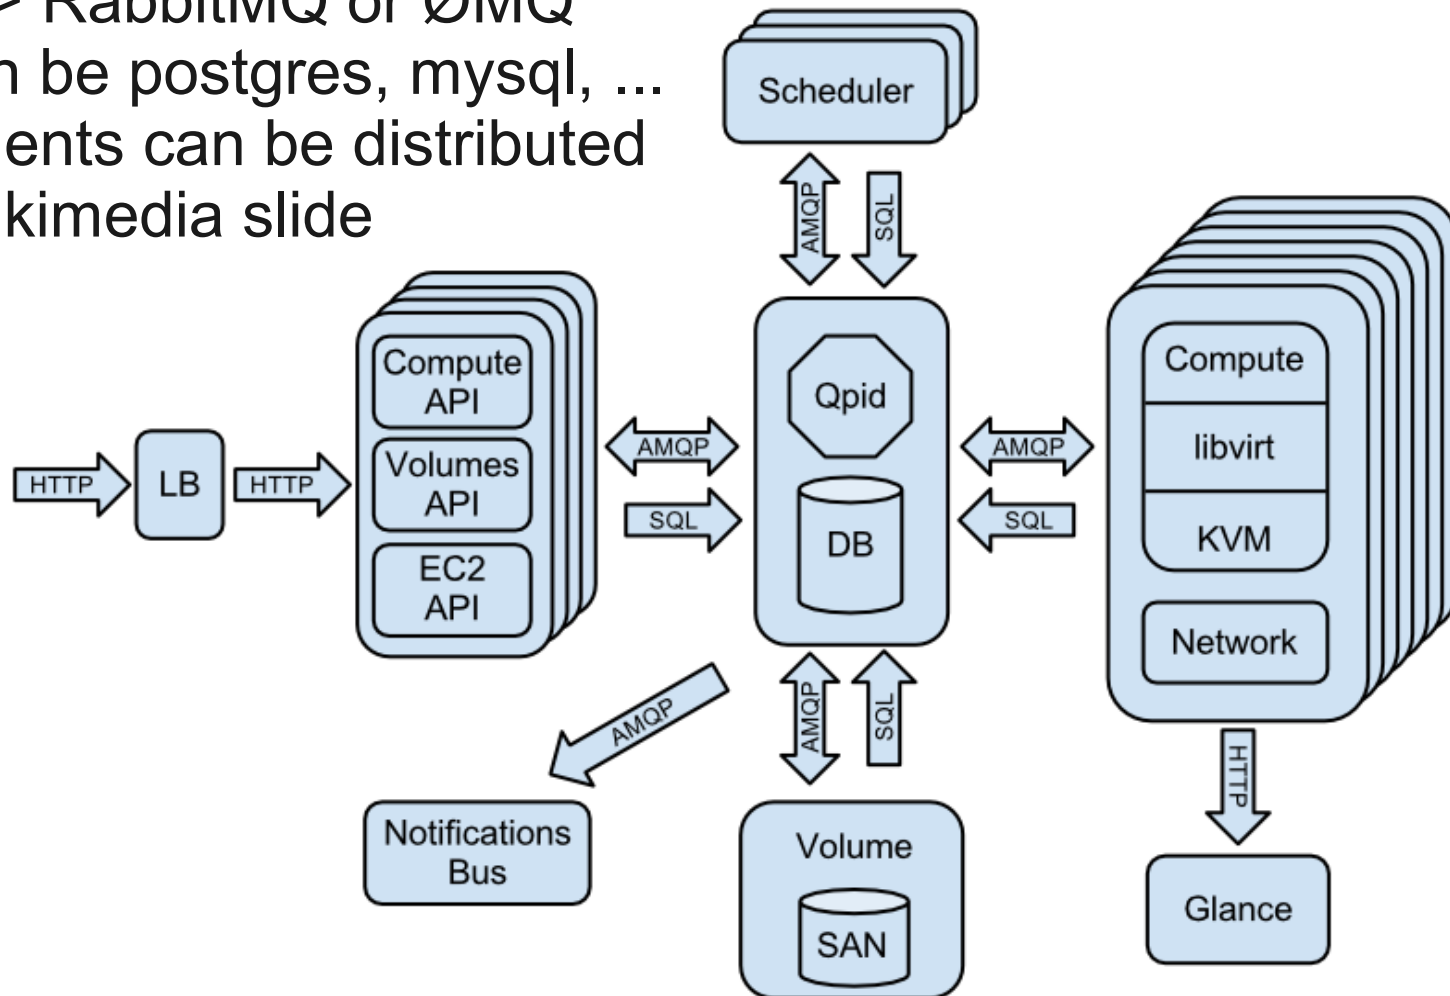
- Nova akin to EC2 – Controls VMs
- Swift akin to S3 – Object storage
- Glance – VM image, registration service
- Keystone – Identity/Auth
- Horizon – Self service and Admin UI
- Quantum – Advanced networking
- Volume akin to EBS – Storage service





Nova Compute Service

- Pluggable components
 - KVM -> XenServer or LXC, ...
 - Qpid -> RabbitMQ or ØMQ
 - DB can be postgres, mysql, ...
- Components can be distributed
 - See wikimedia slide



Scalability notes

- Services use eventlet and greenthread to support concurrency
 - Some explicit locking required in various places
- Can scale almost everything horizontally
- There is a single AMQP/DB though per "cell"
 - Supports 500-1000 hosts
- Several "cells" (zones) can be used for further scalability

Who develops OpenStack

- A rough estimate from the Nova project:

```
$ curl -s https://raw.githubusercontent.com/openstack/nova/master/Authors |  
sed -n 's/.*<.*@\(.*\)>/\1/p' |  
sort | uniq -c | sort -n |  
tail -n20 | column
```

2 midokura.jp	4 ubuntu.com
2 netapp.com	5 pistoncloud.com
2 wikimedia.org	6 isi.edu
3 canonical.com	6 nttdata.co.jp
3 cloudscaling.com	6 openstack.org
3 dreamhost.com	8 hp.com
3 eu.citrix.com	9 redhat.com
3 mirantis.com	10 citrix.com
3 nicira.com	22 rackspace.com
4 lab.ntt.co.jp	31 gmail.com

Python technologies used

```
$ for proj in nova quantum glance keystone horizon; do
  cat $proj/tools/pip-requires
done | sed '/ *#/d; /^ *$/d; s/[>=].*//' |
sort -f | uniq -ci | column | expand
```

1 amqplib	2 iso8601	1 python-cloudfiles
2 anyjson	1 jsonschema	1 python-daemon
1 argparse	2 kombu	1 python-dateutil
1 Babel	1 lockfile	2 python-gflags
2 boto	4 lxml	1 python-novaclient
1 carrot	1 netaddr	4 routes
1 Cheetah	1 pam	4 sqlalchemy
1 Django	1 paramiko	3 sqlalchemy-migrate
4 eventlet	2 passlib	1 suds
1 feedparser	4 paste	4 webob
1 glance	4 PasteDeploy	2 wsgiref
3 greenlet	2 pycrypto	1 xattr
1 httplib2	1 pysendfile	1 python-libvirt

Project Packaging

- Given the number of parts and options there is a large gain with consuming through distros
- Distro package dependencies are a good way to get overall position and leverage of a project

```
yum install rpmorphan graphviz
for pkg in nova glance dashboard keystone quantum quantum-openvswitch;
do
    rpmdep -dot $pkg.dot openstack-$pkg
    dot -Tsvg $pkg.dot -o $pkg.svg
done
```


Release History

- 6 monthly, design summit after each, monthly milestones
- **A**ustin Oct 2010
 - Nova + Swift
- **C**actus Apr 2011
 - Glance
- **D**iablo Sep 2011
 - First “production ready” release
- **E**ssex Apr 2012
 - Horizon + Keystone
- **F**olsom Sep 2012
 - Quantum + Melange + Volume

Development

- Development done on always open trunk
- Releases done to a stable branch
- Standardized tools across all projects
- Git
- Gerrit
 - Standalone patch review server
 - Created by Google for Android development
- Jenkins
 - Continuous integration
 - Gates commits on unit tests

Commit process

- git branch; git commit
 - Commit hook inserts Change-Id: for gerrit tracking
- ./run_tests.sh [subset]
 - Unit tests within a virtualenv
 - Nose used to extend unittest
 - Pep8 enforced <https://github.com/jcrocholl/pep8>
- git-review python tool used to submit to gerrit
 - Developed for OpenStack but not specific to it



Search for status:open

ID	Subject	Owner	Project	Branch	Updated	V	R	A
★ I9e332a79	OVS plugin support for v2 Quantum API	Aaron Rosen	openstack/quantum	master (bug/1011467)	1:10 AM	+1		
★ I3bb1b83b	Implement IP address allocation.	garyk	openstack/quantum	master (bug/1008029)	1:09 AM		✓	✓
★ I1d7304ca	Switch to common logging.	andrewbogott	openstack/nova	master (another)	12:59 AM	+1	✓	
★ I5eee3389	use import_object_ns for compute_driver loading	Sean Dague	openstack/nova	master (bp/virt-driver-cleanup)	12:57 AM	+1	✓	
★ I8d535cb6	Updated files from openstack-common.	andrewbogott	openstack/nova	master (another)	12:56 AM	+1	✓	
★ Ie600b0f0	Adds additional volume.api class for cinder	sleepsonthefloor	openstack/nova	master (8073)	12:56 AM	+1	✓	
★ I80d7a19f	Switch CLI to support underscores and dashes.	Dan Prince	openstack/python-glanceclient	master (bug/1018467)	12:55 AM	+1	✓	
★ I7875c58d	Include all tests in generated tarballs	Thierry Carrez	openstack/glance	master (bug/1018830)	12:50 AM	-1	✓	
★ Ia71d9967	Add check for no domains in libvirt driver.	David Shrewsbury	openstack/nova	master (count_domains)	12:22 AM	+1	✓	
★ I24f189fd	Add links to change pages on /zuul/status page.	Clark Boylan	openstack-ci/zuul	master (change_links_in_html_status)	12:18 AM	+1		
★ If3c62965	implement dhcp agent for quantum	markmcclain	openstack/quantum	master (bp/quantum-dhcp)	12:13 AM	-1	-1	
★ Ic356dcb5	Add 'allocation_pools' parameter to the subnet resource in the Quantum v2 API. H (WORKINPROGRESS)	Salvatore Orlando	openstack/quantum	master (bug/1016308)	Jun 29	+1		
★ I22dd1009	Use ip instead of ifconfig. Should work on F16, F17 and Debian based systems.	ayoung	openstack-dev/devstack	master (master)	Jun 29	+1	-1	
★ Ib362c913	Additional CommandFilters to fix rootwrap on SLES	rhafer	openstack/nova	stable/essex (bug/1013147)	Jun 29	-1	+1	
★ Ia0eec800	If no password in env or command line, try prompting.	Ken Thomas	openstack/python-openstackclient	master (bp/password-prompt)	Jun 29	+1		
★ I34cc9255	Add support for hypervisor-uptime.	Kevin L. Mitchell	openstack/python-novaclient	master (hyper-uptime)	Jun 29	+1		
★ Ie2595897	Get hypervisor uptime.	Kevin L. Mitchell	openstack/nova	master (hyper-uptime)	Jun 29	+1		
★ Iaf77b830	devstack support for v2 nova/quantum integration	Aaron Rosen	openstack-dev/devstack	master (bug/1017760)	Jun 29	+1		
★ Id43175fd	Add tests for volume attach and detach	Dan Smith	openstack/tempest	master (master)	Jun 29	+1	+1	
★ I6596dae7	Redefine the domain's XML on volume attach/detach	Dan Smith	openstack/nova	master (bug/1004791)	Jun 29	+1	-1	
★ Id7d6d547	Do not use pickle for serialization in memcache, but JSON	Vincent Untz	openstack/swift	master (bug/1006414)	Jun 29	-1	-1	
★ I6796b027	change metadata_host to metadata_host_ip	Matt Joyce	openstack/nova	master (master)	Jun 29	+1		
★ I363d54ea	Refactoring code to kernel Dom0 plugin.	Rick Harris	openstack/nova	master (kernel_dom0_plugin)	Jun 29	+1	✓	
★ Icab9551c	Fixes bug 989920	rsokolkov	openstack/nova	stable/essex (bug/989920)	Jun 29	+1	+1	
★ Ib7f9292e	Make git fetching in Gerrit-git-prep more robust.	Clark Boylan	openstack/openstack-ci-puppet	master (retry_fetches)	Jun 29	+1		

Next⇒

Press '?' to view keyboard shortcuts

Change I3e55fcc: Add compatibility for CPU model config with libvirt < 0.9.10

Change-Id:	I3e55fccf38be18454f810efeb5c201aab94ddb0
Owner	Daniel Berrange
Project	openstack/nova
Branch	master
Topic	bug/1003373
Uploaded	Jun 19, 2012 3:47 PM
Updated	Jun 29, 2012 11:13 AM
Status	Review in Progress

Permalink

Add compatibility for CPU model config with libvirt < 0.9.10

Libvirt versions prior to 0.9.10 do not support the simpler 'mode' attribute for choosing CPU model. So with such libvirt releases we must explicitly construct the guest CPU model from the host capabilities info when mode=host-model. We can not support mode=host-passthrough at all with these earlier libvirt versions

This change can be reverted in the future, if the value of the MIN_LIBVIRT_VERSION constant is increased to 0.9.10 or later

Fixes: [bug #1003373](#)
 Implements: [blueprint libvirt-xml-cpu-model](#)
 Change-Id: [I3e55fccf38be18454f810efeb5c201aab94ddb0](#)
 Signed-off-by: Daniel P. Berrange <berrange@redhat.com>

Reviewer	Verified	Code-Review	Approved
Daniel Berrange			
Jenkins	+1		
SmokeStack			
Johannes Erdfelt			
Mark McLoughlin			
Matt Dietz			

- Need Verified
- Need Code-Review
- Need Approved

► Dependencies

Old Version History:


- **Patch Set 1** [0433ffff2d384518a781883b65e34efd9bf91d8f \(gitweb\)](#)
- **Patch Set 2** [7f13d99ae9c3a1c7f73bcd41303ffe5cc308a7c5 \(gitweb\)](#)
- **Patch Set 3** [5bbbe5d01cf0118121de7793ada9882cda08de4c \(gitweb\)](#)
- **Patch Set 4** [ba6dcb4003acf03e79de82ba1348fb5aea65123b \(gitweb\)](#)
- **Patch Set 5** [99067aff3fb2d9ea396f98543cc7b57f83bcf795 \(gitweb\)](#)
- ▼ **Patch Set 6** [36c15fbd05e8509c78e9cdc64cc0cb6725c65311 \(gitweb\)](#)

Review Process

- Gated trunk
 - Reviewers
 - Anyone can ± 1
 - Only project core members can approve
 - Automated tests
 - Through Jenkins
 - Using OpenStack infrastructure itself
 - Sets up test systems and runs gerrit branch
 - Libvirt, Xen, python 2.6, python 2.7, ...
 - If OK, merges gerrit branch to trunk
 - Ancillary jobs like doc builds, pylint reports, more tests, ...

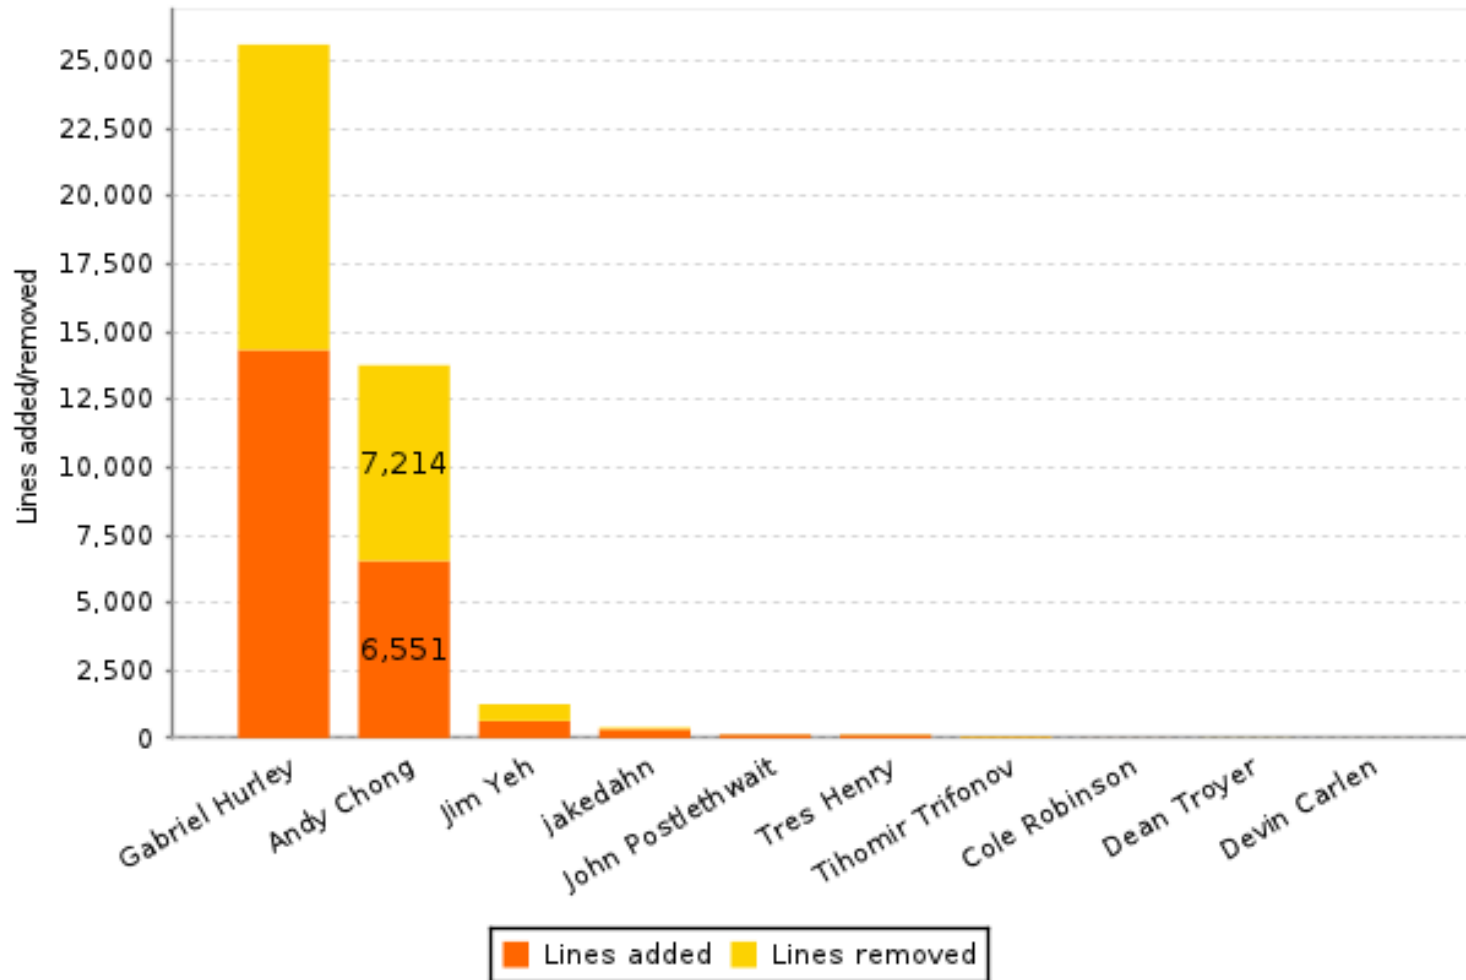
Glance activity week 10 2012

Weekly history of changes per author




Horizon activity week 10 2012

Weekly history of changes per author



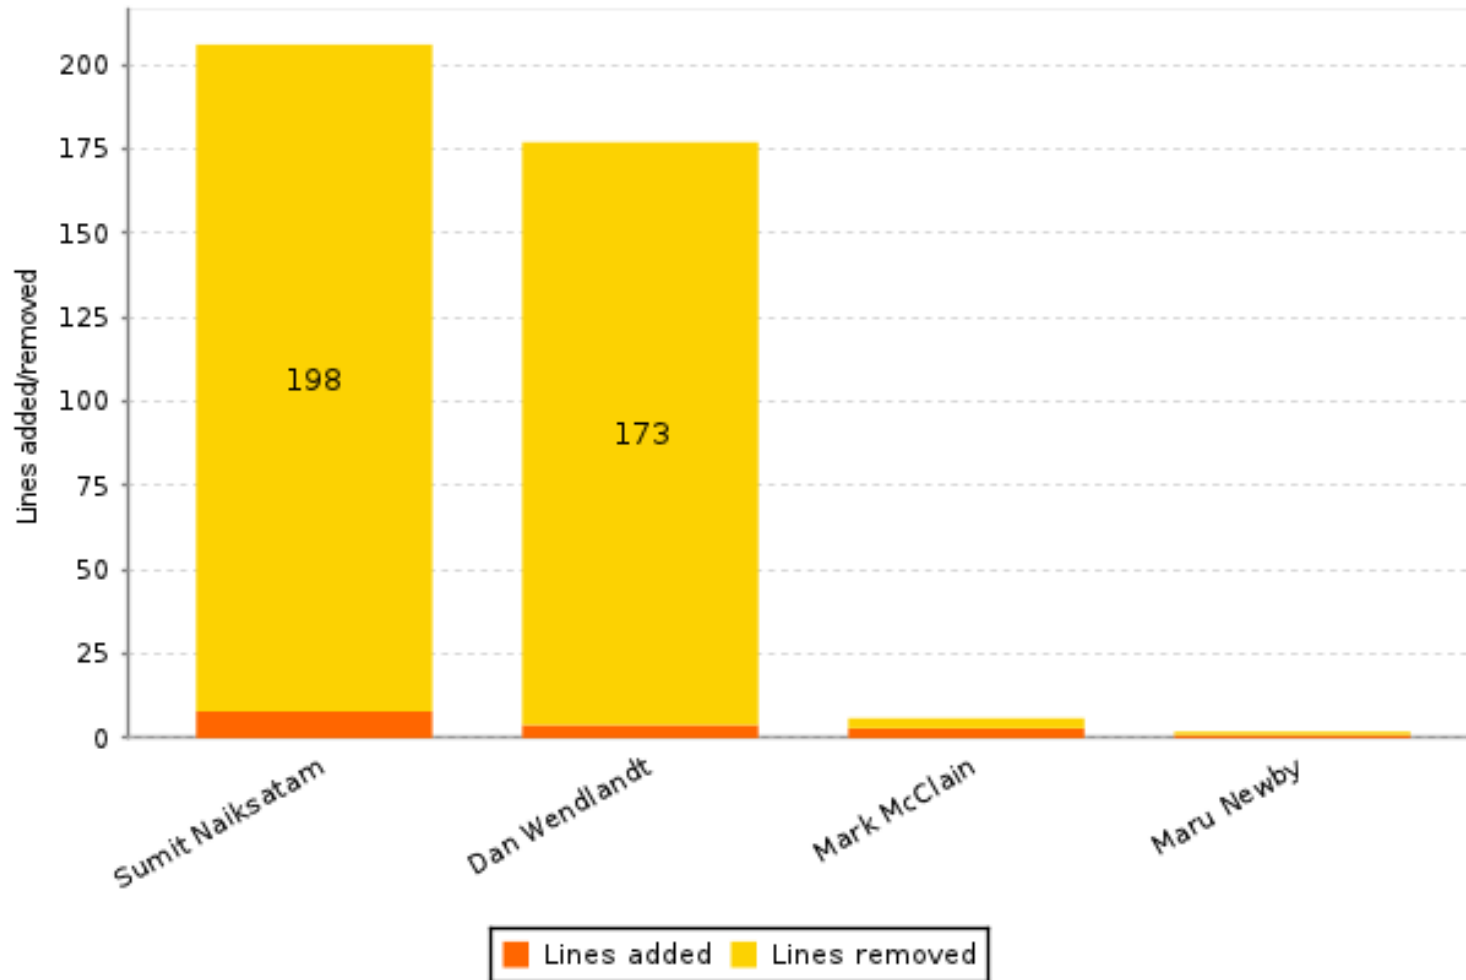
Keystone activity week 10 2012

Weekly history of changes per author




Quantum activity week 10 2012

Weekly history of changes per author




Swift activity week 10 2012




Tempest activity week 10 2012

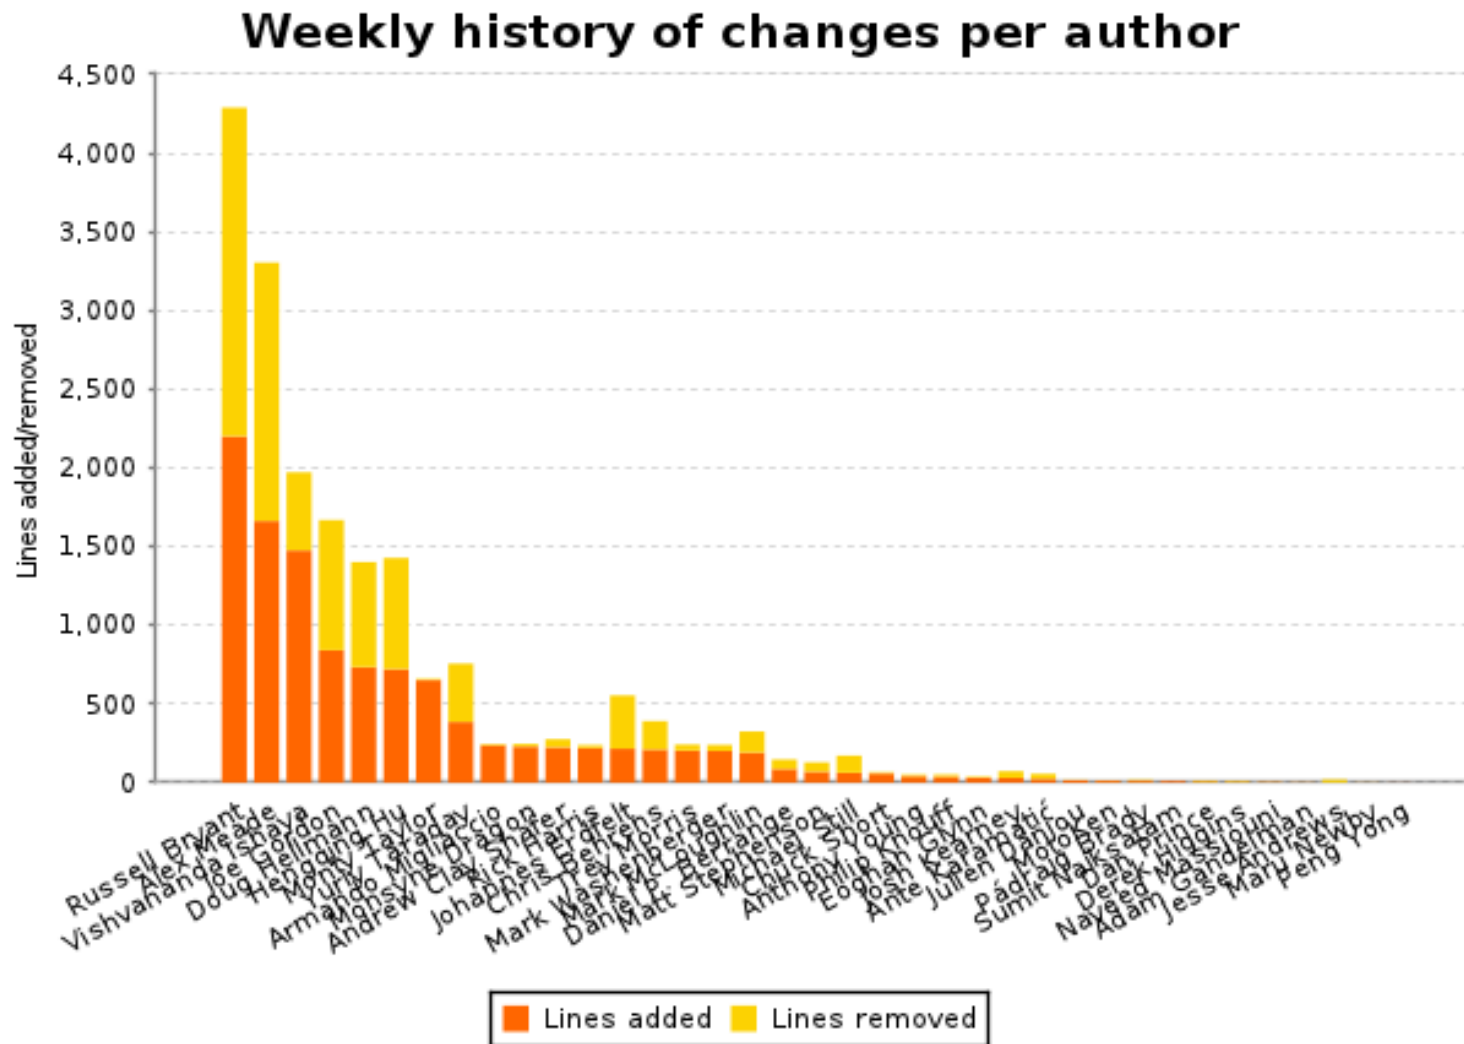
Weekly history of changes per author



Manuals activity week 10 2012



Nova activity week 10 2012



Related python projects

- OZ
 - Generate guest images from install ISOs
 - Optional customization
 - Install packages, run script, ...
 - List of OZ built images from RackSpace
 - <https://github.com/rackerjoe/oz-image-build>
 - ISO -> image -> glance -> nova -> libvirt -> KVM
 - <http://aeolusproject.org/oz.html>

Related python projects

- Heat API
 - AWS Cloudformation type service
 - V4 just released
 - Provisions apps in the cloud
 - And keeps them provisioned
 - Leverages cloud attributes to provide HA
 - <https://github.com/heat-api>

Related python projects

- cloud-init
 - Installed in guests to configure at boot time
 - Directives for package install etc.
 - Support for many distros
 - <https://launchpad.net/cloud-init/>

Similar non python projects

- **oVirt** - java
 - More aligned with private datacenters
- **Eucalyptus** - C
 - Less general, Closed editions.
- **CloudStack** - java
 - Started closed (cloud.com). Closed addons.
- **OpenNebula** - C++
 - Again focusing on datacenters

Trying it out

- <http://trystack.org/>
 - Register to try out a test public OpenStack cloud
- <http://devstack.org/>
 - Script to setup simple/developer installs
- fedoraproject.org/wiki/Getting_started_with_OpenStack
 - Basic setup instructions for machine or VM

Jobs!

- Broad range of technologies
- <http://www.openstack.org/community/jobs/>
 - Currently 90 positions listed
 - Red Hat have 8 location flexible positions there

More info

- <http://docs.openstack.org/>
 - The compute admin guide is a good overview
- <http://ci.openstack.org/>
 - OpenStack Continuous Integration info
- <http://review.openstack.org/>
 - OpenStack gerrit interface
- <http://jenkins.openstack.org/>
 - OpenStack Continuous Integration interface