Apache Libcloud

API Driven Operations

Paul Querna
September 29, 2010

About Me

- Chief Architect at Cloudkick
- Developer on Apache HTTP Server
- Former VP Infrastructure @ ASF
- Libcloud developer!

About the Cloud

About the Cloud

- Awesome.

About the Cloud

- Awesome.
- Really.

About the Cloud

- Awesome.
- Really.
- Maybe not always.

About the Cloud

- Awesome.
- Really.
- Maybe not always.
- But mostly.

Compute

- Amazon EC2
- Rackspace Cloud
- Linode
- GoGrid
- Voxel
- And many more!

I want a server.

I want a server right now.

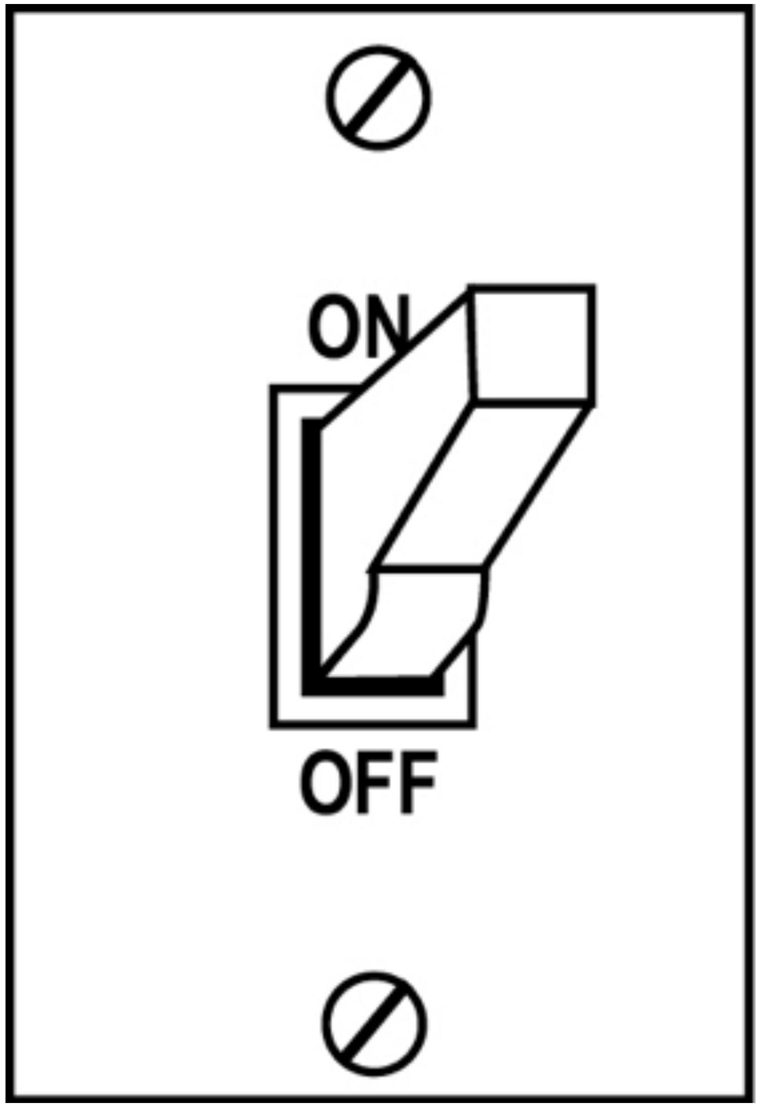

Enter Libcloud

```
from libcloud.types import Provider

from libcloud.providers import get_driver

rs = get_driver(Provider.RACKSPACE)('rackspace-apikey')

rs.create_node('serverA')
```

About Libcloud

- Started in the summer of 2009
- Easy to use.
- Portable.
- Pure Python
  - Java version also underway

Why?

- API Styles:
  - Amazon: XML + Custom HMAC Auth
  - Rackspace: JSON + Auth Tickets
  - SoftLayer: XML RPC + User / Password

16 Providers

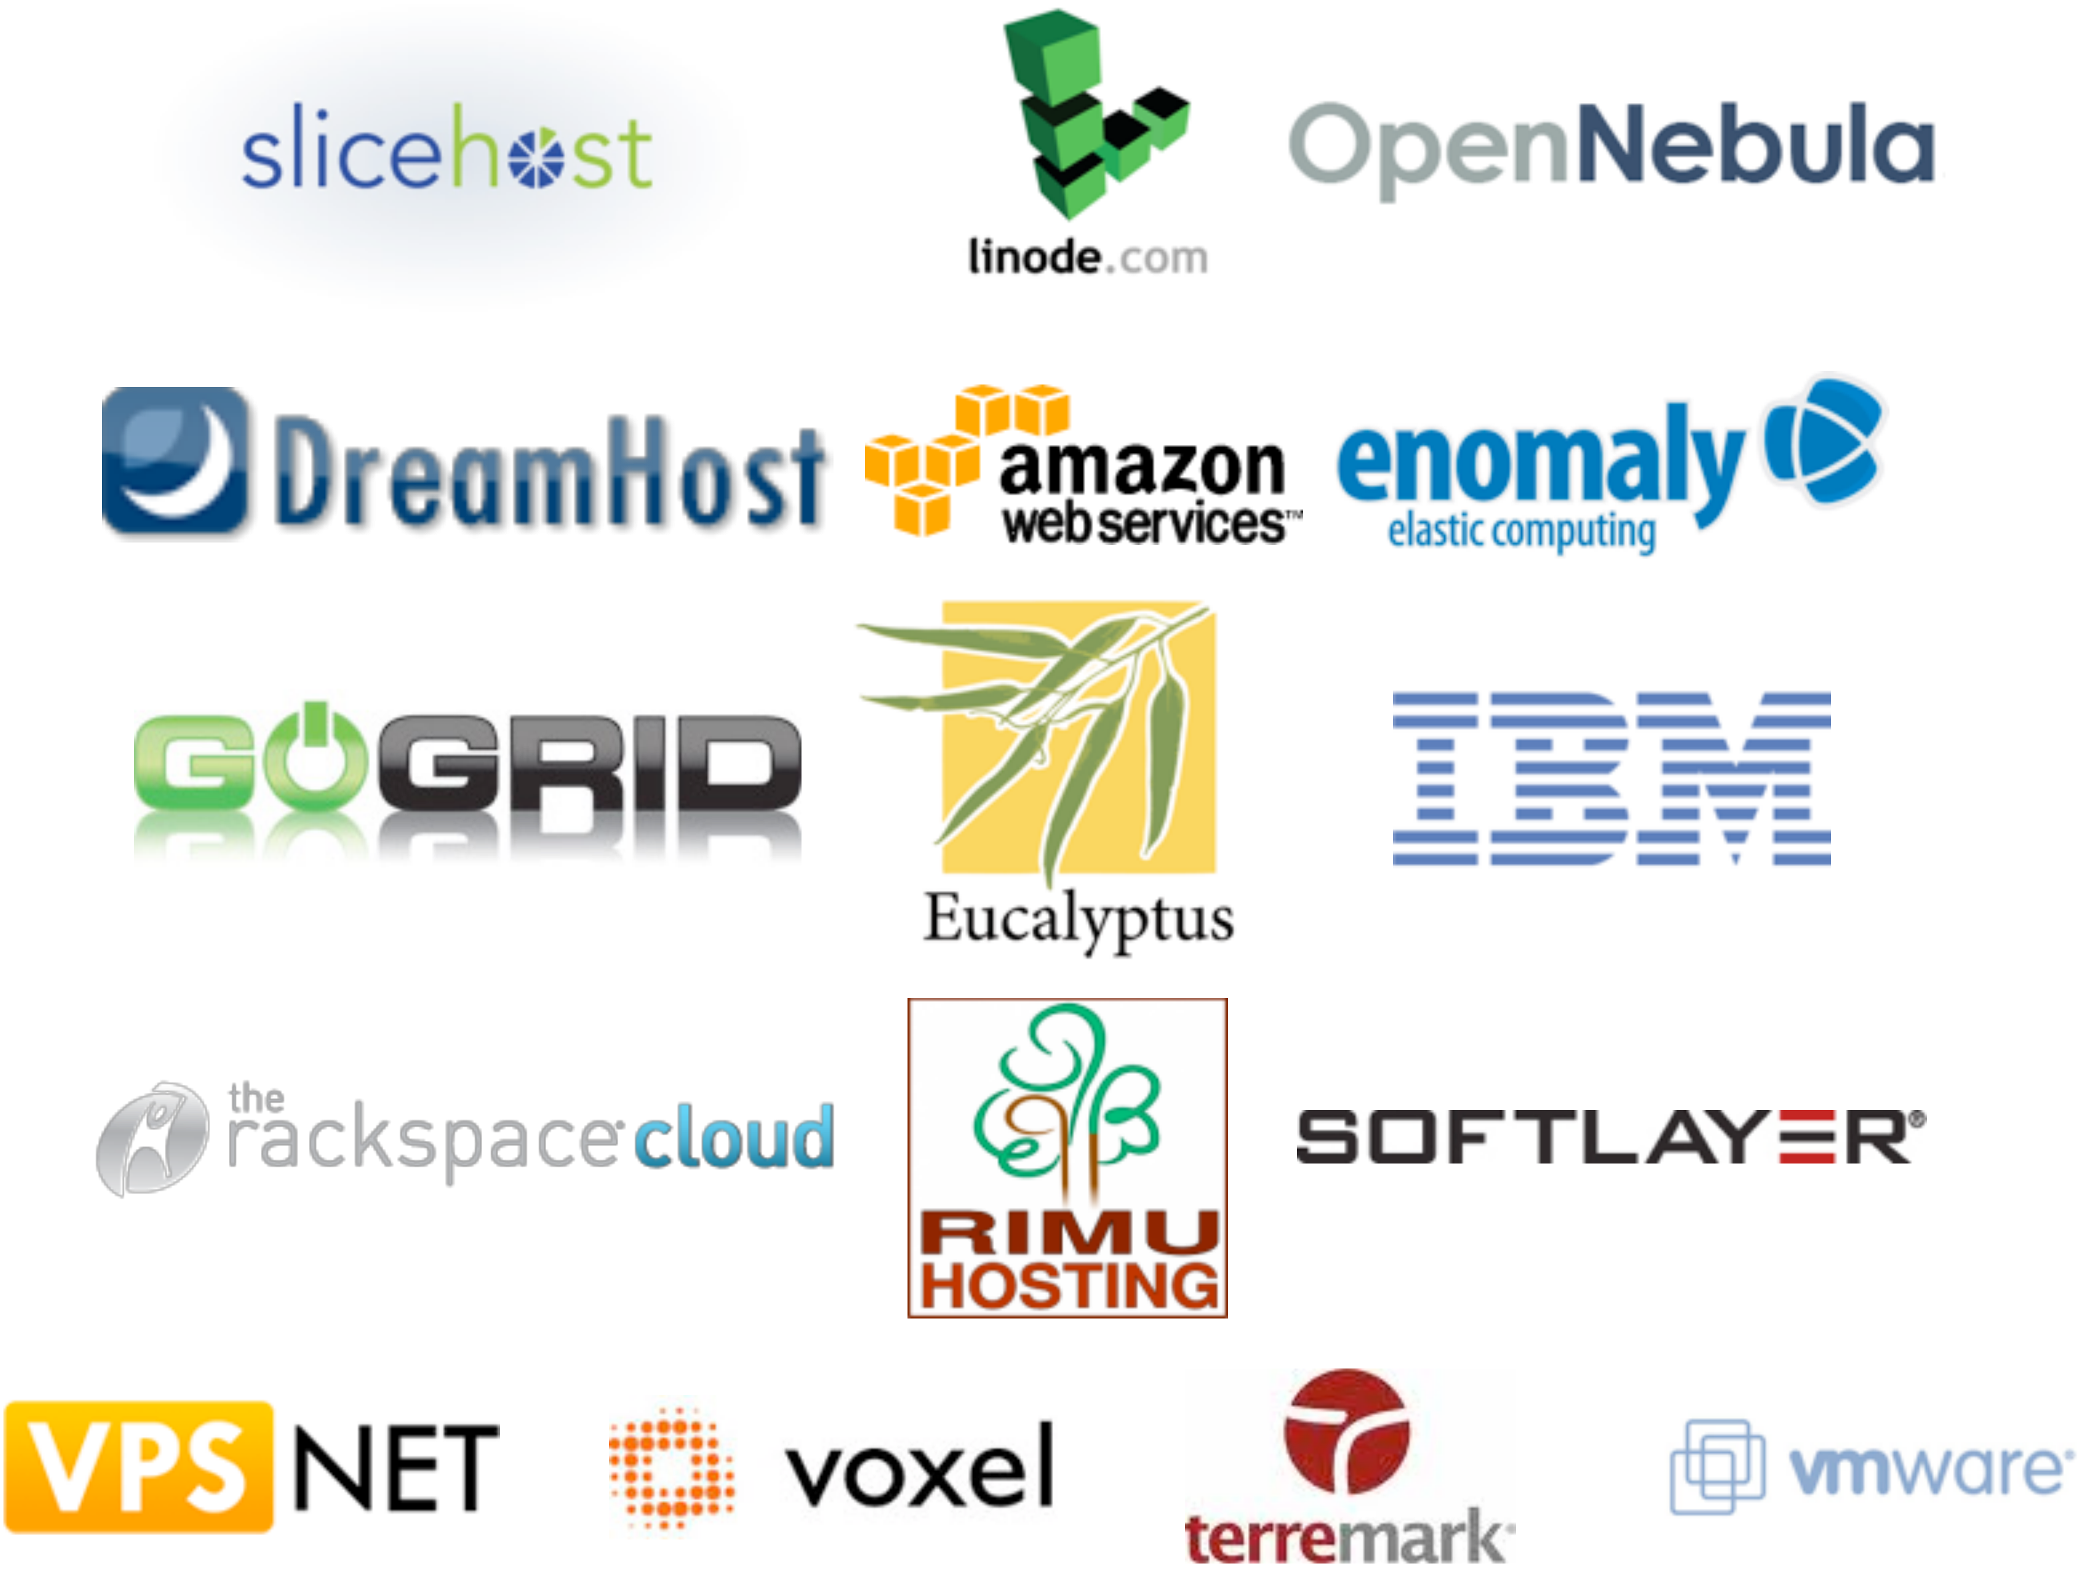

Libcloud APIs

- Originally 6 Core APIs
  - List Nodes
  - List Images
  - List Sizes
  - Create / Destroy / Reboot Node

list_nodes

```
foo = d.list_nodes()
for node in foo:
    print node.id
    print node.public_ip
```

Extended APIs

- Providers inconsistent about services.
- Have a “ex_” prefix, documented per-driver.
- Amazon Security Groups:
    - amz.create_node(‘foo’, ex_securitygroup=‘groupA’)

Getting Started

- easy_install apache-libcloud

Get your Provider Info

- Amazon:
  - http://aws-portal.amazon.com/gp/aws/developer/account/index.html?action=access-key
- Rackspace:
  - https://manage.rackspacecloud.com/APIAccess.do

List your Machines

```
from libcloud.types import Provider
from libcloud.providers import get_driver

d = get_driver(Provider.RACKSPACE)("xxxxxxx")

nodes = d.list_nodes()
for node in nodes:
  print "id: %s  name: %s  public_ips: %s" %
(node.id, node.name, node.public_ip)
```

Tricks!

- Included data:
  - Location
  - Price-per-hour in USD
- Cheapest 4 gigabyte node outside the US

Cheapest 4 gig node outside the US

```
possible = []
for d in drivers:
  loc = filter(lambda x: x.country != 'US',
            d.list_locations())
  for l in loc:
    sizes = filter(lambda x: x.ram >= 4096, d.list_sizes(l))
    for s in sizes:
      possible.append({'size': s,
                    'location': l,
                    'driver': d})

best = sorted(possible, lambda x,y:
              x['size'].price  < y['size'].price)[0]
print best
```

One more thing!

- deploy_node
  - Calls create_node
  - Consistent initial bootstrapping of machines.
    - SSH Keys
    - Configuration Management

Installing Puppet

```
skey = SSHKeyDeployment(key)
sd = ScriptDeployment("apt-get
install -y puppet")
msd = MultiStepDeployment([skey, sd])
node = d.deploy_node(name="lc-test",
deploy=msd)
```

Running commands

- Fabric
  - http://www.fabfile.org/

Integrating with Fabric

```
env.hosts = [x.public_ip[0] for x in
d.list_nodes()]

def hostname():
  run('hostname')
```

```
$ fab hostname
[173.45.245.33] run: hostname
[173.45.245.33] out:  lctest3.k1k.me
[173.45.245.32] run: hostname
[173.45.245.32] out:  lctest2.k1k.me

Done.
Disconnecting from 173.45.245.33... done.
Disconnecting from 173.45.245.32... done.
```

Up next for Libcloud?

Higher Level Apps

- Silver Lining: Python Deployment on the Cloud
    - http://cloudsilverlining.org/
- Pantheon Mercury: Drupal Deployment on the Cloud
    - http://getpantheon.com/

Applications; Not just VMs

- Cast:
  - http://github.com/cloudkick/cast
- Galaxy:
  - http://github.com/ning/galaxy

Related Projects

- JClouds
  - Java
- Apache Deltacloud
  - Ruby, little different ideas

Questions?

- Apache Libcloud:

  http://libcloud.org/

  #libcloud on Freenode IRC

- Slides online:

  http://paul.querna.org/slides