OpenStack Compute

IWAMOTO Toshihiro

OSS Technology Section II
OSS Platform Technology Center
Business Strategy Group
SCSK

OpenStackについて

- Rackspace Hosting社と NASA により2010年に設立された
- Community based な開発
- Apache 2.0 ライセンス
- “Essex” が2012年4月にリリースされた
  - 次は “Folsom”
  - その次は “Grizzly”

Cloud Computingつて

(いろんな人が定義をしてると思いますが)

- Web-ベース GUI から仮想マシンを(いつでも、任意の数だけ)起動できる
- 仮想マシンの不揮発ストレージ
- 対障害性

  RegionやAvailability Zoneの概念
- 便利機能(DB, firewall,固定IPアドレス、モニタリング、auto scaling, ...)

OpenStack Computeとは

- クラウド構築用ソフトウェア
- Amazon EC2 互換APIあり
- マルチテナント
- 複数ハイパバイザ対応
  - KVM, LXC, QEMU, UML, VMWare, Xen
- 設計の特徴
  - Component based, Highly available, Fault-Tolerant
- Nova はプロジェクト名

User Stories » OpenStack Open Source Cloud Computing Software - Iceweasel

www.openstack.org/user-stories/

Iceweasel | User Stories » OpenStack Open ...

OpenStack User Stories

Read stories about companies using OpenStack to get real work done. (Would you like to be listed here?)

Explore

NeCTAR

Government & Research

Explore

San Diego Supercomputer Center

Academic & Research

Explore

MercadoLibre, Inc.

e-Commerce

Explore

Safew Labs

IT services, software development, cloud services

More companies using OpenStack:

Mirantis

Professional Services & Training

Choopa.com

Hosting & Cloud Computing

Canonical

Open Source Software & Services

eNovance

Cloud & Managed Services Provider

OpenStack関連製品

- Rackspace Open Cloud

  public, private, hybrid

- HP Cloud

  フリートライアル (http://hpcloud.com/)

- SuSE Cloud (private cloud)

- Piston Enterprise OpenStack (private cloud)/

- Airframe (無償版)

OpenStack Computeの構成

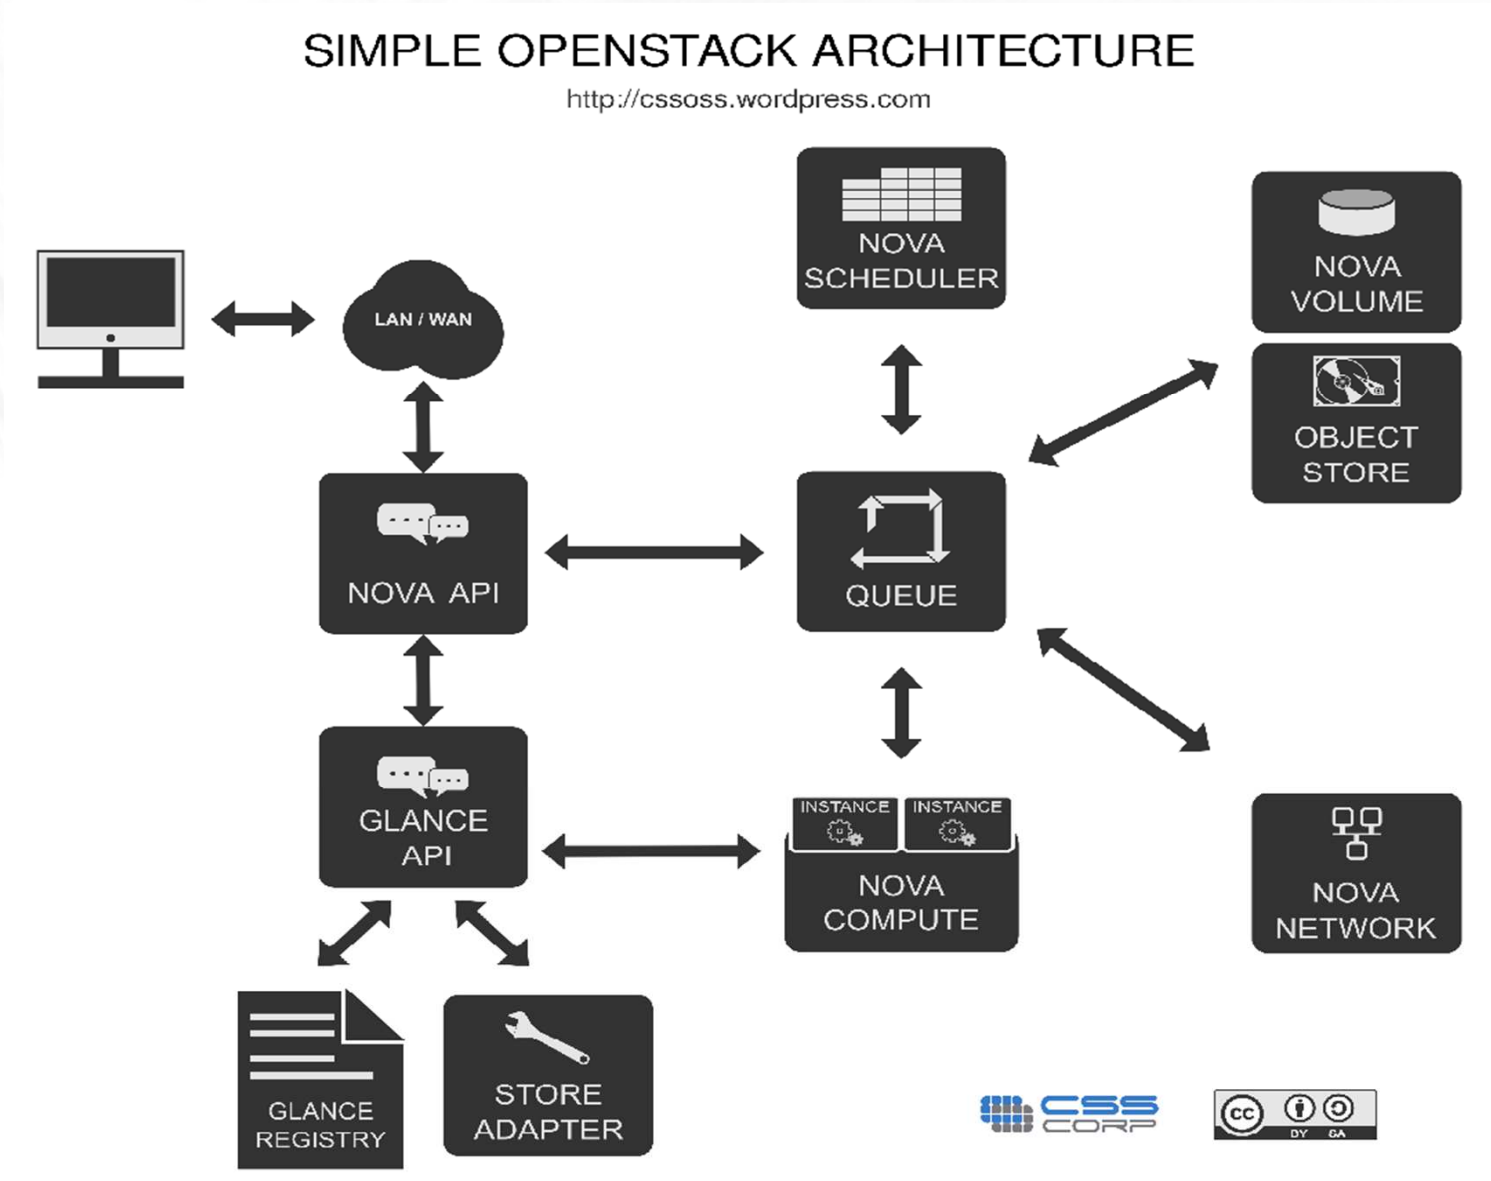

OpenStack のプロジェクト

- Object Storage (Swift)
- Compute (Nova)
- Image Service (Glance)
- Dashboard (Horizon)
- Identity (Keystone)
- Quantum
- Melange
- 詳しくは後程...

OpenStackプロジェクトの開発スタイル (1)

- Launchpad.net
  - Blueprint (開発する機能の管理)
  - バグ管理
- ソースコードはGithub
  - 昔は bzr だった
- “Design Summit”
  - 次のリリースでの機能の議論
  - http://wiki.openstack.org/Summit

OpenStackプロジェクトの開発スタイル(2)

- Gerrit (code review system)
- Jenkins (continuous integration)
  - Gerrit に patch set が追加される度に Jenkins から unit test が実行される
- Tempest
- メイリングリストがあるが実際にはIRCで物事が決定されているらしい

ドメイン別commiter数

Nova

59 @gmail.com
32 @rackspace.com
15 @hp.com
11 @citrix.com
8 @redhat.com
8 @intel.com
7 @griddynamics.com
6 @openstack.org
6 @isi.edu
5 @ubuntu.com
5 @pistoncloud.com
5 @lab.ntt.co.jp
5 @canonical.com

Glance

17 @rackspace.com
16 @gmail.com
6 @hp.com
6 @redhat.com
3 @suse.de
3 @griddynamics.com
3 @canonical.com

Quantum

20 @gmail.com
8 @cisco.com
7 @nicira.com
4 @thoughtworks.com
3 @internap.com

gerrit

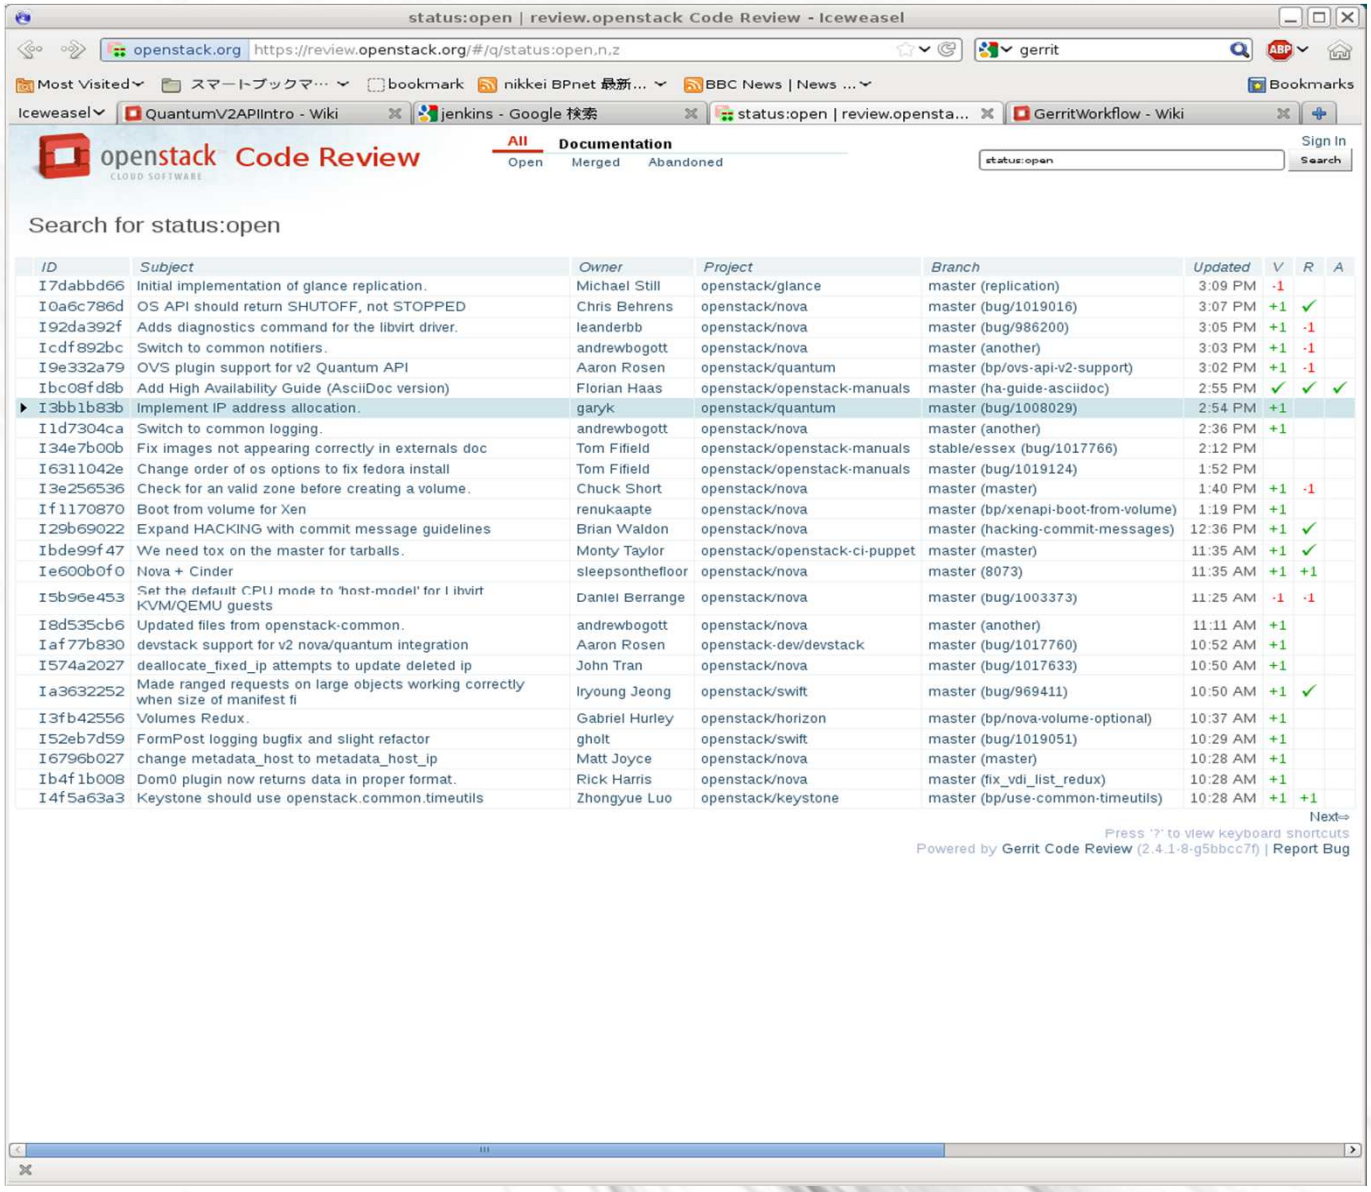
status:open | review.openstack Code Review - Iceweasel
openstack.org https://review.openstack.org/#/q/status:open,n,z
gerrit
Most Visited
bookmark
nikkei BPnet 最新...
BBC News | News ...
Bookmarks
Iceweasel
QuantumV2APIIntro - Wiki
jenkins - Google 検索
status:open | review.opensta...
GerritWorkflow - Wiki
openstack Code Review
All
Documentation
Open
Merged
Abandoned
Sign In
status:open
Search
Search for status:open
ID
Subject
Owner
Project
Branch
Updated
V
R
A
I7dabbd66 Initial implementation of glance replication. Michael Still openstack/glance master (replication) 3:09 PM -1
I0a6c786d OS API should return SHUTOFF, not STOPPED Chris Behrens openstack/nova master (bug/1019016) 3:07 PM +1 ✓
I92da392f Adds diagnostics command for the libvirt driver. leanderbb openstack/nova master (bug/986200) 3:05 PM +1 -1
Icdf892bc Switch to common notifiers. andrewbogott openstack/nova master (another) 3:03 PM +1 -1
I9e332a79 OVS plugin support for v2 Quantum API Aaron Rosen openstack/quantum master (bp/ovs-api-v2-support) 3:02 PM +1 -1
Ibc08fd8b Add High Availability Guide (AsciiDoc version) Florian Haas openstack/openstack-manuals master (ha-guide-asciidoc) 2:55 PM ✓ ✓ ✓
I3bb1b83b Implement IP address allocation. garyk openstack/quantum master (bug/1008029) 2:54 PM +1
I1d7304ca Switch to common logging. andrewbogott openstack/nova master (another) 2:36 PM +1
I34e7b00b Fix images not appearing correctly in externals doc Tom Fifield openstack/openstack-manuals stable/essex (bug/1017766) 2:12 PM
I6311042e Change order of os options to fix fedora install Tom Fifield openstack/openstack-manuals master (bug/1019124) 1:52 PM
I3e256536 Check for an valid zone before creating a volume. Chuck Short openstack/nova master (master) 1:40 PM +1 -1
If1170870 Boot from volume for Xen renukaapte openstack/nova master (bp/xenapi-boot-from-volume) 1:19 PM +1
I29b69022 Expand HACKING with commit message guidelines Brian Waldon openstack/nova master (hacking-commit-messages) 12:36 PM +1 ✓
Ibde99f47 We need tox on the master for tarballs. Monty Taylor openstack/openstack-ci-puppet master (master) 11:35 AM +1 ✓
Ie600b0f0 Nova + Cinder sleepsonthefloor openstack/nova master (8073) 11:35 AM +1 +1
I5b96e453 Set the default CPU mode to 'host-model' for Libvirt KVM/QEMU guests Daniel Berrange openstack/nova master (bug/1003373) 11:25 AM -1 -1
I8d535cb6 Updated files from openstack-common. andrewbogott openstack/nova master (another) 11:11 AM +1
Iaf77b830 devstack support for v2 nova/quantum integration Aaron Rosen openstack-dev/devstack master (bug/1017760) 10:52 AM +1
I574a2027 deallocate_fixed_ip attempts to update deleted ip John Tran openstack/nova master (bug/1017633) 10:50 AM +1
Ia3632252 Made ranged requests on large objects working correctly when size of manifest fi lryoung Jeong openstack/swift master (bug/969411) 10:50 AM +1 ✓
I3fb42556 Volumes Redux. Gabriel Hurley openstack/horizon master (bp/nova-volume-optional) 10:37 AM +1
I52eb7d59 FormPost logging bugfix and slight refactor gholt openstack/swift master (bug/1019051) 10:29 AM +1
I6796b027 change metadata_host to metadata_host_ip Matt Joyce openstack/nova master (master) 10:28 AM +1
Ib4f1b008 Dom0 plugin now returns data in proper format. Rick Harris openstack/nova master (fix_vdi_list_redux) 10:28 AM +1
I4f5a63a3 Keystone should use openstack.common.timeutils Zhongyue Luo openstack/keystone master (bp/use-common-timeutils) 10:28 AM +1 +1
Next⇒
Press '?' to view keyboard shortcuts
Powered by Gerrit Code Review (2.4.1-8-g5bbcc7f) | Report Bug

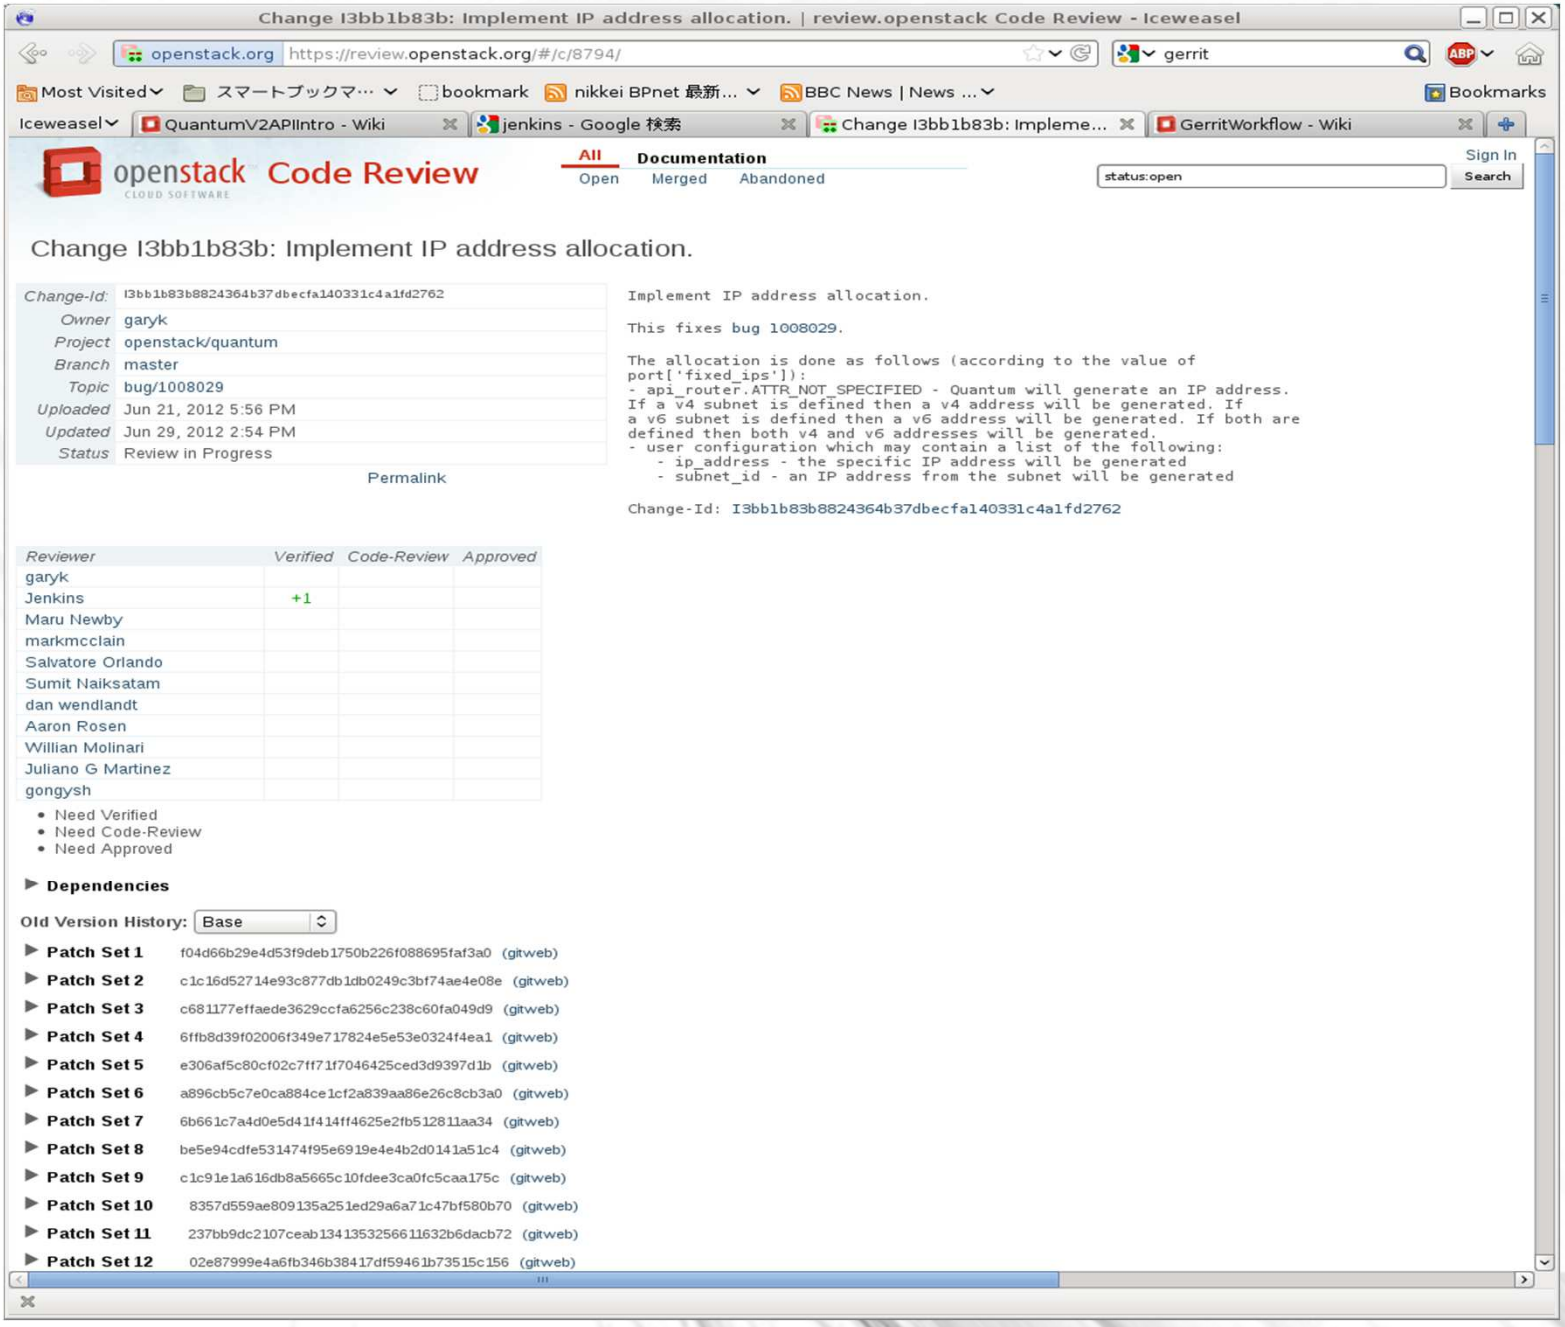
Change I3bb1b83b: Implement IP address allocation. | review.openstack Code Review - Iceweasel
openstack.org https://review.openstack.org/#/c/8794/
gerrit
Most Visited
スマートブックマ…
bookmark
nikkei BPnet 最新…
BBC News | News …
Bookmarks
Iceweasel
QuantumV2APIIntro - Wiki
jenkins - Google 検索
Change I3bb1b83b: Impleme…
GerritWorkflow - Wiki
openstack Code Review
CLOUD SOFTWARE
All
Documentation
Open
Merged
Abandoned
Sign In
status:open
Search
Change I3bb1b83b: Implement IP address allocation.
Change-Id: I3bb1b83b8824364b37dbecfa140331c4a1fd2762
Owner garyk
Project openstack/quantum
Branch master
Topic bug/1008029
Uploaded Jun 21, 2012 5:56 PM
Updated Jun 29, 2012 2:54 PM
Status Review in Progress
Permalink
Implement IP address allocation.
This fixes bug 1008029.
The allocation is done as follows (according to the value of
port['fixed_ips']):
- api_router.ATTR_NOT_SPECIFIED - Quantum will generate an IP address.
If a v4 subnet is defined then a v4 address will be generated. If
a v6 subnet is defined then a v6 address will be generated. If both are
defined then both v4 and v6 addresses will be generated.
- user configuration which may contain a list of the following:
- ip_address - the specific IP address will be generated
- subnet_id - an IP address from the subnet will be generated
Change-Id: I3bb1b83b8824364b37dbecfa140331c4a1fd2762
Reviewer Verified Code-Review Approved
garyk
Jenkins +1
Maru Newby
markmcclain
Salvatore Orlando
Sumit Naiksatam
dan wendlandt
Aaron Rosen
Willian Molinari
Juliano G Martinez
gongysh
Need Verified
Need Code-Review
Need Approved
Dependencies
Old Version History: Base
Patch Set 1 f04d66b29e4d53f9deb1750b226f088695faf3a0 (gitweb)
Patch Set 2 c1c16d52714e93c877db1db0249c3bf74ae4e08e (gitweb)
Patch Set 3 c681177effaede3629ccfa6256c238c60fa049d9 (gitweb)
Patch Set 4 6ffb8d39f02006f349e717824e5e53e0324f4ea1 (gitweb)
Patch Set 5 e306af5c80cf02c7ff71f7046425ced3d9397d1b (gitweb)
Patch Set 6 a896cb5c7e0ca884ce1cf2a839aa86e26c8cb3a0 (gitweb)
Patch Set 7 6b661c7a4d0e5d41f414ff4625e2fb512811aa34 (gitweb)
Patch Set 8 be5e94cdfe531474f95e6919e4e4b2d0141a51c4 (gitweb)
Patch Set 9 c1c91e1a616db8a5665c10fdee3ca0fc5caa175c (gitweb)
Patch Set 10 8357d559ae809135a251ed29a6a71c47bf580b70 (gitweb)
Patch Set 11 237bb9dc2107ceab1341353256611632b6dacb72 (gitweb)
Patch Set 12 02e87999e4a6fb346b38417df59461b73515c156 (gitweb)

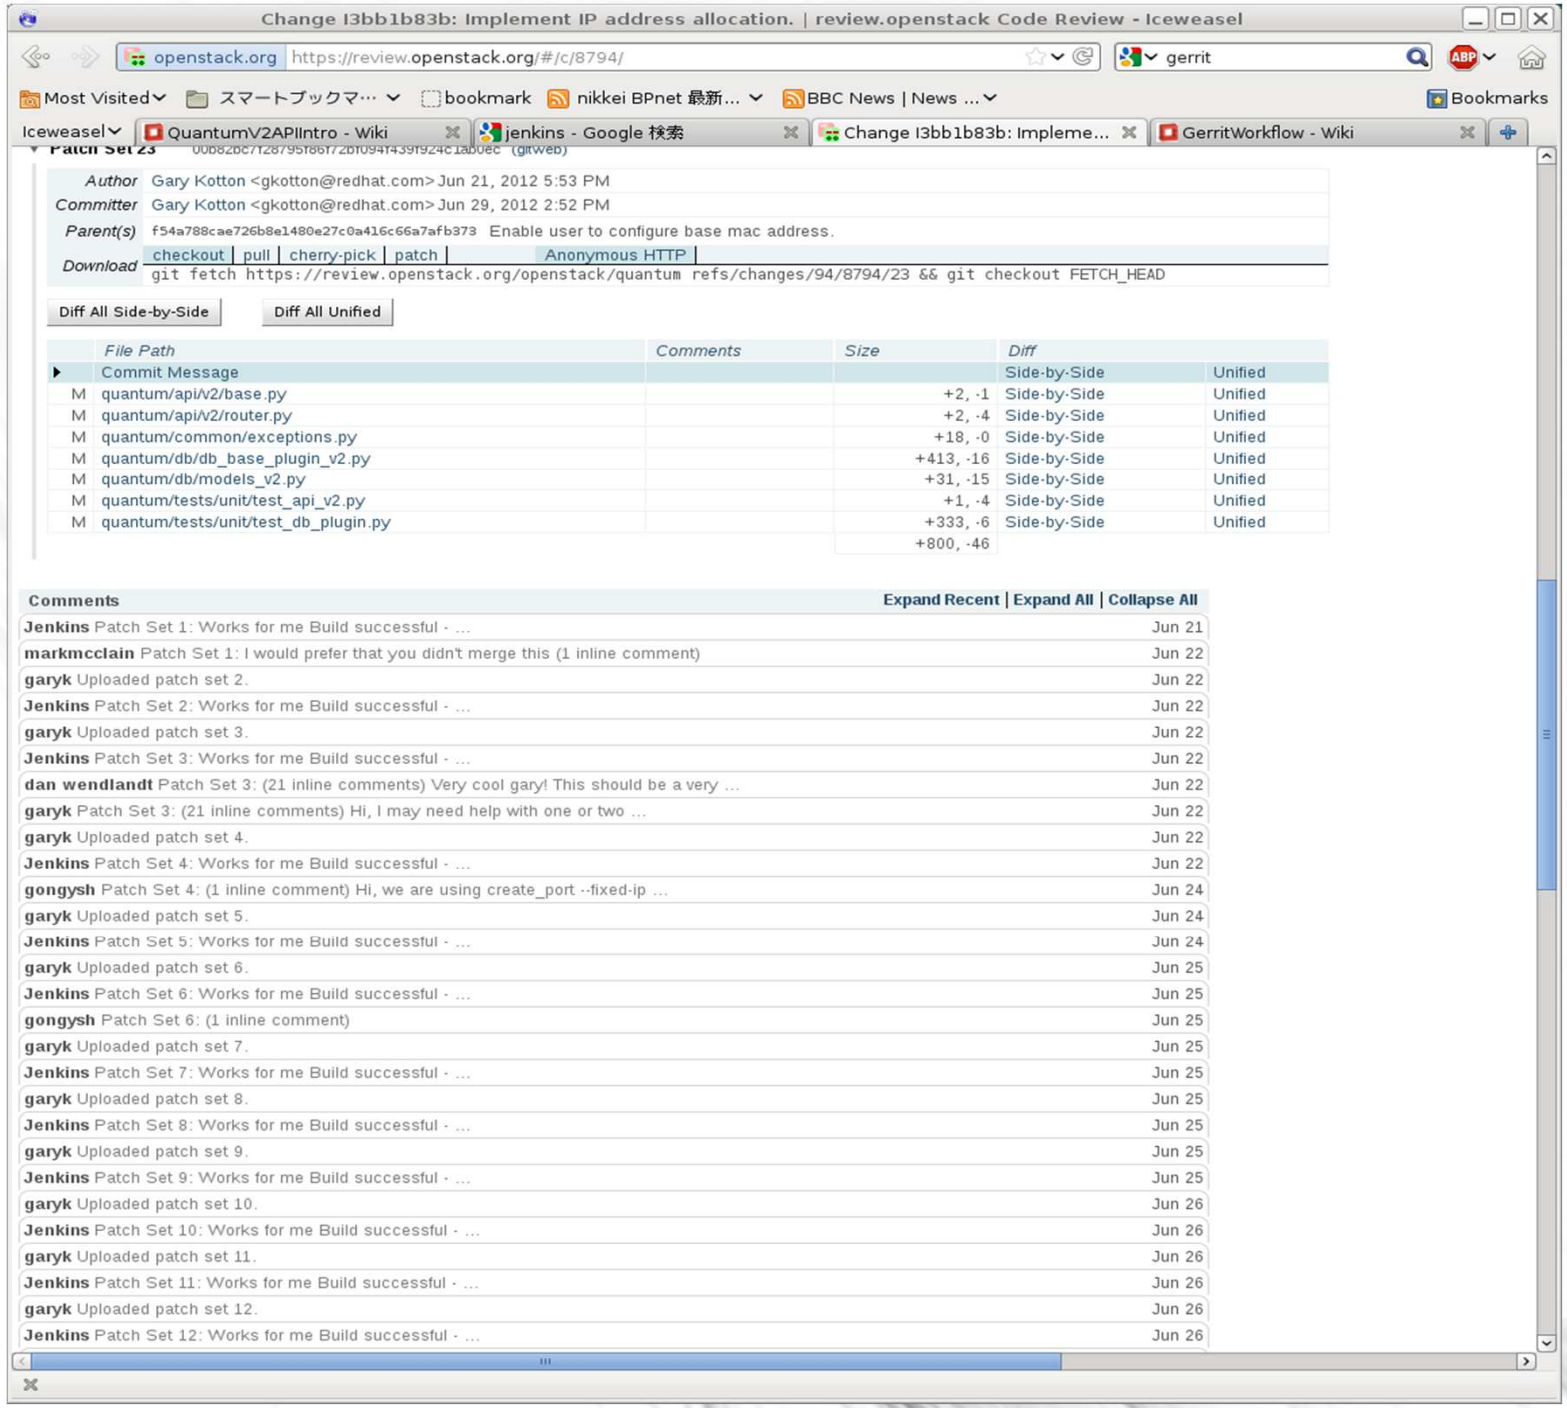
Change I3bb1b83b: Implement IP address allocation. | review.openstack Code Review - Iceweasel
openstack.org https://review.openstack.org/#/c/8794/
gerrit
Most Visited
スマートブックマ…
bookmark
nikkei BPnet 最新…
BBC News | News …
Bookmarks
Iceweasel
QuantumV2APIIntro - Wiki
jenkins - Google 検索
Change I3bb1b83b: Impleme…
GerritWorkflow - Wiki
Patch Set 23
Author Gary Kotton <gkotton@redhat.com> Jun 21, 2012 5:53 PM
Committer Gary Kotton <gkotton@redhat.com> Jun 29, 2012 2:52 PM
Parent(s) f54a788cae726b8e1480e27c0a416c66a7afb373 Enable user to configure base mac address.
Download checkout | pull | cherry-pick | patch | Anonymous HTTP |
git fetch https://review.openstack.org/openstack/quantum refs/changes/94/8794/23 && git checkout FETCH_HEAD
Diff All Side-by-Side
Diff All Unified
File Path Comments Size Diff
Commit Message Side-by-Side Unified
M quantum/api/v2/base.py +2, -1 Side-by-Side Unified
M quantum/api/v2/router.py +2, -4 Side-by-Side Unified
M quantum/common/exceptions.py +18, -0 Side-by-Side Unified
M quantum/db/db_base_plugin_v2.py +413, -16 Side-by-Side Unified
M quantum/db/models_v2.py +31, -15 Side-by-Side Unified
M quantum/tests/unit/test_api_v2.py +1, -4 Side-by-Side Unified
M quantum/tests/unit/test_db_plugin.py +333, -6 Side-by-Side Unified
+800, -46
Comments Expand Recent | Expand All | Collapse All
Jenkins Patch Set 1: Works for me Build successful - … Jun 21
markmcclain Patch Set 1: I would prefer that you didn't merge this (1 inline comment) Jun 22
garyk Uploaded patch set 2. Jun 22
Jenkins Patch Set 2: Works for me Build successful - … Jun 22
garyk Uploaded patch set 3. Jun 22
Jenkins Patch Set 3: Works for me Build successful - … Jun 22
dan wendlandt Patch Set 3: (21 inline comments) Very cool gary! This should be a very … Jun 22
garyk Patch Set 3: (21 inline comments) Hi, I may need help with one or two … Jun 22
garyk Uploaded patch set 4. Jun 22
Jenkins Patch Set 4: Works for me Build successful - … Jun 22
gongysh Patch Set 4: (1 inline comment) Hi, we are using create_port --fixed-ip … Jun 24
garyk Uploaded patch set 5. Jun 24
Jenkins Patch Set 5: Works for me Build successful - … Jun 24
garyk Uploaded patch set 6. Jun 25
Jenkins Patch Set 6: Works for me Build successful - … Jun 25
gongysh Patch Set 6: (1 inline comment) Jun 25
garyk Uploaded patch set 7. Jun 25
Jenkins Patch Set 7: Works for me Build successful - … Jun 25
garyk Uploaded patch set 8. Jun 25
Jenkins Patch Set 8: Works for me Build successful - … Jun 25
garyk Uploaded patch set 9. Jun 25
Jenkins Patch Set 9: Works for me Build successful - … Jun 25
garyk Uploaded patch set 10. Jun 26
Jenkins Patch Set 10: Works for me Build successful - … Jun 26
garyk Uploaded patch set 11. Jun 26
Jenkins Patch Set 11: Works for me Build successful - … Jun 26
garyk Uploaded patch set 12. Jun 26
Jenkins Patch Set 12: Works for me Build successful - … Jun 26

開発動向

- Essexの新機能
  - 共通認証モジュール Keystone
  - Horizon 初リリース
- 次期リリース(Folsom)では
  - Quantumが “core”に。V2 API
  - Glance V2 API
  - Nova-volume が Cinder に
  - no-db-messaging

今日話す内容

- OpenStackの概要
- devstackを使った動作例
- OpenStack Computeを構成するプロセスの説明
- プロセス間の処理の流れ
  - VM起動
- ネットワークについて

devstackを用いた環境構築

- とりあえず動かすにはこれを使うべき(構成要素が多すぎるので手でやるとたいへん)
- http://devstack.org
- 必要なものをまとめてインストールして Compute 各サービスを起動してくれる
  - 設定ファイルも作成してくれる

devstackを用いた環境構築(2)

1. 1台マシン(KVMを使用するため、vt機能は必要)を用意し、ubuntu 12.04をクリーンインストール
2. $ sudo apt-get install git
3. $ git clone git://github.com/openstack-dev/devstack.git
4. $ cd devstack
5. $ git checkout stable/essex
6. $ ./stack.sh

- (大雑把に言って) *api という名前のものはHTTPを受信し、それ以外はAMQP で通信
  - AMQP == Advanced Message Queuing Protocol

devstackを用いた環境構築(3)

- 実行するとこうなる

```
$ ps x|grep -v bash¥$
  PID TTY      STAT   TIME COMMAND
27213 ?        S      0:00 sshd: stack@pts/0
28470 ?        Sl     0:00 /usr/sbin/apache2 -k start
28471 ?        Sl     0:00 /usr/sbin/apache2 -k start
28473 ?        Sl     0:00 /usr/sbin/apache2 -k start
29355 pts/2    S+     0:00 python bin/glance-registry --config-file=/etc/glance/glance-registry.conf
29458 pts/4    S+     0:00 python bin/glance-api --config-file=/etc/glance/glance-api.conf
29630 pts/5    S+     0:01 python /opt/stack-gh/keystone/bin/keystone-all --config-file /etc/keystone/keystone.conf --log-config /etc/keystone/logging.conf -d --debug
29889 pts/6    S+     0:00 python /opt/stack-gh/nova/bin/nova-api
30029 pts/7    S+     0:00 python /opt/stack-gh/nova/bin/nova-compute
30146 pts/8    S+     0:00 python /opt/stack-gh/nova/bin/nova-cert
30273 pts/9    S+     0:00 python /opt/stack-gh/nova/bin/nova-volume
30385 pts/10   S+     0:00 python /opt/stack-gh/nova/bin/nova-network
30527 pts/11   S+     0:00 python /opt/stack-gh/nova/bin/nova-scheduler
30633 pts/12   S+     0:00 python ./utils/nova-novncproxy --config-file /etc/nova/nova.conf --web .
30734 pts/13   S+     0:00 python ./bin/nova-xvpvncproxy --config-file /etc/nova/nova.conf
30836 pts/14   S+     0:00 python ./bin/nova-consoleauth
31039 pts/16   S+     0:00 python /opt/stack-gh/nova/bin/nova-objectstore
```

- （大雑把に言って）*api という名前のものはHTTPを受信し、それ以外はAMQP で通信
  - AMQP == Advanced Message Queuing Protocol

devstackを用いた環境構築(4)

- コード規模（pythonのみの行数）

```
stack@iwamotonova:/opt/stack-gh$ for d in *; do echo -n $d; find $d -name ¥*.py |xargs wc -l|tail -1; done
glance                    35351 total
horizon                   20921 total
keystone                  16105 total
melange                   15406 total
nova                     149938 total
noVNC                      1451 total
python-glanceclient        1896 total
python-keystoneclient      4917 total
python-novaclient          9431 total
stack@iwamotonova:/opt/stack-gh$ for d in *; do echo -n $d; find $d -name ¥*.py -path ¥*tests¥*|xargs wc -l|tail -1; done
glance                    19709 total
horizon                    6048 total
keystone                   4631 total
melange                    7940 total
nova                      60944 total
python-glanceclient         422 total
python-keystoneclient      2118 total
python-novaclient          3165 total
```

使ってみる(1)

- Web browser で port 80に接続する。接続先は dashboard(horizon)

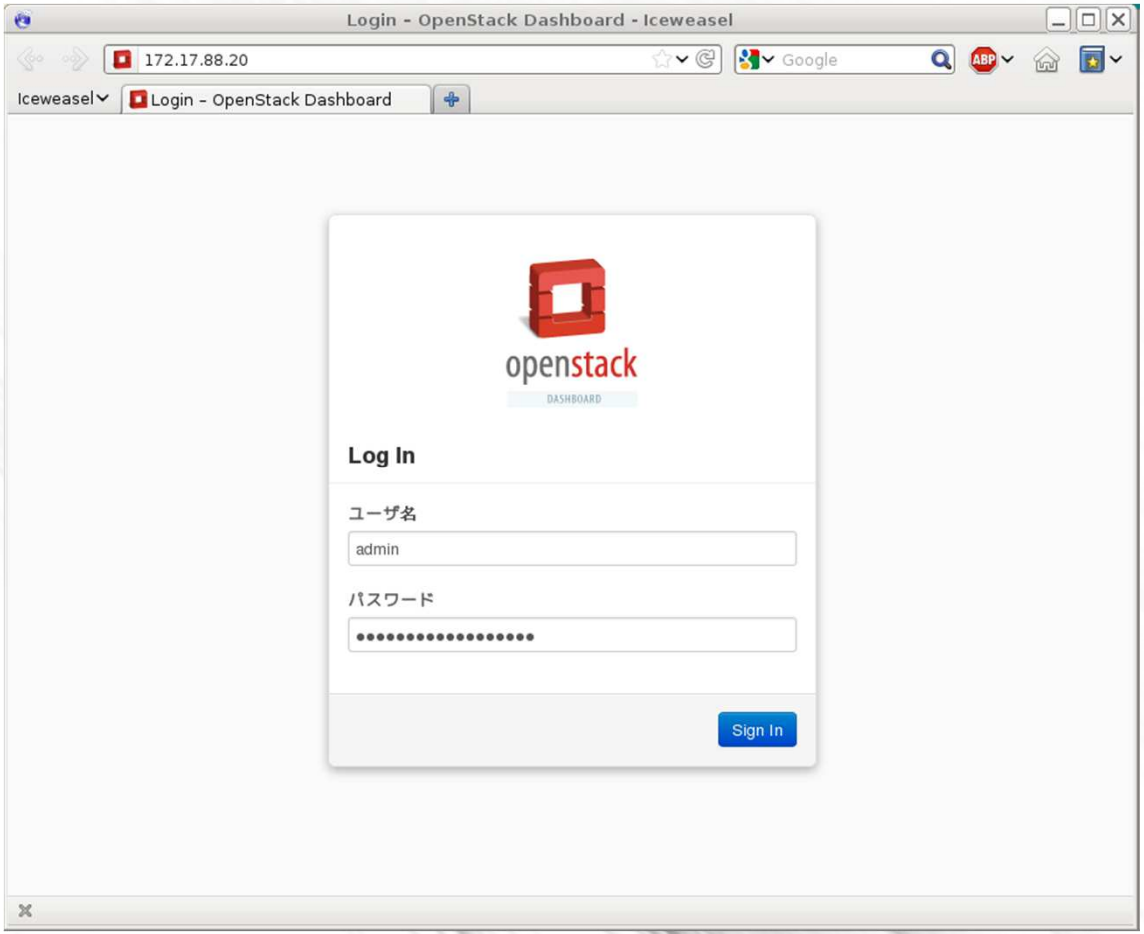

使ってみる(2)

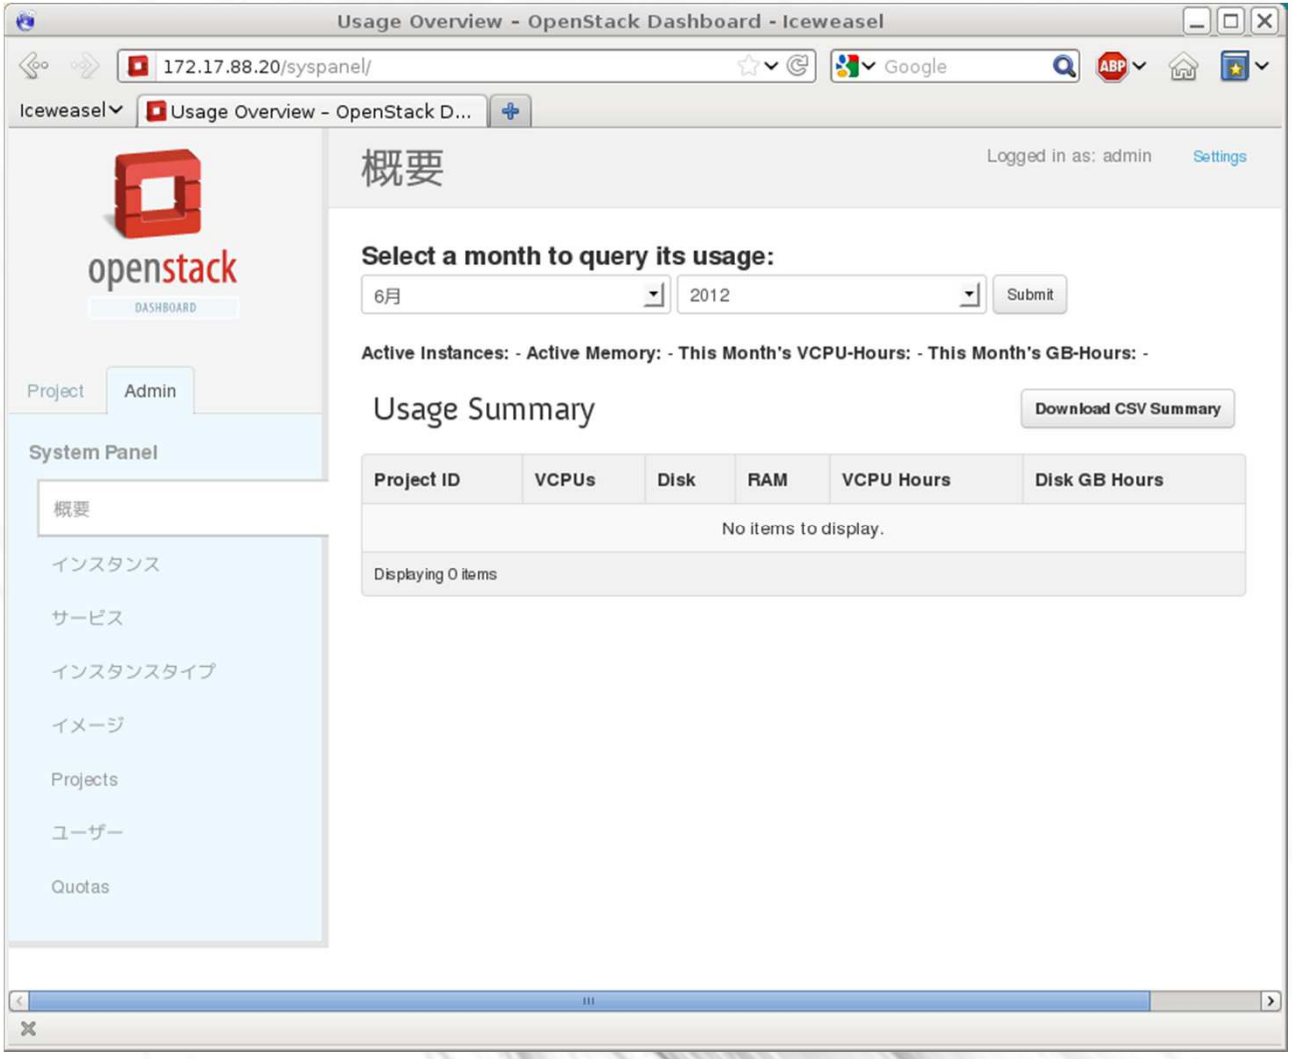
Usage Overview - OpenStack Dashboard - Iceweasel
172.17.88.20/syspanel/
Google
Iceweasel
Usage Overview - OpenStack D...
openstack
DASHBOARD
概要
Logged in as: admin
Settings
Select a month to query its usage:
6月
2012
Submit
Active Instances: - Active Memory: - This Month's VCPU-Hours: - This Month's GB-Hours: -
Usage Summary
Download CSV Summary
Project ID
VCPUs
Disk
RAM
VCPU Hours
Disk GB Hours
No items to display.
Displaying 0 items
Project
Admin
System Panel
概要
インスタンス
サービス
インスタンスタイプ
イメージ
Projects
ユーザー
Quotas

使ってみる(3)

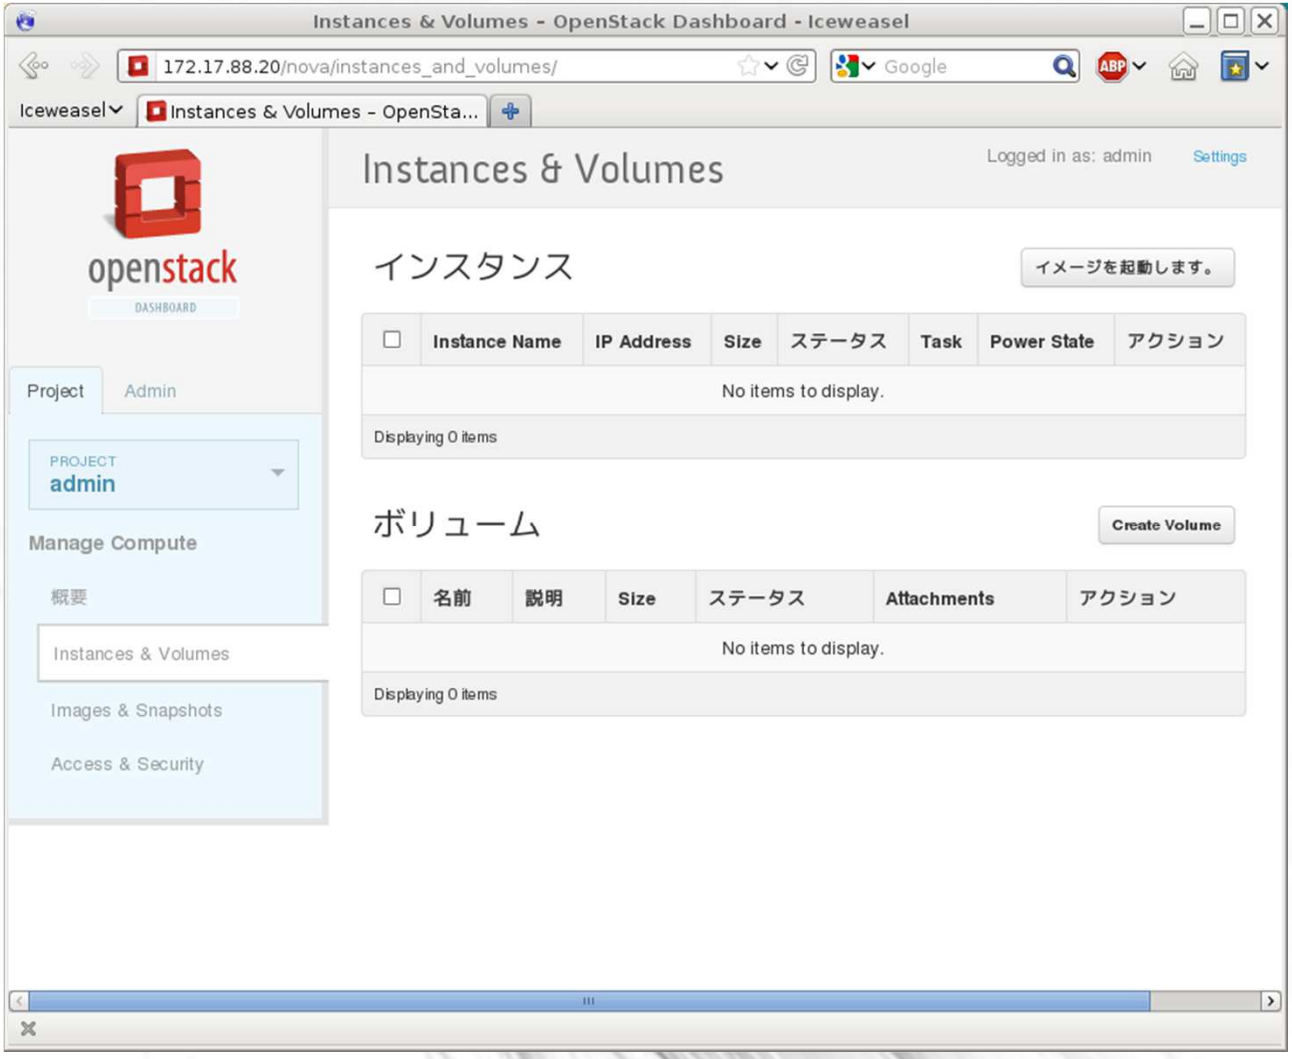
Instances & Volumes - OpenStack Dashboard - Iceweasel
172.17.88.20/nova/instances_and_volumes/
Google
Iceweasel
Instances & Volumes - OpenSta...
openstack
DASHBOARD
Project
Admin
PROJECT
admin
Manage Compute
概要
Instances & Volumes
Images & Snapshots
Access & Security
Instances & Volumes
Logged in as: admin
Settings
インスタンス
イメージを起動します。
Instance Name
IP Address
Size
ステータス
Task
Power State
アクション
No items to display.
Displaying 0 items
ボリューム
Create Volume
名前
説明
Size
ステータス
Attachments
アクション
No items to display.
Displaying 0 items

使ってみる(4)

- VMを起動する（devstackが起動イメージを用意してくれている）

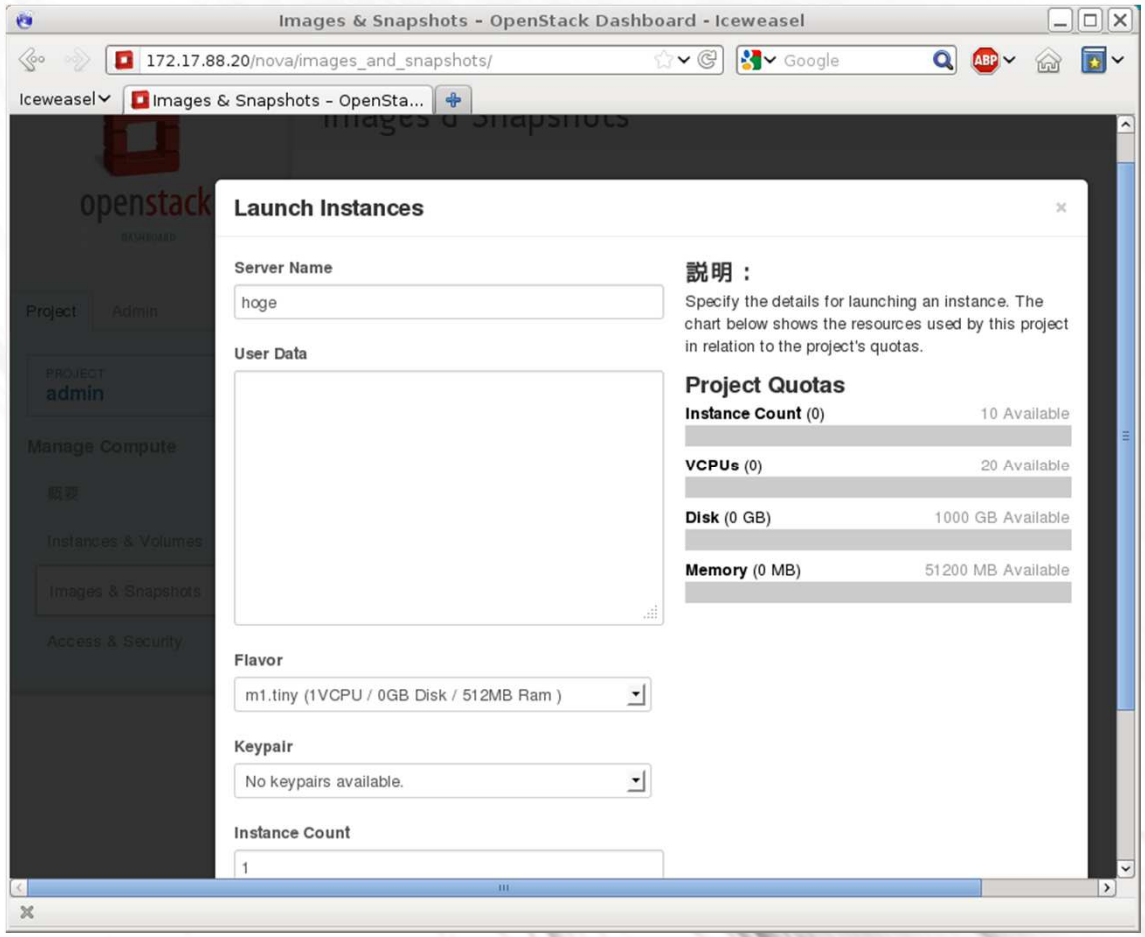

使ってみる(5)

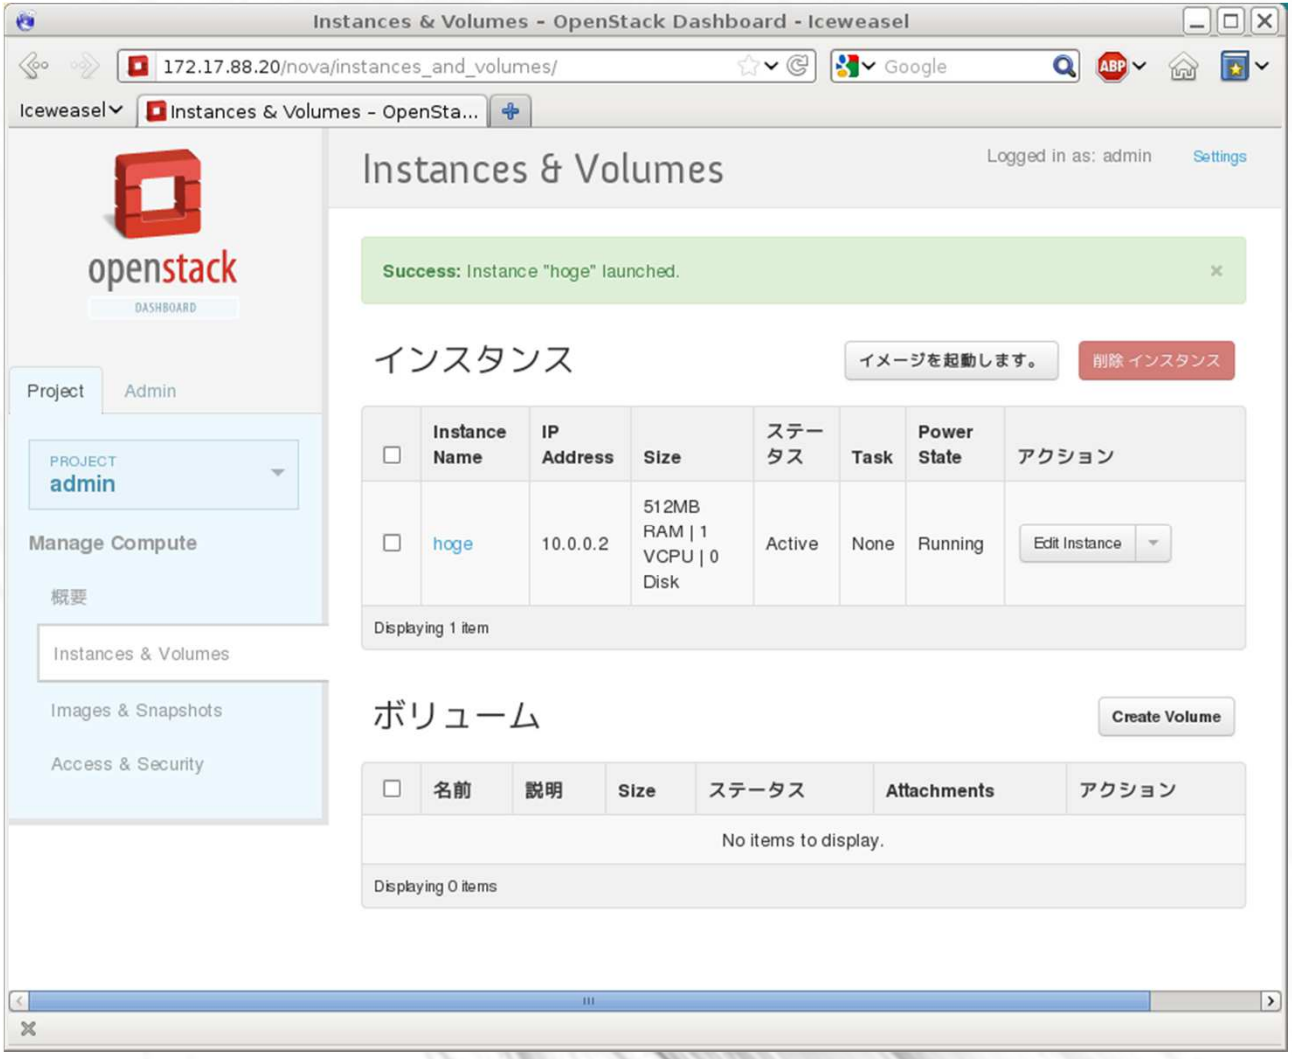
Instances & Volumes - OpenStack Dashboard - Iceweasel
172.17.88.20/nova/instances_and_volumes/
Google
Iceweasel
Instances & Volumes - OpenSta...
openstack
DASHBOARD
Project
Admin
PROJECT
admin
Manage Compute
概要
Instances & Volumes
Images & Snapshots
Access & Security
Instances & Volumes
Logged in as: admin
Settings
Success: Instance "hoge" launched.
インスタンス
イメージを起動します。
削除 インスタンス
Instance Name
IP Address
Size
ステータス
Task
Power State
アクション
hoge
10.0.0.2
512MB RAM | 1 VCPU | 0 Disk
Active
None
Running
Edit Instance
Displaying 1 item
ボリューム
Create Volume
名前
説明
Size
ステータス
Attachments
アクション
No items to display.
Displaying 0 items

使ってみる(6)

- VNC console でVMにログイン

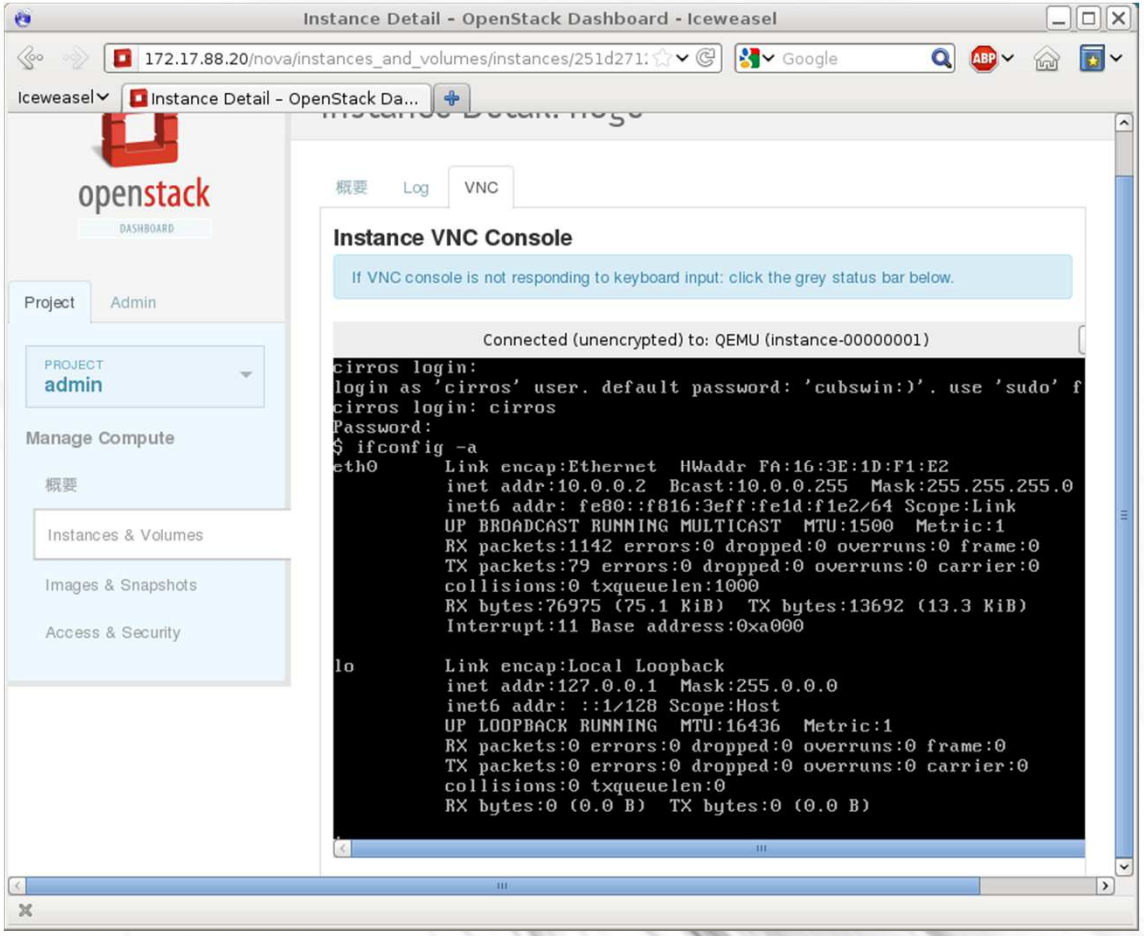

使ってみる(7)

- この他にも以下の操作などが可能
  - VMの停止
  - nova volume の作成、削除、スナップショット作成
- 起動イメージ
  - http://docs.openstack.org/trunk/openstack-compute/admin/content/starting-images.html
- ネットワーク作成はGUIからはできない

使ってみる(8)

- Horizonから使える機能は全体の一部
- devstack/exercises/　にクライアントコマンドの使用例あり
  - クライアントコマンドはpython
  - Jsonボディを組みたててHTTP　RESTリクエストを送信するもの
- euca2ools も使える(上のスクリプトでも使用)
- 試してみてください

何が起きたか(1)

- ユーザがブラウザで Horizon にコマンドを送信
- Webアプリ(horizon)は
  - 受け取ったコマンドをもとにnova-apiにコマンドを送信
  - VMの状態などをユーザに返す
  - django で書かれている

何が起きたか(2)

- nova-apiは nova-scheduler に
- nova-scheduler は nova-compute に
  - 複数台構成のときはここで適当なものが選ばれる
  - 起動時に実装が選べる
    - SimpleScheduler
    - ChanceScheduler
    - FilterScheduler (devstack の default)
    - MultiScheduler (default)

何が起きたか(3)

- Nova-computeは…
  - AMQPから起動命令を受信
  - DatabaseからVMの情報を取得
  - Image serviceから起動イメージを取得
  - ネットワーク作成
  - VM起動
    - libvirt.xml の作成など
  - DatabaseのVMの状態をACTIVEに更新

OpenStack Computeの構成(再掲)

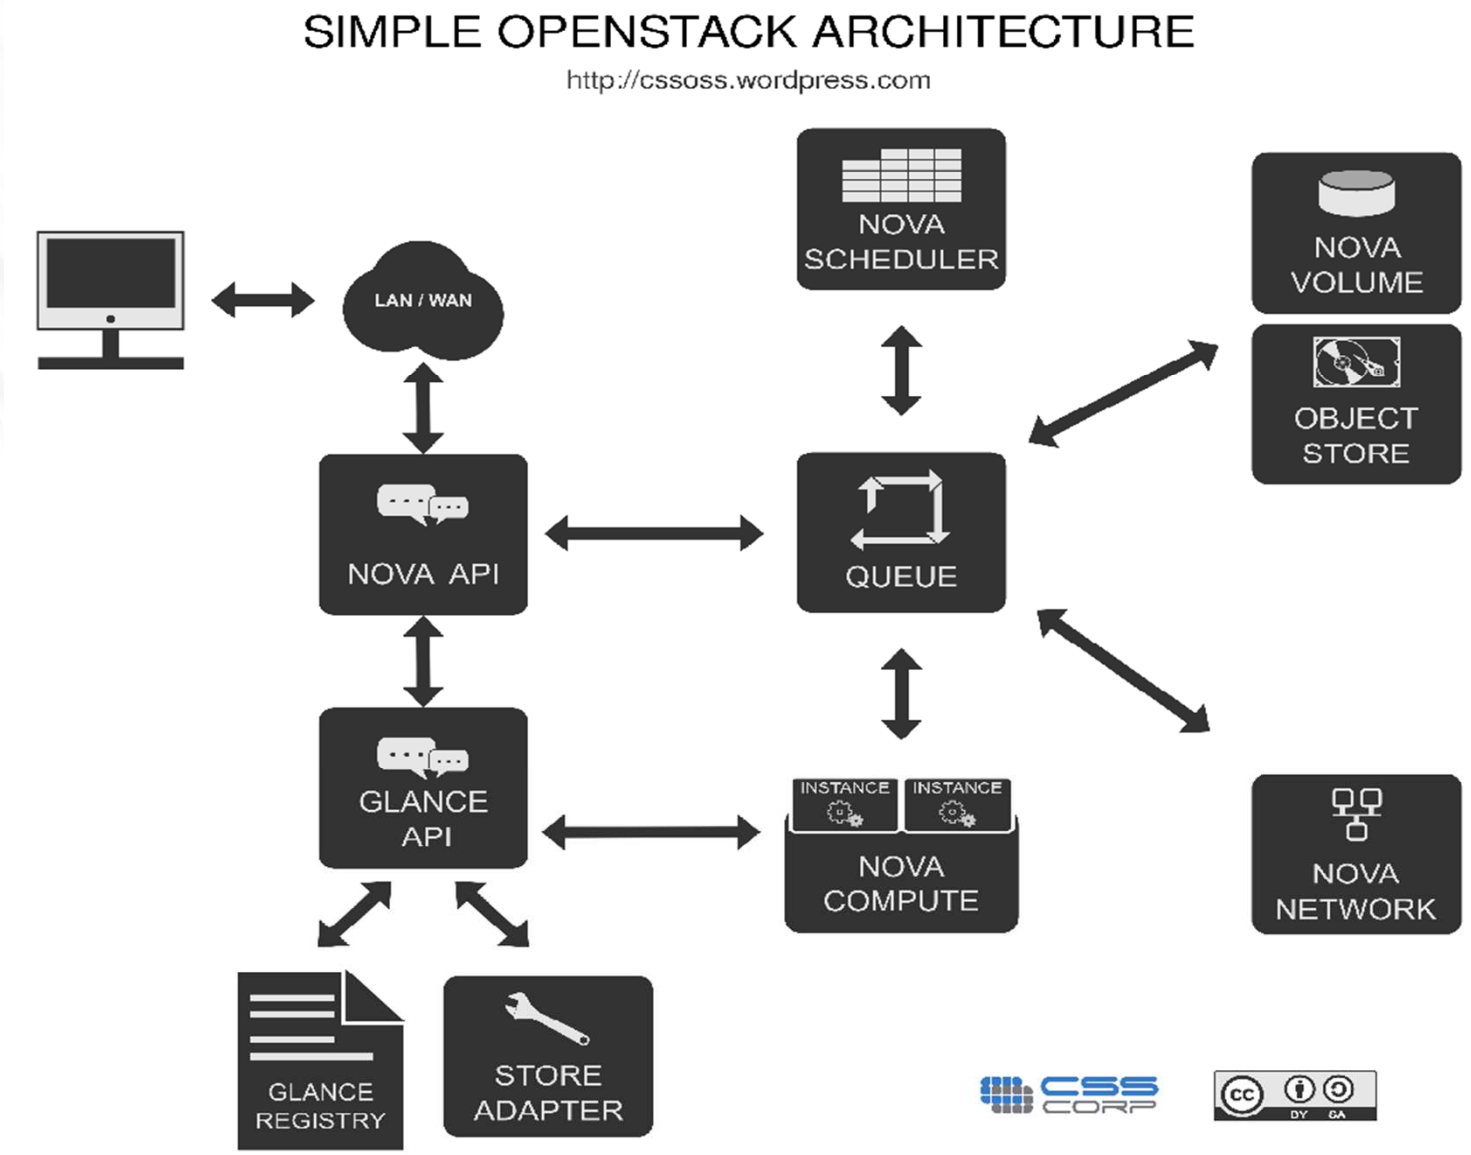

Database のtable一覧(nova)

agent_builds
aggregate_hosts
aggregate_metadata
aggregates
auth_tokens
block_device_mapping
bw_usage_cache
cells
certificates
compute_nodes
console_pools
consoles
dns_domains
fixed_ips
floating_ips
instance_actions
instance_faults
instance_info_caches
instance_metadata
instance_type_extra_specs
instance_types
instances
iscsi_targets
key_pairs
migrate_version
migrations
networks
projects
provider_fw_rules
quotas
s3_images
security_group_instance_association
security_group_rules
security_groups
services
sm_backend_config
sm_flavors
sm_volume
snapshots
user_project_association
user_project_role_association
user_role_association
users
virtual_interfaces
virtual_storage_arrays
volume_metadata
volume_type_extra_specs
volume_types
volumes

何が起きたか(4)

- VM起動コマンド

```
POST /v2/2f3e7954849c4208b63e50a75be85d42/servers HTTP/1.1
Host: 172.17.88.20:8774
Content-Length: 173
x-auth-project-id: 2f3e7954849c4208b63e50a75be85d42
accept-encoding: gzip, deflate
accept: application/json
x-auth-token: 7fa84cde93a54529a71a14a410bb34e2
user-agent: python-novaclient
content-type: application/json

{"server": {"name": "hoge2", "imageRef": "675e2b20-5754-4e6d-a8c0-a4fec4a1ce5a", "flavorRef": "1", "max_count": 1, "min_count": 1, "security_groups": [{"name": "default"}]}}
HTTP/1.1 202 Accepted
X-Compute-Request-Id: req-422b02dd-bf82-4ed0-a98e-53cceb9c1c2f
Location: http://172.17.88.20:8774/v2/2f3e7954849c4208b63e50a75be85d42/servers/18a84742-c222-4e9b-bfb4-ac946f4d06b3
Content-Type: application/json
Content-Length: 398
Date: Fri, 29 Jun 2012 08:43:11 GMT

{"server": {"OS-DCF:diskConfig": "MANUAL", "id": "18a84742-c222-4e9b-bfb4-ac946f4d06b3", "links": [{"href": "http://172.17.88.20:8774/v2/2f3e7954849c4208b63e50a75be85d42
/servers/18a84742-c222-4e9b-bfb4-ac946f4d06b3", "rel": "self"}, {"href": "http://172.17.88.20:8774/2f3e7954849c4208b63e50a75be85d42/servers/18a84742-c222-4e9b-bfb4-ac946
f4d06b3", "rel": "bookmark"}], "adminPass": "27FP9peAsxuS"}}
```

何が起きたか(5)

- ネットワークはnova-networkが担当
  - FLAGS.network_manager で実装が選べる

```
    cfg.StrOpt('network_manager',
               default='nova.network.manager.VlanManager',
               help='full class name for the Manager for network'),
```

```
class ComputeManager(manager.SchedulerDependentManager):
    """Manages the running instances from creation to destruction."""

    def __init__(self, compute_driver=None, *args, **kwargs):
        """Load configuration options and connect to the hypervisor."""
...
        self.network_manager = utils.import_object(FLAGS.network_manager)
```

文字列からクラスが初期化される

何が起きたか(5)

- ネットワークはnova-networkが担当
  - FLAGS.network_manager で実装が選べる

```
    cfg.StrOpt('network_manager',
               default='nova.network.manager.VlanManager',
               help='full class name for the Manager for network'),
```

```
class ComputeManager(manager.SchedulerDependentManager):
    """Manages the running instances from creation to destruction."""

    def __init__(self, compute_driver=None, *args, **kwargs):
        """Load configuration options and connect to the hypervisor."""
...
        self.network_manager = utils.import_object(FLAGS.network_manager)
```

文字列からクラスが初期化される

何が起きたか(6)

- devstack 起動時にアドレス領域の作成

```
# create a small network
$NOVA_DIR/bin/nova-manage network create private 10.0.0.0/24 1 256
```

- 仮想インターフェイスの作成
  - IP, MAC address の払い出し
  - bridgeに接続
  - DHCP でIPアドレスをVMに配布
  - iptables NATルールの作成

ネットワークについて(1)

- ネットワークは大変
  - 資源管理(IP, MAC address)
  - 仮想インターフェイスの作成削除
  - 仮想NICの配線
    - プライベートネットワーク、テナント間の分離
    - 複数台ホストがあるとき
  - ファイアウォール
  - QoS

ネットワークについて(2)

例1

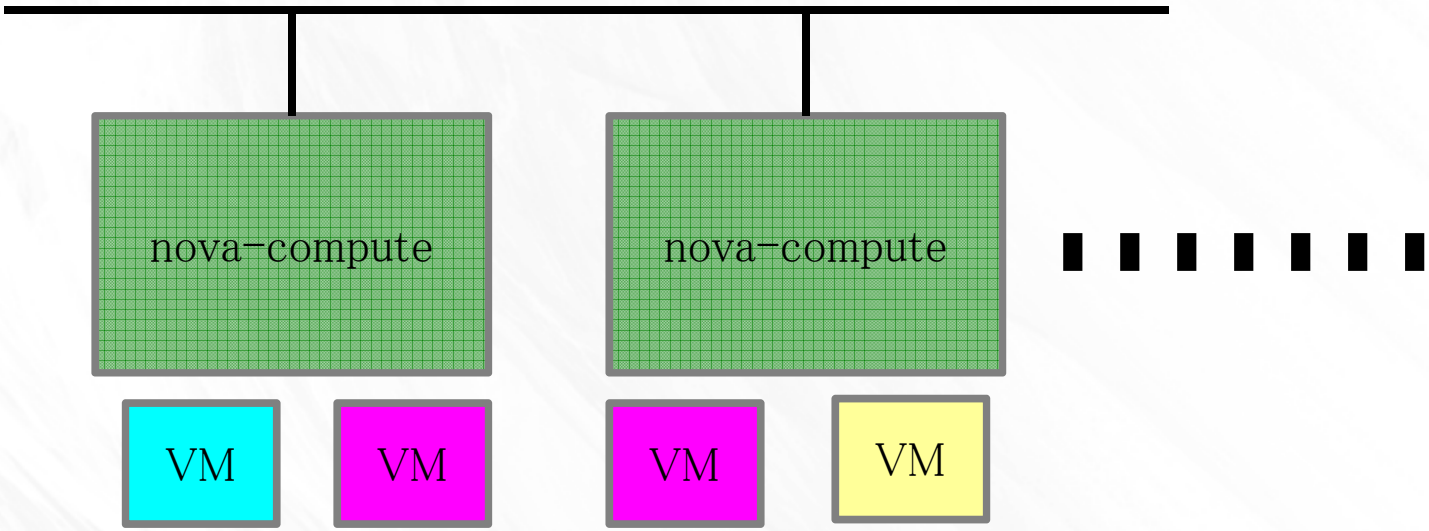

ネットワークについて(3)

例2

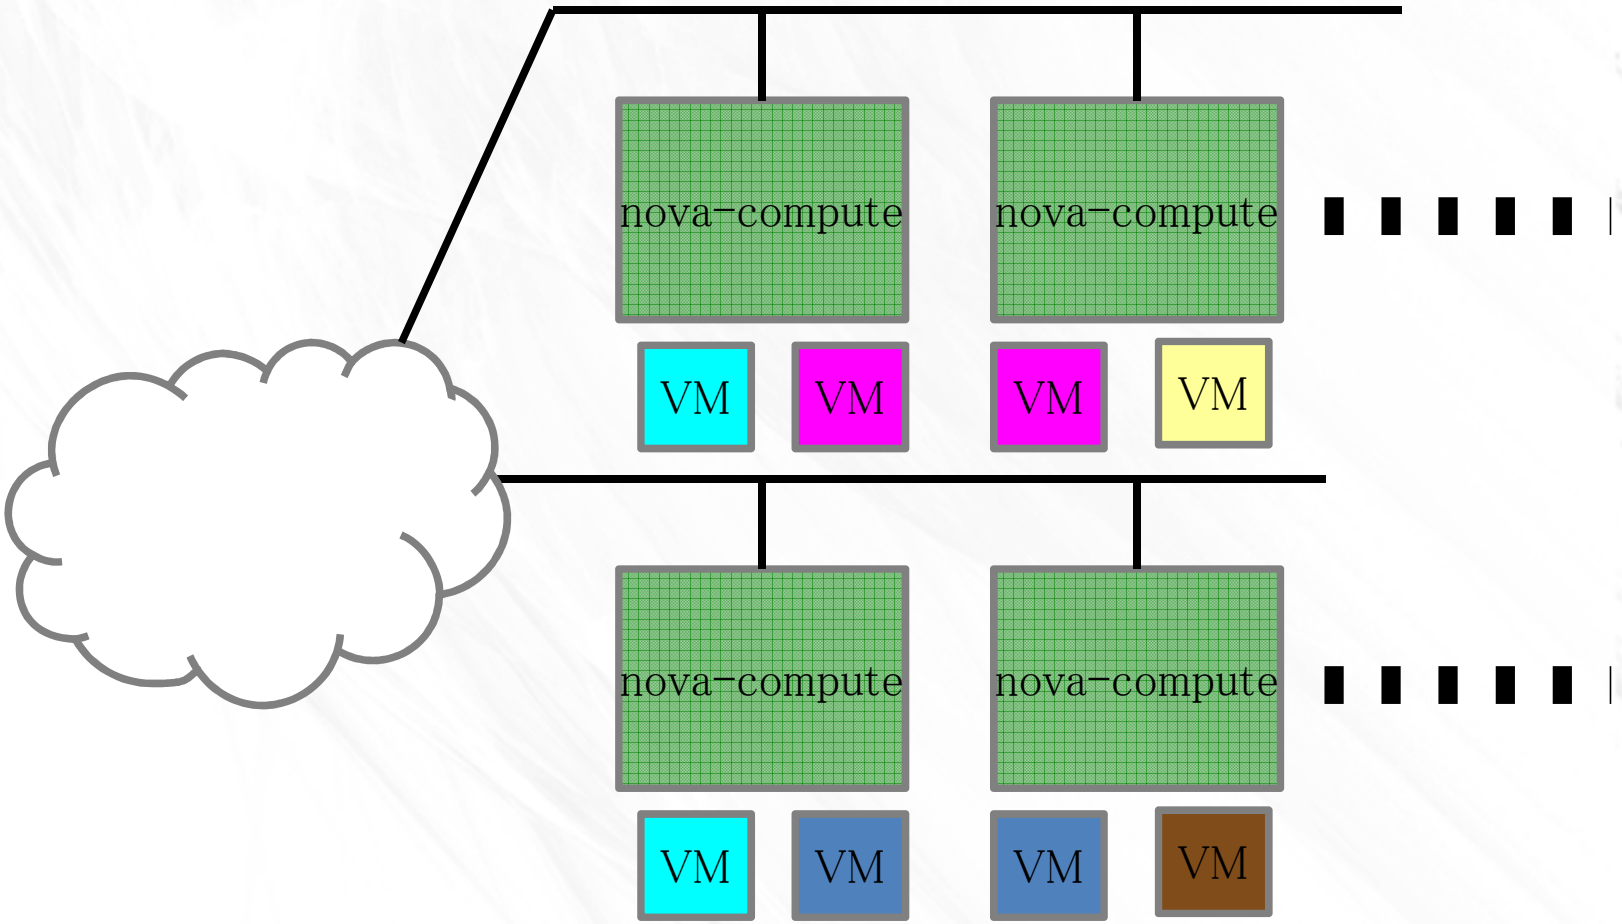

ネットワークについて(4)

- network_manager は以下の実装がある
  - FlatManager
  - FlatDHCPManager (devstack のdefault)
  - VlanManager (default)

    VLAN IDによるテナント間分離あり

    cloudpipe (openvpnでプライベートネットワークにアクセス可能)
  - QuantumManager

    novaと別プロジェクトなので分離するため

Quantum(1)

- プラグイン構成で拡張性を持たせた
  - Open vSwitch Plugin
  - Cisco UCS/Nexus Plugin
  - Linux Bridge Plugin
  - Nicira Network Virtualization Platform (NVP)Plugin
  - Ryu OpenFlow Controller Plugin
- 複数のプライベートネットワーク
- L2-over-L3 など

Quantum(2)

- devstackで openvswitch pluginが使える
  - Nova側にQuantumManager
  - 各 nova-compute ホストに quantum-agent
- 次期API
- http://wiki.openstack.org/QuantumV2APIIntro

Quantum in Essex

nova-network経由でQuantumが使われる

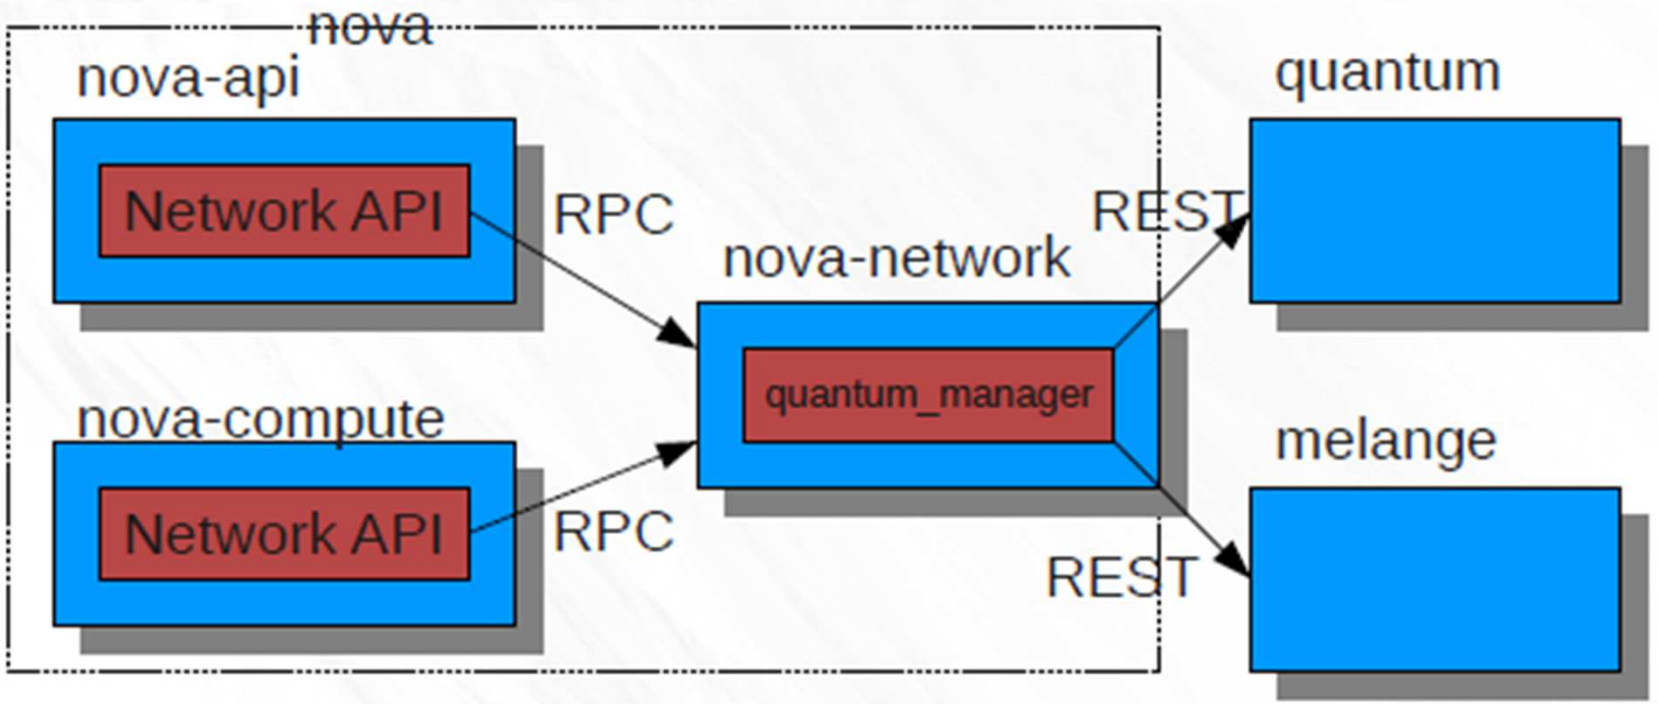

Quantum in Folsom

- nova-networkはなくなった
- Quantum APIはv2
- IPアドレス管理追加(melangeは不要)

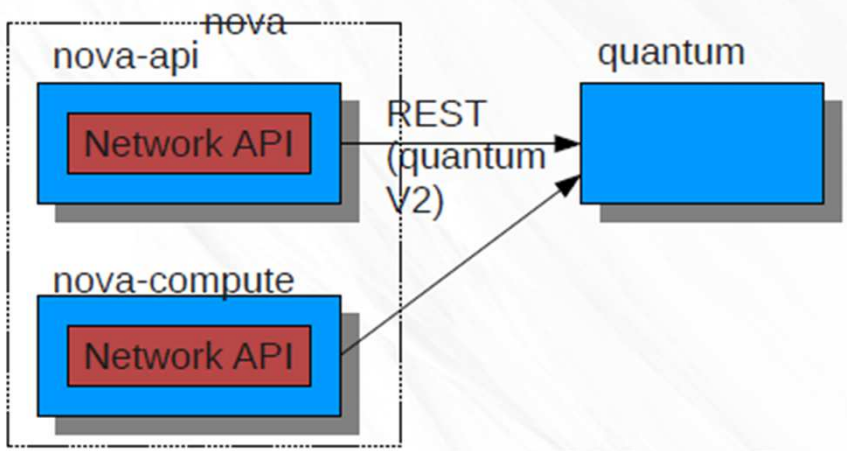

nova-volumes(1)

- 普通にVMを起動するとdiskは揮発性(VMを終了するとなくなる)
  - 実体は nova-compute の動くホストのlocal storage
- Nova-volumesを使うと永続的なdisk(volume)を作れる
  - volumeから起動もできる

nova-volumes(2)

- 実体は　nova-volumes　が動作するホスト上のLinux LVM上の論理ボリューム
  - iSCSI で nova-compute に見せている
  - ドライバを用意すれば iSCSI の箱も使える
- Snapshotとかとれる
- nova本体から分離することになった
  - http://wiki.openstack.org/Cinder

まとめ

- Openstackの開発手法(ツール)
  - 開発動向
- Openstack Computeの大まかな構成
- VM起動時の処理の流れ
- ネットワークの処理・課題
- Nova-volumes

コマンド一覧(1)

```
usage: nova [--debug] [--os_username OS_USERNAME] [--os_password OS_PASSWORD]
            [--os_tenant_name OS_TENANT_NAME] [--os_auth_url OS_AUTH_URL]
            [--os_region_name OS_REGION_NAME] [--service_type SERVICE_TYPE]
            [--service_name SERVICE_NAME] [--endpoint_type ENDPOINT_TYPE]
            [--os_compute_api_version OS_COMPUTE_API_VERSION]
            [--username USERNAME] [--region_name REGION_NAME]
            [--apikey APIKEY] [--projectid PROJECTID] [--url URL]
            <subcommand> ...

Command-line interface to the OpenStack Nova API.

Positional arguments:
  <subcommand>
    absolute-limits     Print a list of absolute limits for a user
    actions             Retrieve server actions.
    add-fixed-ip        Add new IP address to network.
    add-floating-ip     Add a floating IP address to a server.
    aggregate-add-host  Add the host to the specified aggregate.
    aggregate-create    Create a new aggregate with the specified details.
    aggregate-delete    Delete the aggregate by its id.
    aggregate-details   Show details of the specified aggregate.
    aggregate-list      Print a list of all aggregates.
    aggregate-remove-host
                        Remove the specified host from the specfied aggregate.
    aggregate-set-metadata
                        Update the metadata associated with the aggregate.
    aggregate-update    Update the aggregate's name and optionally
                        availability zone.
```

コマンド一覧(2)

```
boot                 Boot a new server.
cloudpipe-create     Create a cloudpipe instance for the given project
cloudpipe-list       Print a list of all cloudpipe instances.
console-log          Get console log output of a server.
credentials          Show user credentials returned from auth
delete               Immediately shut down and delete a server.
describe-resource    Show details about a resource
diagnostics          Retrieve server diagnostics.
dns-create           Create a DNS entry for domain, name and ip.
dns-create-private-domain
                     Create the specified DNS domain.
dns-create-public-domain
                     Create the specified DNS domain.
dns-delete           Delete the specified DNS entry.
dns-delete-domain    Delete the specified DNS domain.
dns-domains          Print a list of available dns domains.
dns-list             List current DNS entries for domain and ip or domain
                     and name.
endpoints            Discover endpoints that get returned from the
                     authenticate services
flavor-create        Create a new flavor
flavor-delete        Delete a specific flavor
flavor-list          Print a list of available 'flavors' (sizes of
                     servers).
floating-ip-create   Allocate a floating IP for the current tenant.
floating-ip-delete   De-allocate a floating IP.
floating-ip-list     List floating ips for this tenant.
floating-ip-pool-list
                     List all floating ip pools.
get-vnc-console      Get a vnc console to a server.
host-action          Perform a power action on a host.
```

コマンド一覧(3)

```
host-update          Update host settings.
image-create         Create a new image by taking a snapshot of a running
                     server.
image-delete         Delete an image.
image-list           Print a list of available images to boot from.
image-meta           Set or Delete metadata on an image.
image-show           Show details about the given image.
keypair-add          Create a new key pair for use with instances
keypair-delete       Delete keypair by its id
keypair-list         Print a list of keypairs for a user
list                 List active servers.
live-migration       Migrates a running instance to a new machine.
lock                 Lock a server.
meta                 Set or Delete metadata on a server.
migrate              Migrate a server.
pause                Pause a server.
quota-class-show     List the quotas for a quota class.
quota-class-update   Update the quotas for a quota class.
quota-defaults       List the default quotas for a tenant.
quota-show           List the quotas for a tenant.
quota-update         Update the quotas for a tenant.
rate-limits          Print a list of rate limits for a user
reboot               Reboot a server.
rebuild              Shutdown, re-image, and re-boot a server.
remove-fixed-ip      Remove an IP address from a server.
remove-floating-ip   Remove a floating IP address from a server.
rename               Rename a server.
rescue               Rescue a server.
resize               Resize a server.
resize-confirm       Confirm a previous resize.
resize-revert        Revert a previous resize (and return to the previous
                     VM).
```

コマンド一覧(4)

```
resume               Resume a server.
root-password        Change the root password for a server.
secgroup-add-group-rule
                     Add a source group rule to a security group.
secgroup-add-rule    Add a rule to a security group.
secgroup-create      Create a security group.
secgroup-delete      Delete a security group.
secgroup-delete-group-rule
                     Delete a source group rule from a security group.
secgroup-delete-rule
                     Delete a rule from a security group.
secgroup-list        List security groups for the current tenant.
secgroup-list-rules
                     List rules for a security group.
show                 Show details about the given server.
ssh                  SSH into a server.
suspend              Suspend a server.
unlock               Unlock a server.
unpause              Unpause a server.
unrescue             Unrescue a server.
usage-list           List usage data for all tenants
```

コマンド一覧(5)

```
volume-attach       Attach a volume to a server.
volume-create       Add a new volume.
volume-delete       Remove a volume.
volume-detach       Detach a volume from a server.
volume-list         List all the volumes.
volume-show         Show details about a volume.
volume-snapshot-create
                    Add a new snapshot.
volume-snapshot-delete
                    Remove a snapshot.
volume-snapshot-list
                    List all the snapshots.
volume-snapshot-show
                    Show details about a snapshot.
volume-type-create  Create a new volume type.
volume-type-delete  Delete a specific flavor
volume-type-list    Print a list of available 'volume types'.
x509-create-cert    Create x509 cert for a user in tenant
x509-get-root-cert  Fetches the x509 root cert.
bash-completion     Prints all of the commands and options to stdout so
                    that the
help                Display help about this program or one of its
                    subcommands.
```