Python at

cloudkick

Tomaz Muraus
tomaz.muraus@rackspace.com
June 22, 2011

EUROPYTHON 2011
FLORENCE, JUNE 20–26

rackspace

Agenda

- Who am I
- What is Cloudkick
- Cloudkick Architecture
- Why Python?
- Python at Cloudkick

Agenda

- Web Application & Django
- Backend / Network Services & Twisted
- Service Communication & RPC
- Testing
- Other places
- Questions

Who am I

- Tomaz Muraus, @KamiSLO
- Software developer at Cloudkick / Rackspace
- Author of multiple Python libraries and Django apps – http://github.com/Kami
- FOSS supporter & lover
- Apache Libcloud Committer

What is Cloudkick

- Server management and monitoring SaaS
- Manage all your cloud and physical servers from a single control panel
- Set up monitors and alerts – only get waken up when the stuff actually breaks
- Annotate graphs
- and more...

What is Cloudkick

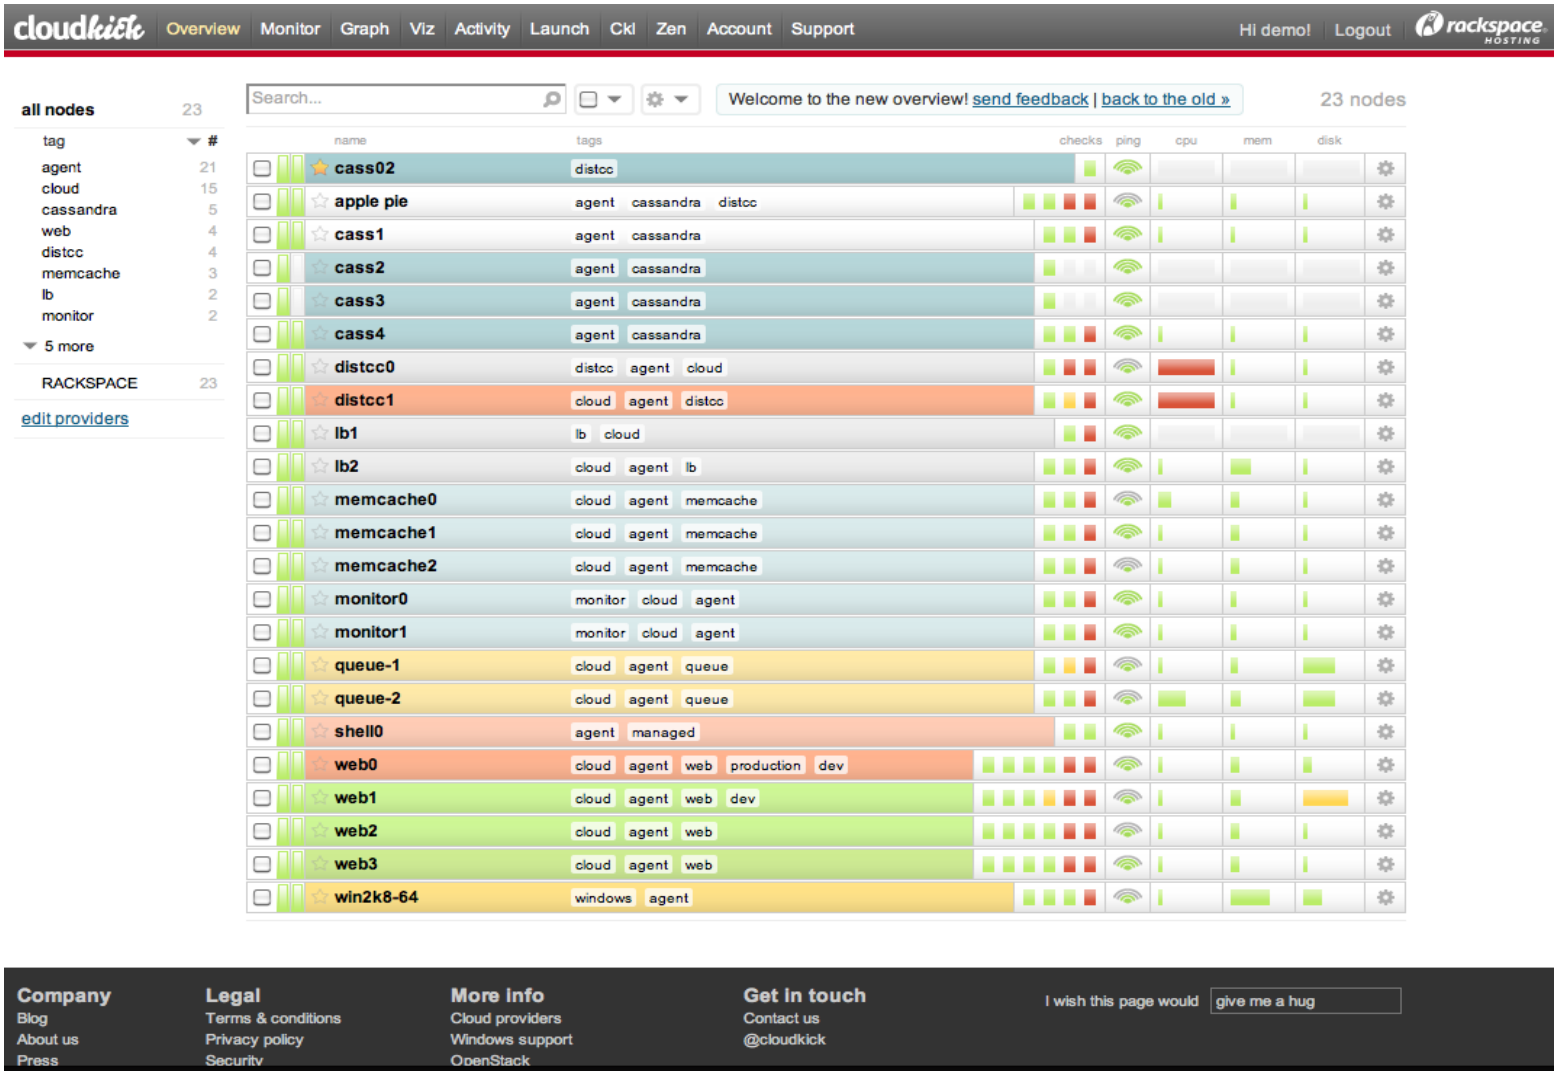

EUROPYTHON 2011
FLORENCE, JUNE 20–26

rackspace

What is Cloudkick

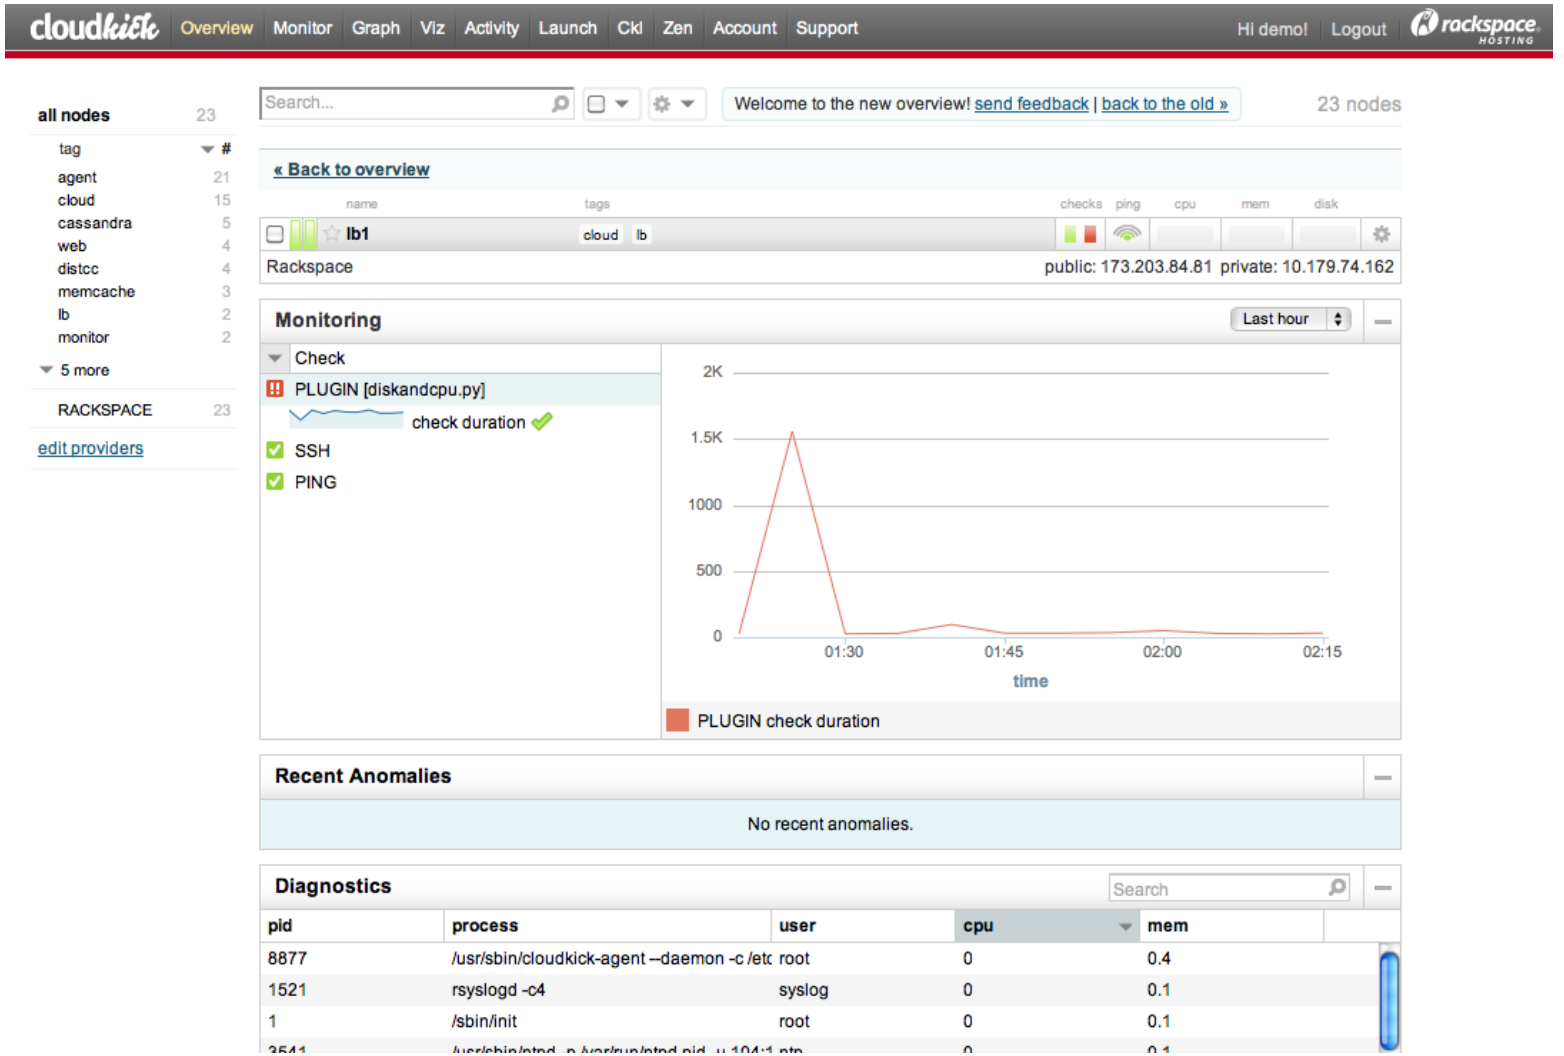

What is Cloudkick

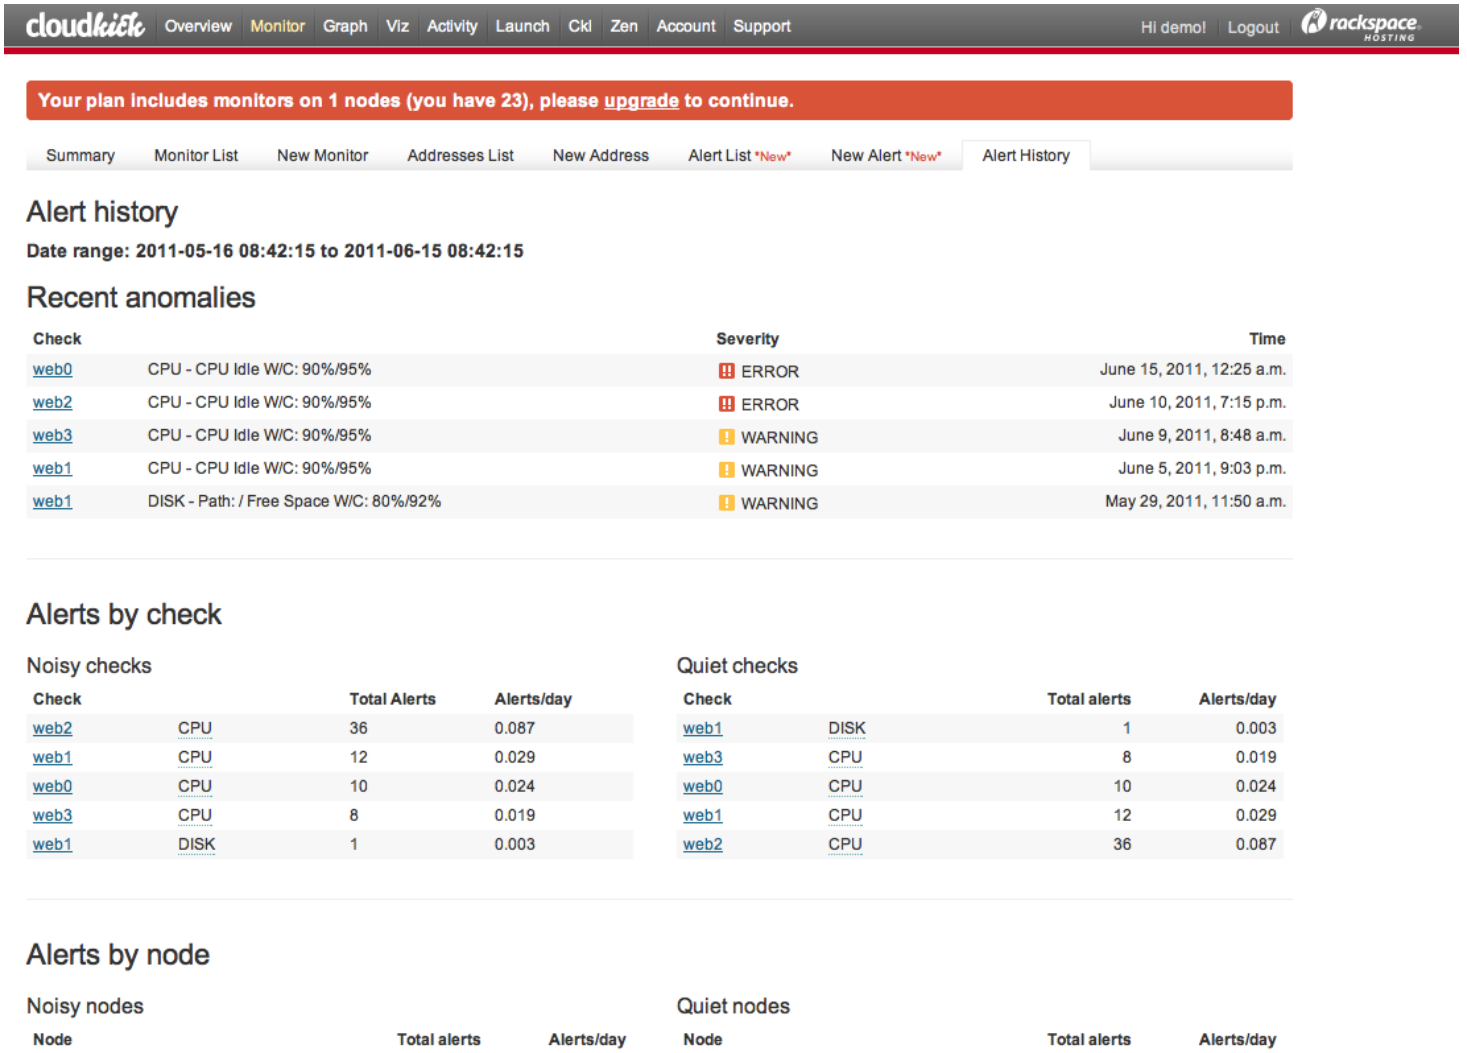

EUROPYTHON 2011
FLORENCE, JUNE 20–26

rackspace

What is Cloudkick

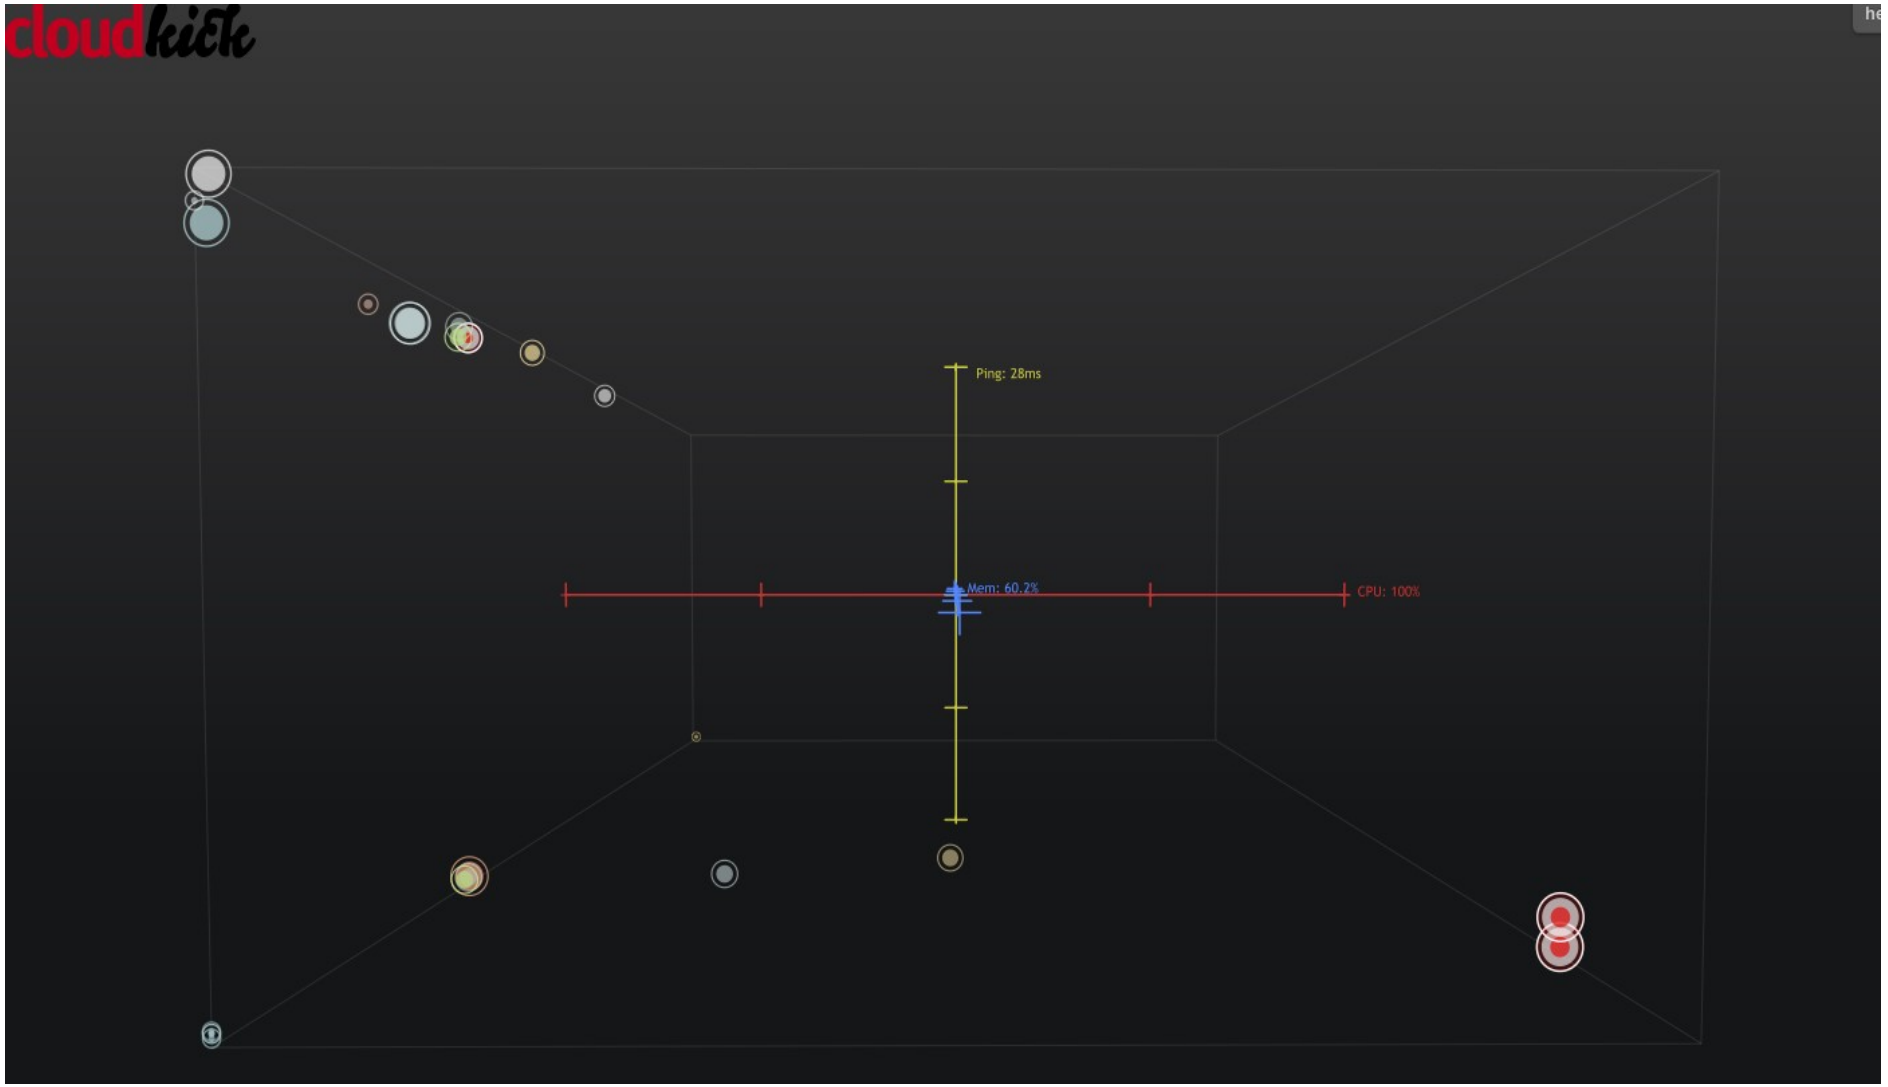

Cloudkick Architecture

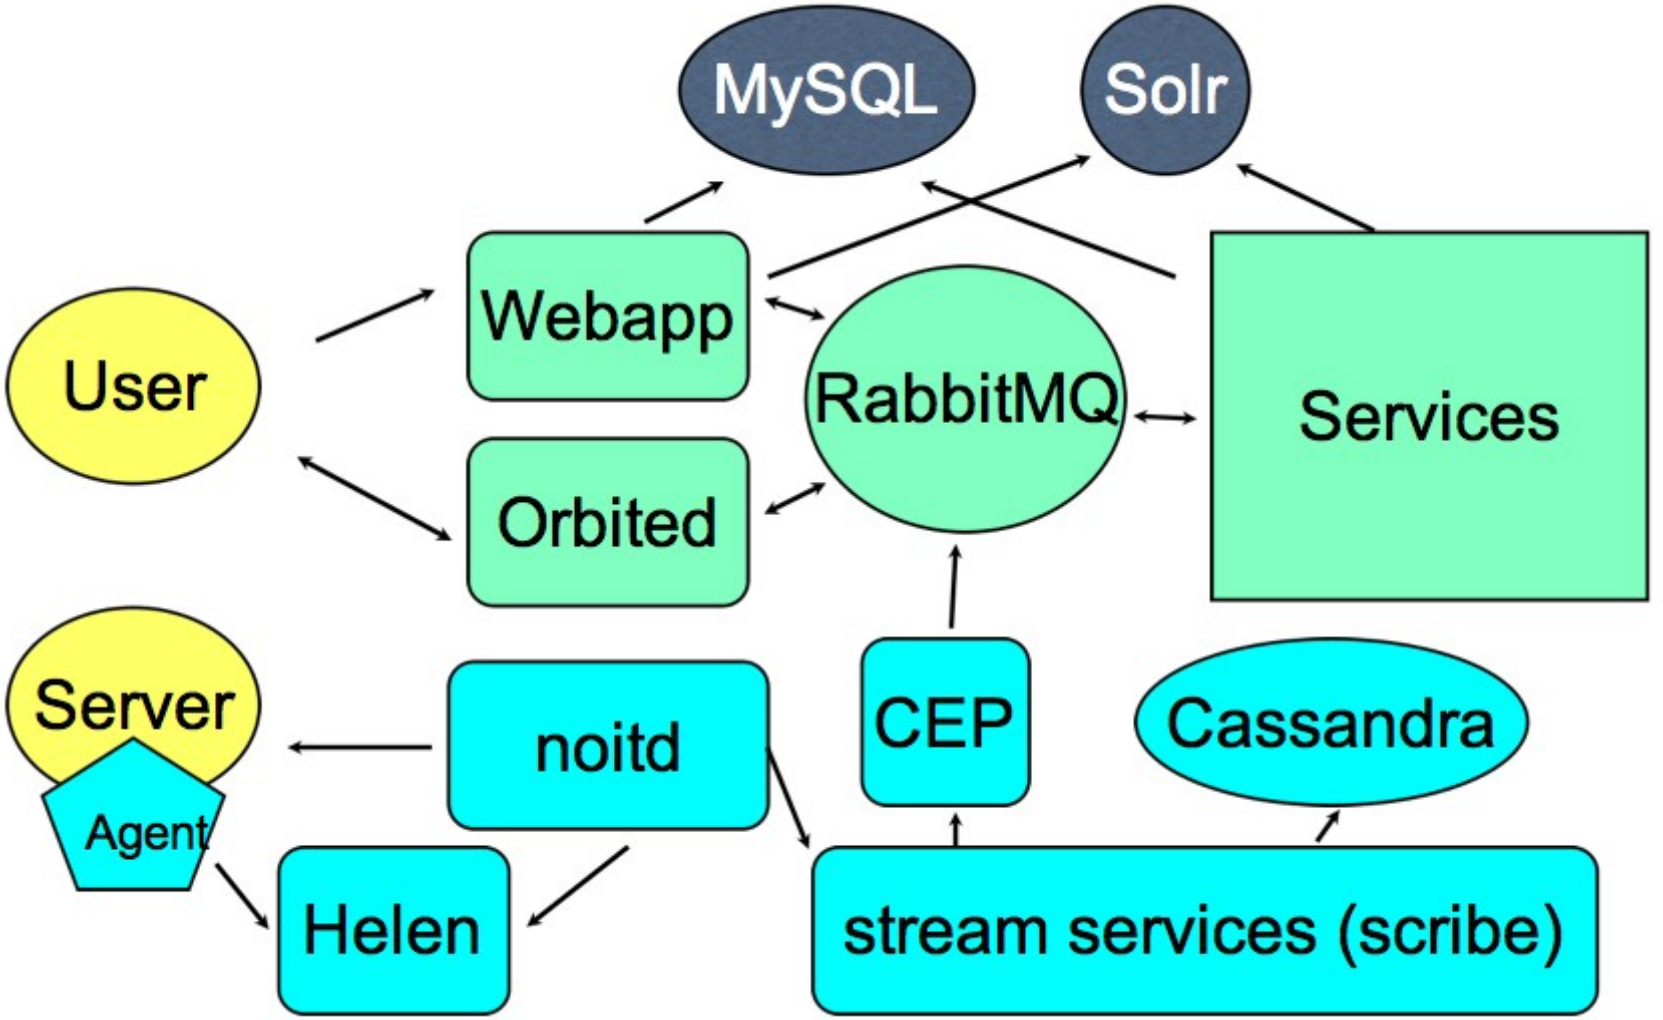

Why Python?

- Expressive
- Developers love it
- Quality library for just about anything you can think of
- Large and active community
- Whitespace matters!

Python at Cloudkick

- Used in many different places
- Web application
- Backend / Network services
- Random scripties

Web application & Django

- Web framework - Django
- Migrations - south
- API – piston
- Cassandra – Custom library & ORM
- Solr - solrpy

Web application & Django

- Profiling – ProfilingMiddleware, python-profiler
- Exception logging - django-sentry

Web application & Django

```
         100011 function calls (92344 primitive calls) in 1.997 CPU seconds

   Ordered by: internal time, call count

    ncalls  tottime  percall  cumtime  percall filename:lineno(function)
 373/332    0.132    0.000    0.615    0.002 /data/virtualenv/lib/python2.6/site-packages/django/utils/im
 347/229    0.087    0.000    0.256    0.001 /data/virtualenv/lib/python2.6/sre_parse.py:385(_parse)
    5187    0.062    0.000    0.099    0.000 /data/virtualenv/lib/python2.6/sre_parse.py:207(get)
2863/336    0.058    0.000    0.154    0.000 /usr/lib/python2.6/copy.py:144(deepcopy)
     379    0.056    0.000    0.122    0.000 /data/virtualenv/lib/python2.6/site-packages/django/utils/re
    7551    0.048    0.000    0.048    0.000 /data/virtualenv/lib/python2.6/site-packages/django/utils/re
    5884    0.045    0.000    0.045    0.000 /data/virtualenv/lib/python2.6/sre_parse.py:188(__next)
   234/9    0.041    0.000    0.334    0.037 /data/virtualenv/lib/python2.6/site-packages/django/template
    29/1    0.038    0.001    0.837    0.837 /data/virtualenv/lib/python2.6/site-packages/django/core/url
      22    0.029    0.001    0.038    0.002 /data/virtualenv/lib/python2.6/site-packages/django/template
    1140    0.026    0.000    0.070    0.000 /data/virtualenv/lib/python2.6/site-packages/django/utils/da


DB connection to default
Top used SQL queries:
2 x SELECT `inventory_nodeagent`.`id`, `inventory_nodeagent`.`node_id`, `inventory_nodeagent`.`state`, `in
2 x SELECT `inventory_generationnodeagent`.`id`, `inventory_generationnodeagent`.`node_id`, `inventory_gen
2 x SELECT `auth_user`.`id`, `auth_user`.`username`, `auth_user`.`first_name`, `auth_user`.`last_name`, `a
2 x SELECT `account_accountuser`.`id`, `account_accountuser`.`account_id`, `account_accountuser`.`user_id`
1 x SELECT `tagging_taggednode`.`id`, `tagging_taggednode`.`node_id`, `tagging_taggednode`.`tag_id`, `tagg
Top SQL queries by time:
[{'sql': u'SELECT `account_accountuser`.`id`, `account_accountuser`.`account_id`, `account_accountuser`.`u
  'time': '0.012'},
 {'sql': u'SELECT `monitoring_check`.`id`, `monitoring_check`.`created_at`, `monitoring_check`.`updated_at
  'time': '0.005'},
 {'sql': u'SELECT `inventory_node`.`id`, `inventory_node`.`created_at`, `inventory_node`.`updated_at`, `in
  'time': '0.003'},
```

EUROPYTHON 2011
FLORENCE, JUNE 20–26

rackspace

Backend Services & Twisted

- Twisted is used heavily
- Old, battle-tested, (mostly) works
- Supports almost any protocol you can think of
- Writing good Twisted code is not that easy

Backend Services & Twisted

- Uncaught exceptions = memory leaks
- Be careful with threads.deferToThread
- emailLogObserver for sending tracebacks to email

Backend Services & Twisted

- Uncaught exceptions = memory leaks
- Be careful with threads.deferToThread
- emailLogObserver for sending tracebacks to email

Backend Services & Twisted

- We use a simple stats library for recording service-level metrics
- Metrics are exposed over HTTP in a JSON format
- Similar to https://github.com/codahale/metrics

Backend Services & Twisted

- counter.add(key)
- counter.inc_ops(key)
- counter.dec_ops(key)
- counter.add_avg(key, value)
- counter.bind(key, type, func, *args, **kwargs)
- @count_calls

Backend Services & Twisted

```
counter = Counter()

def my_function2(*args, **kwargs):
  try:
    # Increase operation count
    counter.inc_ops('execute_some_op')
    some_op()
  except Exception, e:
    # Increase the counter
    counter.add('some_op_failed')
  finally:
    # Signalize that the function has finished
    # pending -= 1
    counter.dec_ops('execute_come_op')
```

Backend Services & Twisted

```
counter = Counter()

def my_function3(*args, **kwargs):
  start = time.time()
  my_function()
  # Measure how long function execution took
  counter.add_avg('my_function_execute_ms',
                  (time.time() - start))
```

Backend Services & Twisted

```
{'metrics': [{'name': 'execute_some_op_pending',
              'type': 'int', 'value': 3},
             {'name': 'execute_some_op_total',
              'type': 'gauge', 'value': 4},
             {'name': 'my_function_execute_msg_avg',
              'type': 'float',
              'value': 22.666},
             {'name': 'my_function_execute_msg_max',
              'type': 'float',
              'value': 100},
             {'name': 'my_function_execute_msg_min',
              'type': 'float',
              'value': 4.2},
            {'name': 'some_op_failed', 'type': 'float', 'value': 3},
            {'name': 'uptime', 'type': 'float', 'value': 123.08233}],
 'state': 'ok',
 'status': 'service is good'}
```

Service Communication and RPC

- Event framework
- Scribe
- Thrift

Service Communication and RPC

- Event framework
- Multiple policies:
  - Fan out
  - Round Robin
  - Scoped Round Robin
- Used in Django and Twisted land

Service Communication and RPC

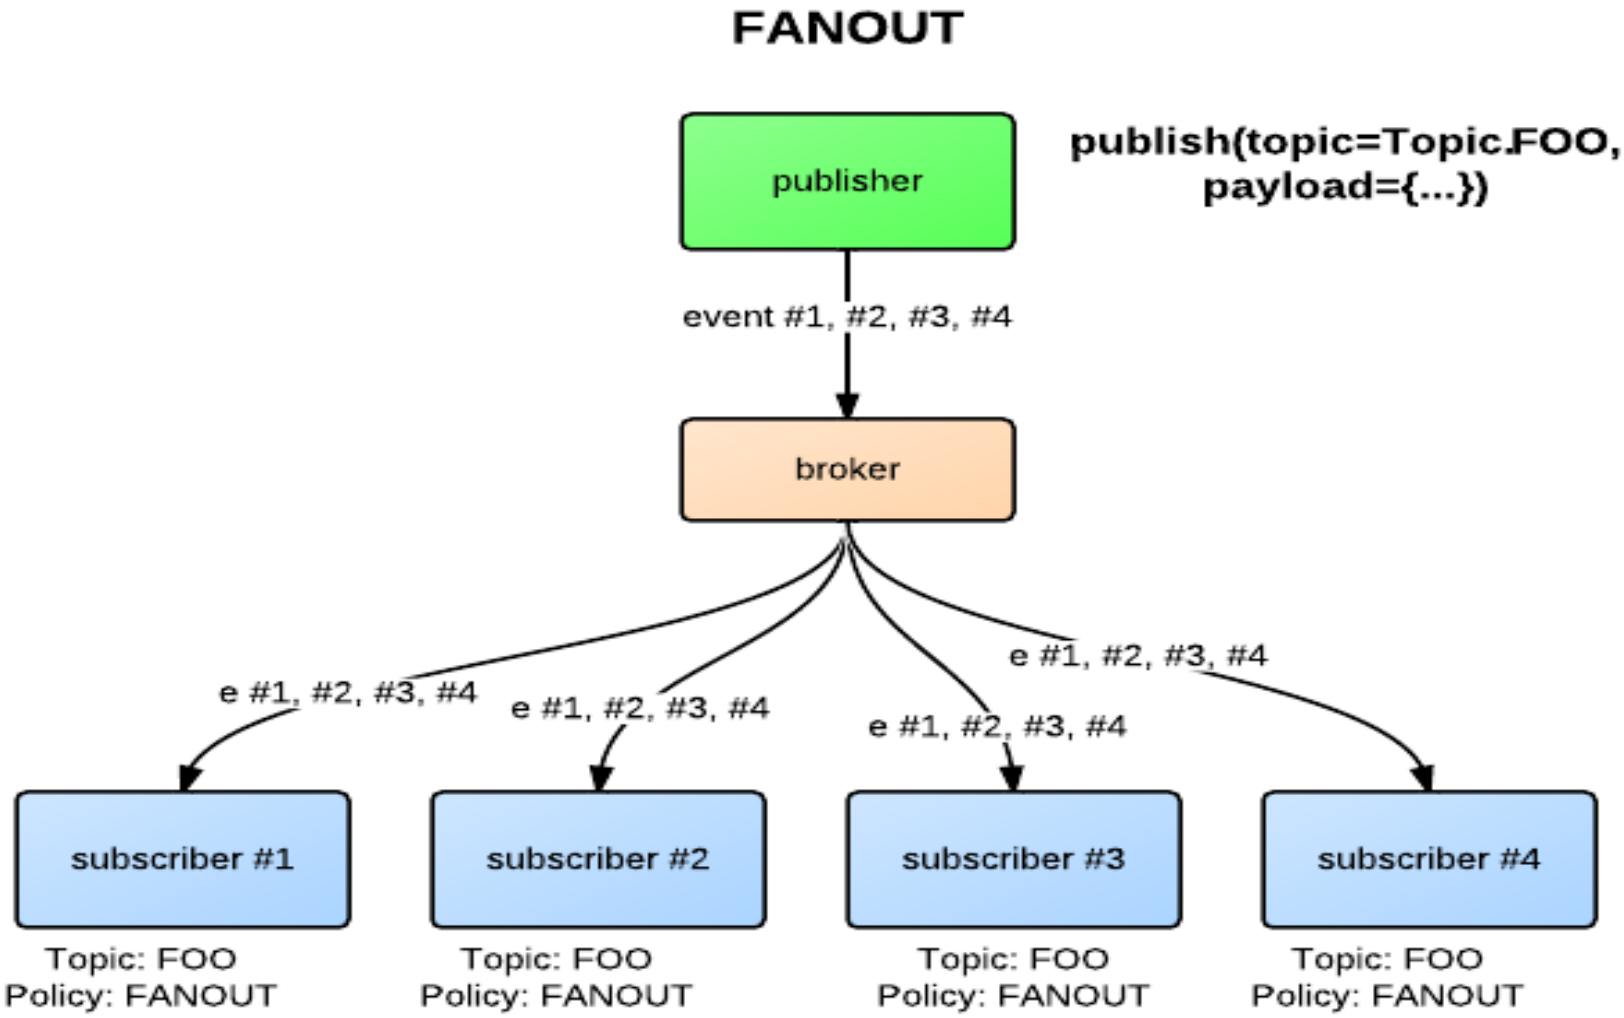

Service Communication and RPC

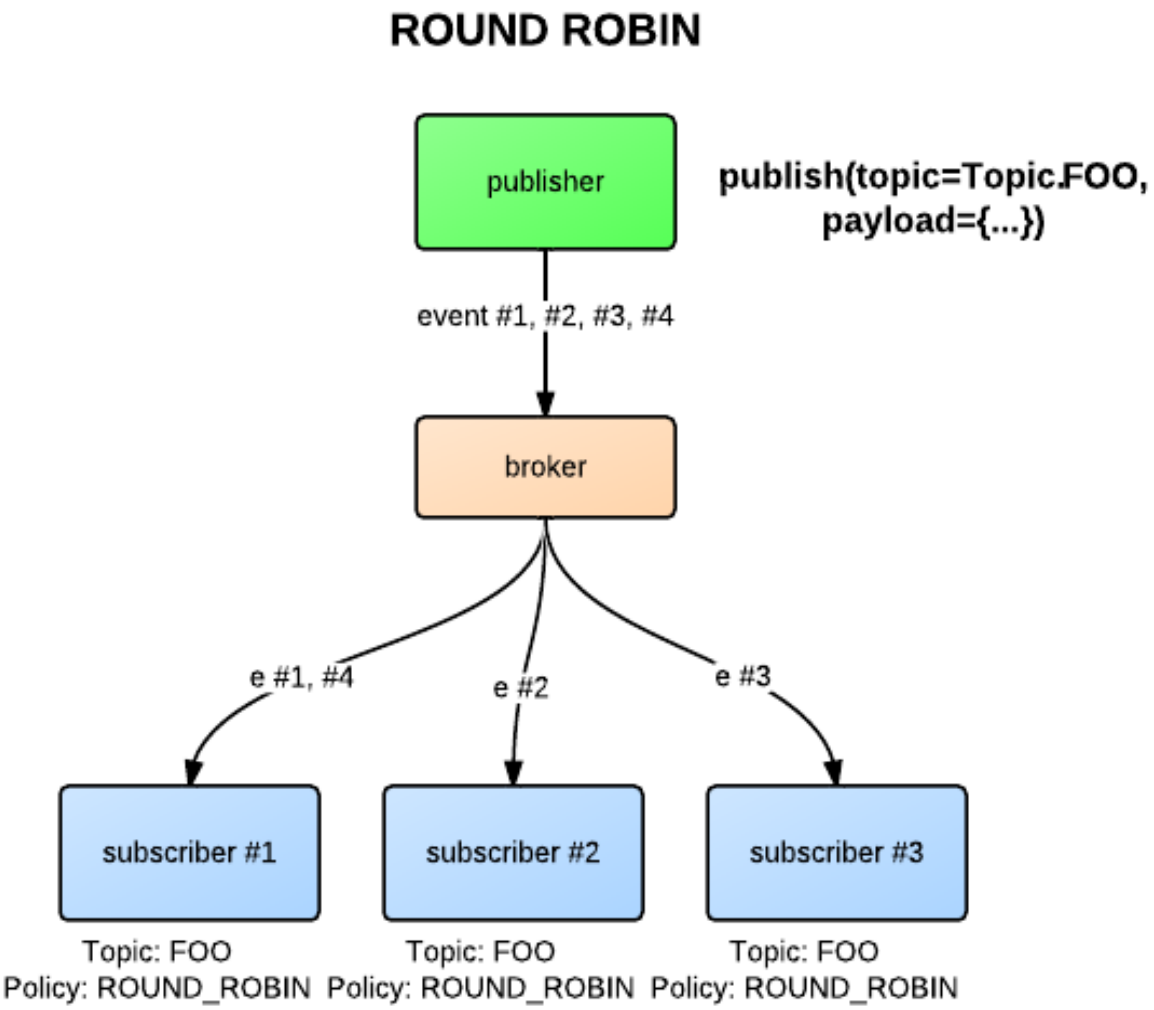

Service Communication and RPC

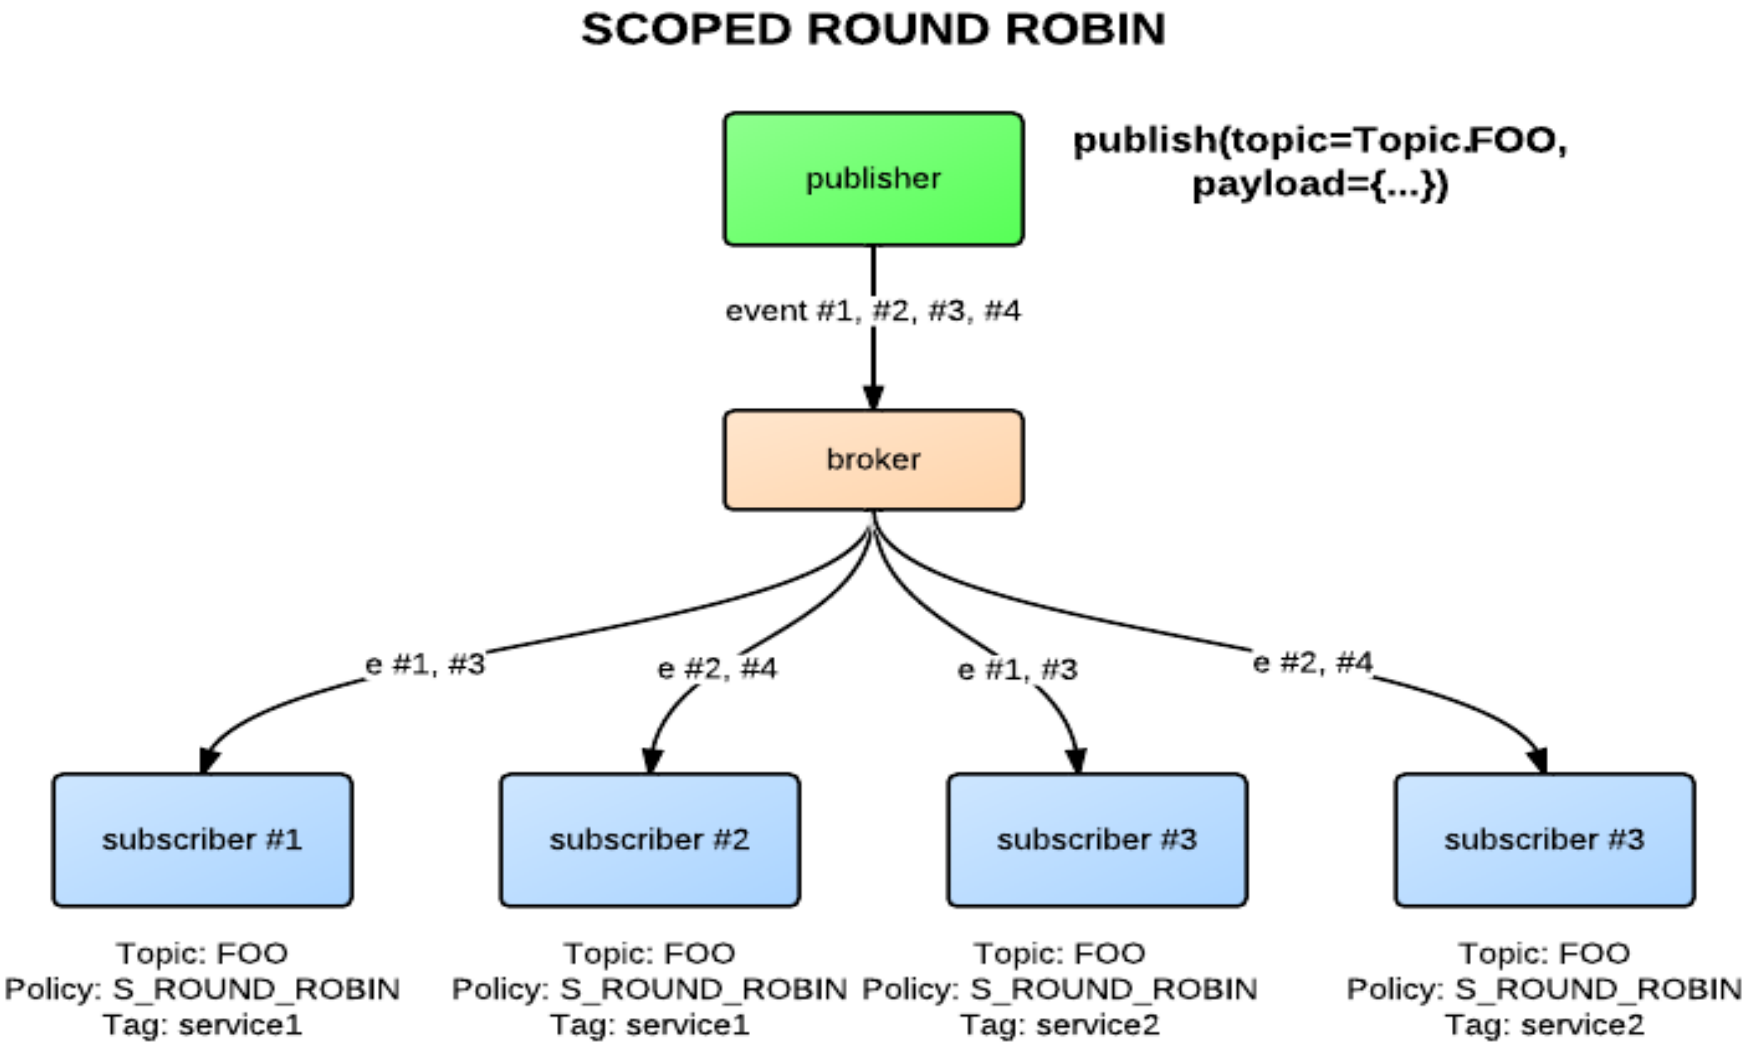

Service Communication and RPC

```
# Subscribe example

from cloudkick.events.async import dispatcher

def readyhook():
  dispatcher.register(Topic.TEST1, Policy.FANOUT,
                      handle_event1, topic)
  dispatcher.register(Topic.TEST2, Policy.ROUND_ROBIN,
                      handle_event2, topic)

dispatcher.ready_hook(readyhook)
```

Service Communication and RPC

```
# Publish example

from cloudkick.events.async import dispatcher
from cloudkick.events.path_defines import Topics

dispatcher.start()
event = event_ttypes.NodeNameUpdate(node=_id=1234,
                        old_name='foo', new_name='bar')
dispatcher.publish(Topics.NODE_UPDATE, event,
                       account_id=1, user_id=2)
```

Testing

- Not that fun
- Someone has to do it
- Functional tests – Django test framework & unittest
- Integration tests – Twisted test framework (trial)

Testing

- Custom Django and Twisted parallel test runner
- Up to 50% faster

Other places

- Deployment – fabric
- Continuous Integration – buildbot
- ...
- Snakes are everywhere!11

Questions?

Questions?

- Thanks

- P.S. We are looking for Python devs - http://rackertalent.com/san-francisco/