Advanced Aspects of the Django Ecosystem

Haystack, Celery & Fabric

Simon Willison, @simonw

EuroPython, 21st June 2011

http://lanyrd.com/sftzq

Secret Weapons

- **Haystack**
  - Top quality full-text search in seconds
- **Celery**
  - Offline processing in an instant
- **Fabric**
  - Automated deployments for lazy people

Lanyrd.com

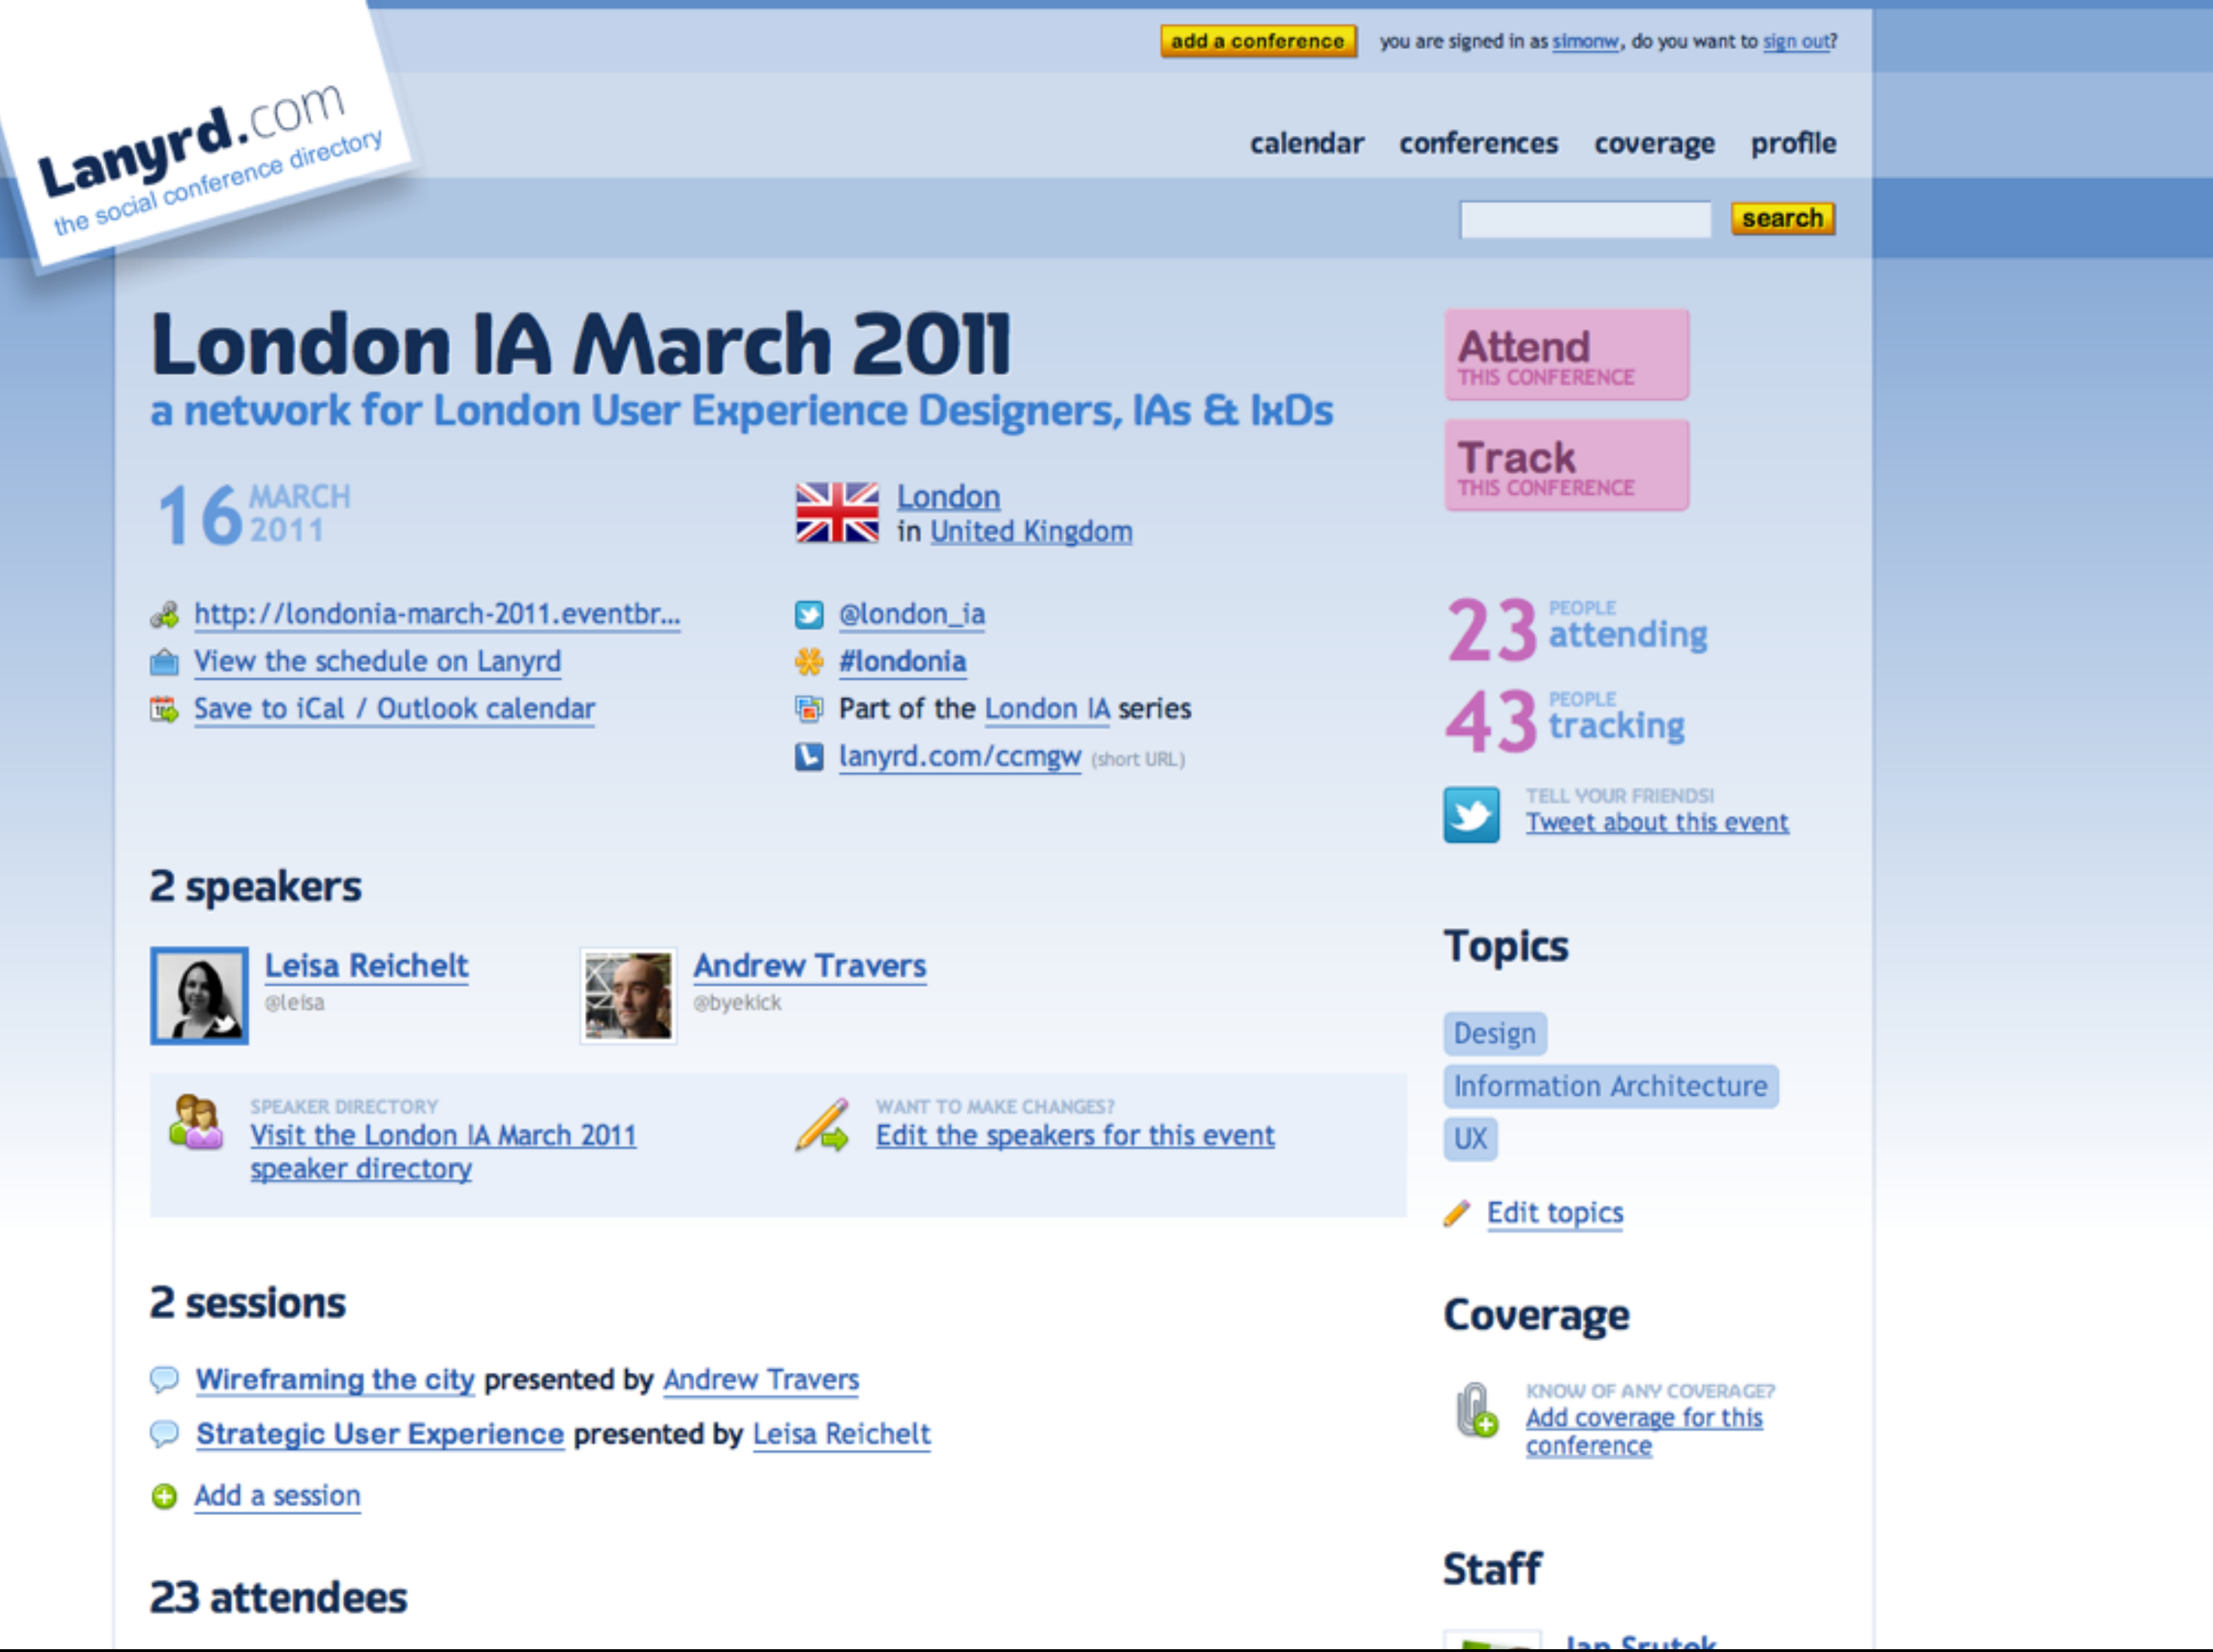

Lanyrd.com
the social conference directory

add a conference you are signed in as simonw, do you want to sign out?

calendar conferences coverage profile

search

London IA March 2011

a network for London User Experience Designers, IAs & IxDs

16 MARCH 2011

London
in United Kingdom

http://londonia-march-2011.eventbr...
View the schedule on Lanyrd
Save to iCal / Outlook calendar

@london_ia
#londonia
Part of the London IA series
lanyrd.com/ccmgw (short URL)

2 speakers

Leisa Reichelt
@leisa

Andrew Travers
@byekick

SPEAKER DIRECTORY
Visit the London IA March 2011 speaker directory

WANT TO MAKE CHANGES?
Edit the speakers for this event

2 sessions

Wireframing the city presented by Andrew Travers

Strategic User Experience presented by Leisa Reichelt

Add a session

23 attendees

Attend
THIS CONFERENCE

Track
THIS CONFERENCE

23 PEOPLE attending

43 PEOPLE tracking

TELL YOUR FRIENDS!
Tweet about this event

Topics

Design

Information Architecture

UX

Edit topics

Coverage

KNOW OF ANY COVERAGE?
Add coverage for this conference

Staff

Ian Srutek

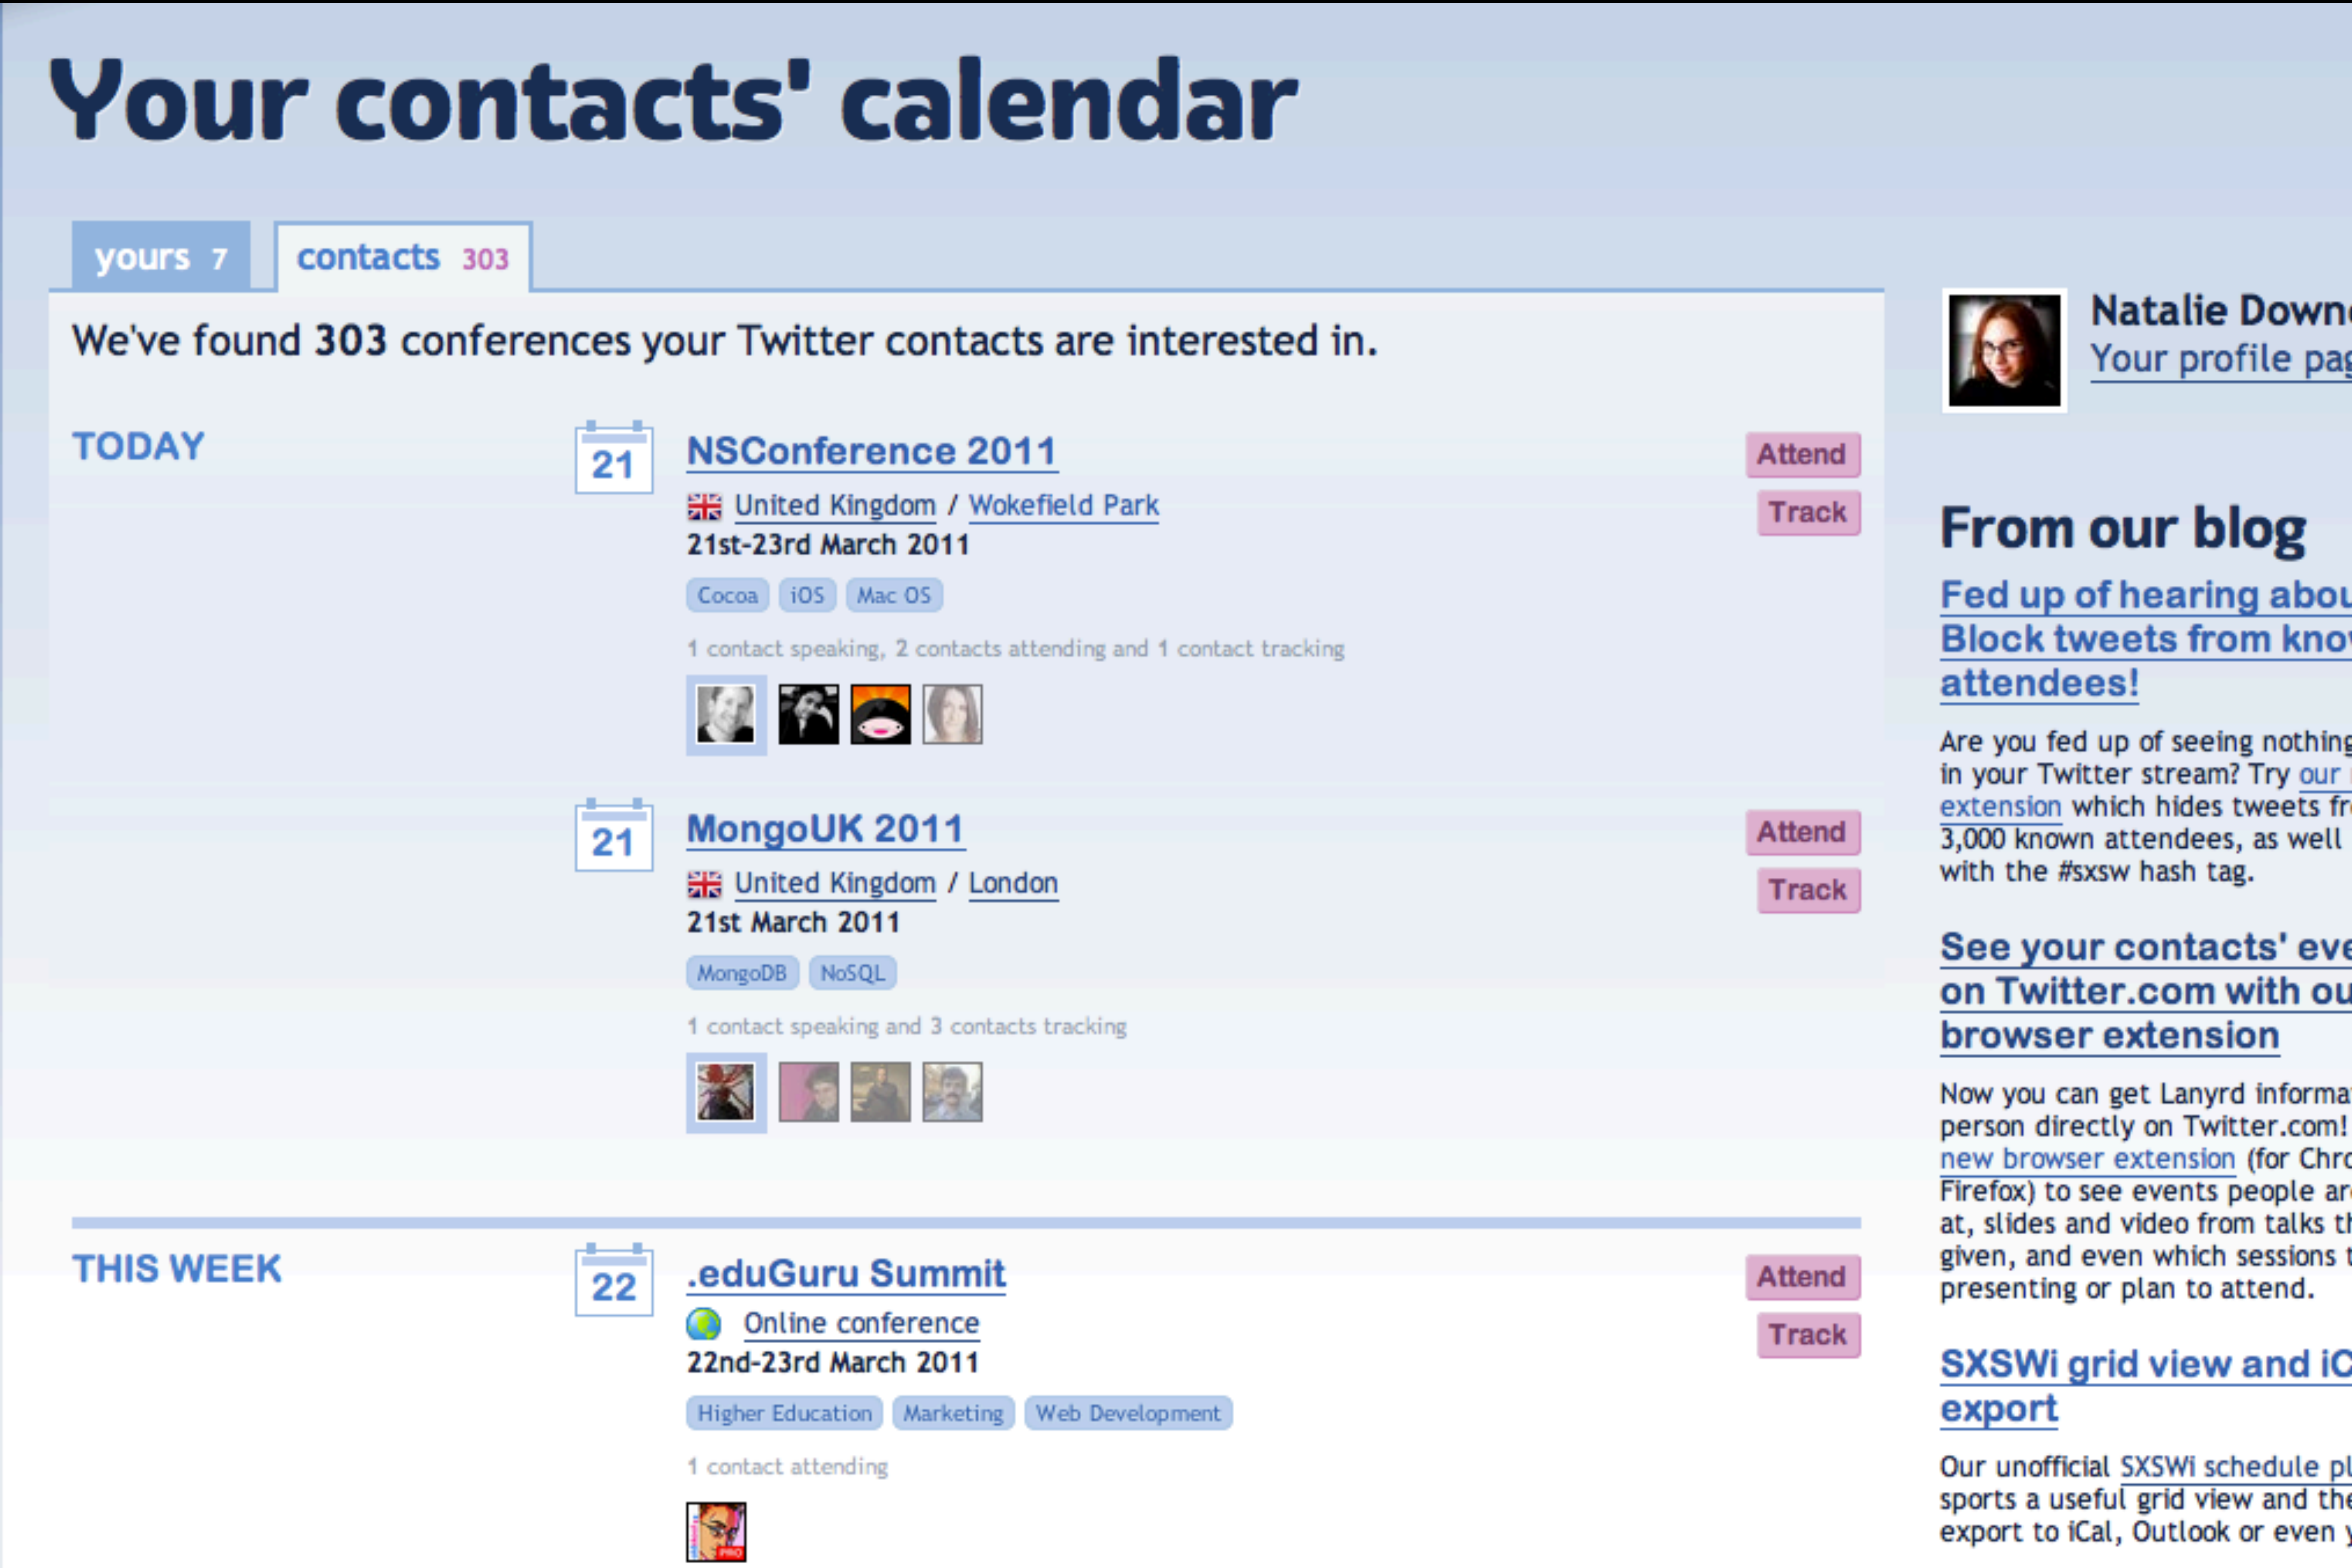

Your contacts' calendar

yours 7 | contacts 303

We've found **303** conferences your Twitter contacts are interested in.

TODAY

21 **NSConference 2011**
United Kingdom / Wokefield Park
21st-23rd March 2011
Cocoa · iOS · Mac OS
1 contact speaking, 2 contacts attending and 1 contact tracking
Attend · Track

21 **MongoUK 2011**
United Kingdom / London
21st March 2011
MongoDB · NoSQL
1 contact speaking and 3 contacts tracking
Attend · Track

THIS WEEK

22 **.eduGuru Summit**
Online conference
22nd-23rd March 2011
Higher Education · Marketing · Web Development
1 contact attending
Attend · Track

Natalie Down
Your profile pag

From our blog

Fed up of hearing abou Block tweets from kno attendees!

Are you fed up of seeing nothing in your Twitter stream? Try our extension which hides tweets fr 3,000 known attendees, as well with the #sxsw hash tag.

See your contacts' eve on Twitter.com with ou browser extension

Now you can get Lanyrd informa person directly on Twitter.com! new browser extension (for Chro Firefox) to see events people ar at, slides and video from talks th given, and even which sessions t presenting or plan to attend.

SXSWi grid view and iC export

Our unofficial SXSWi schedule pl sports a useful grid view and the export to iCal, Outlook or even y

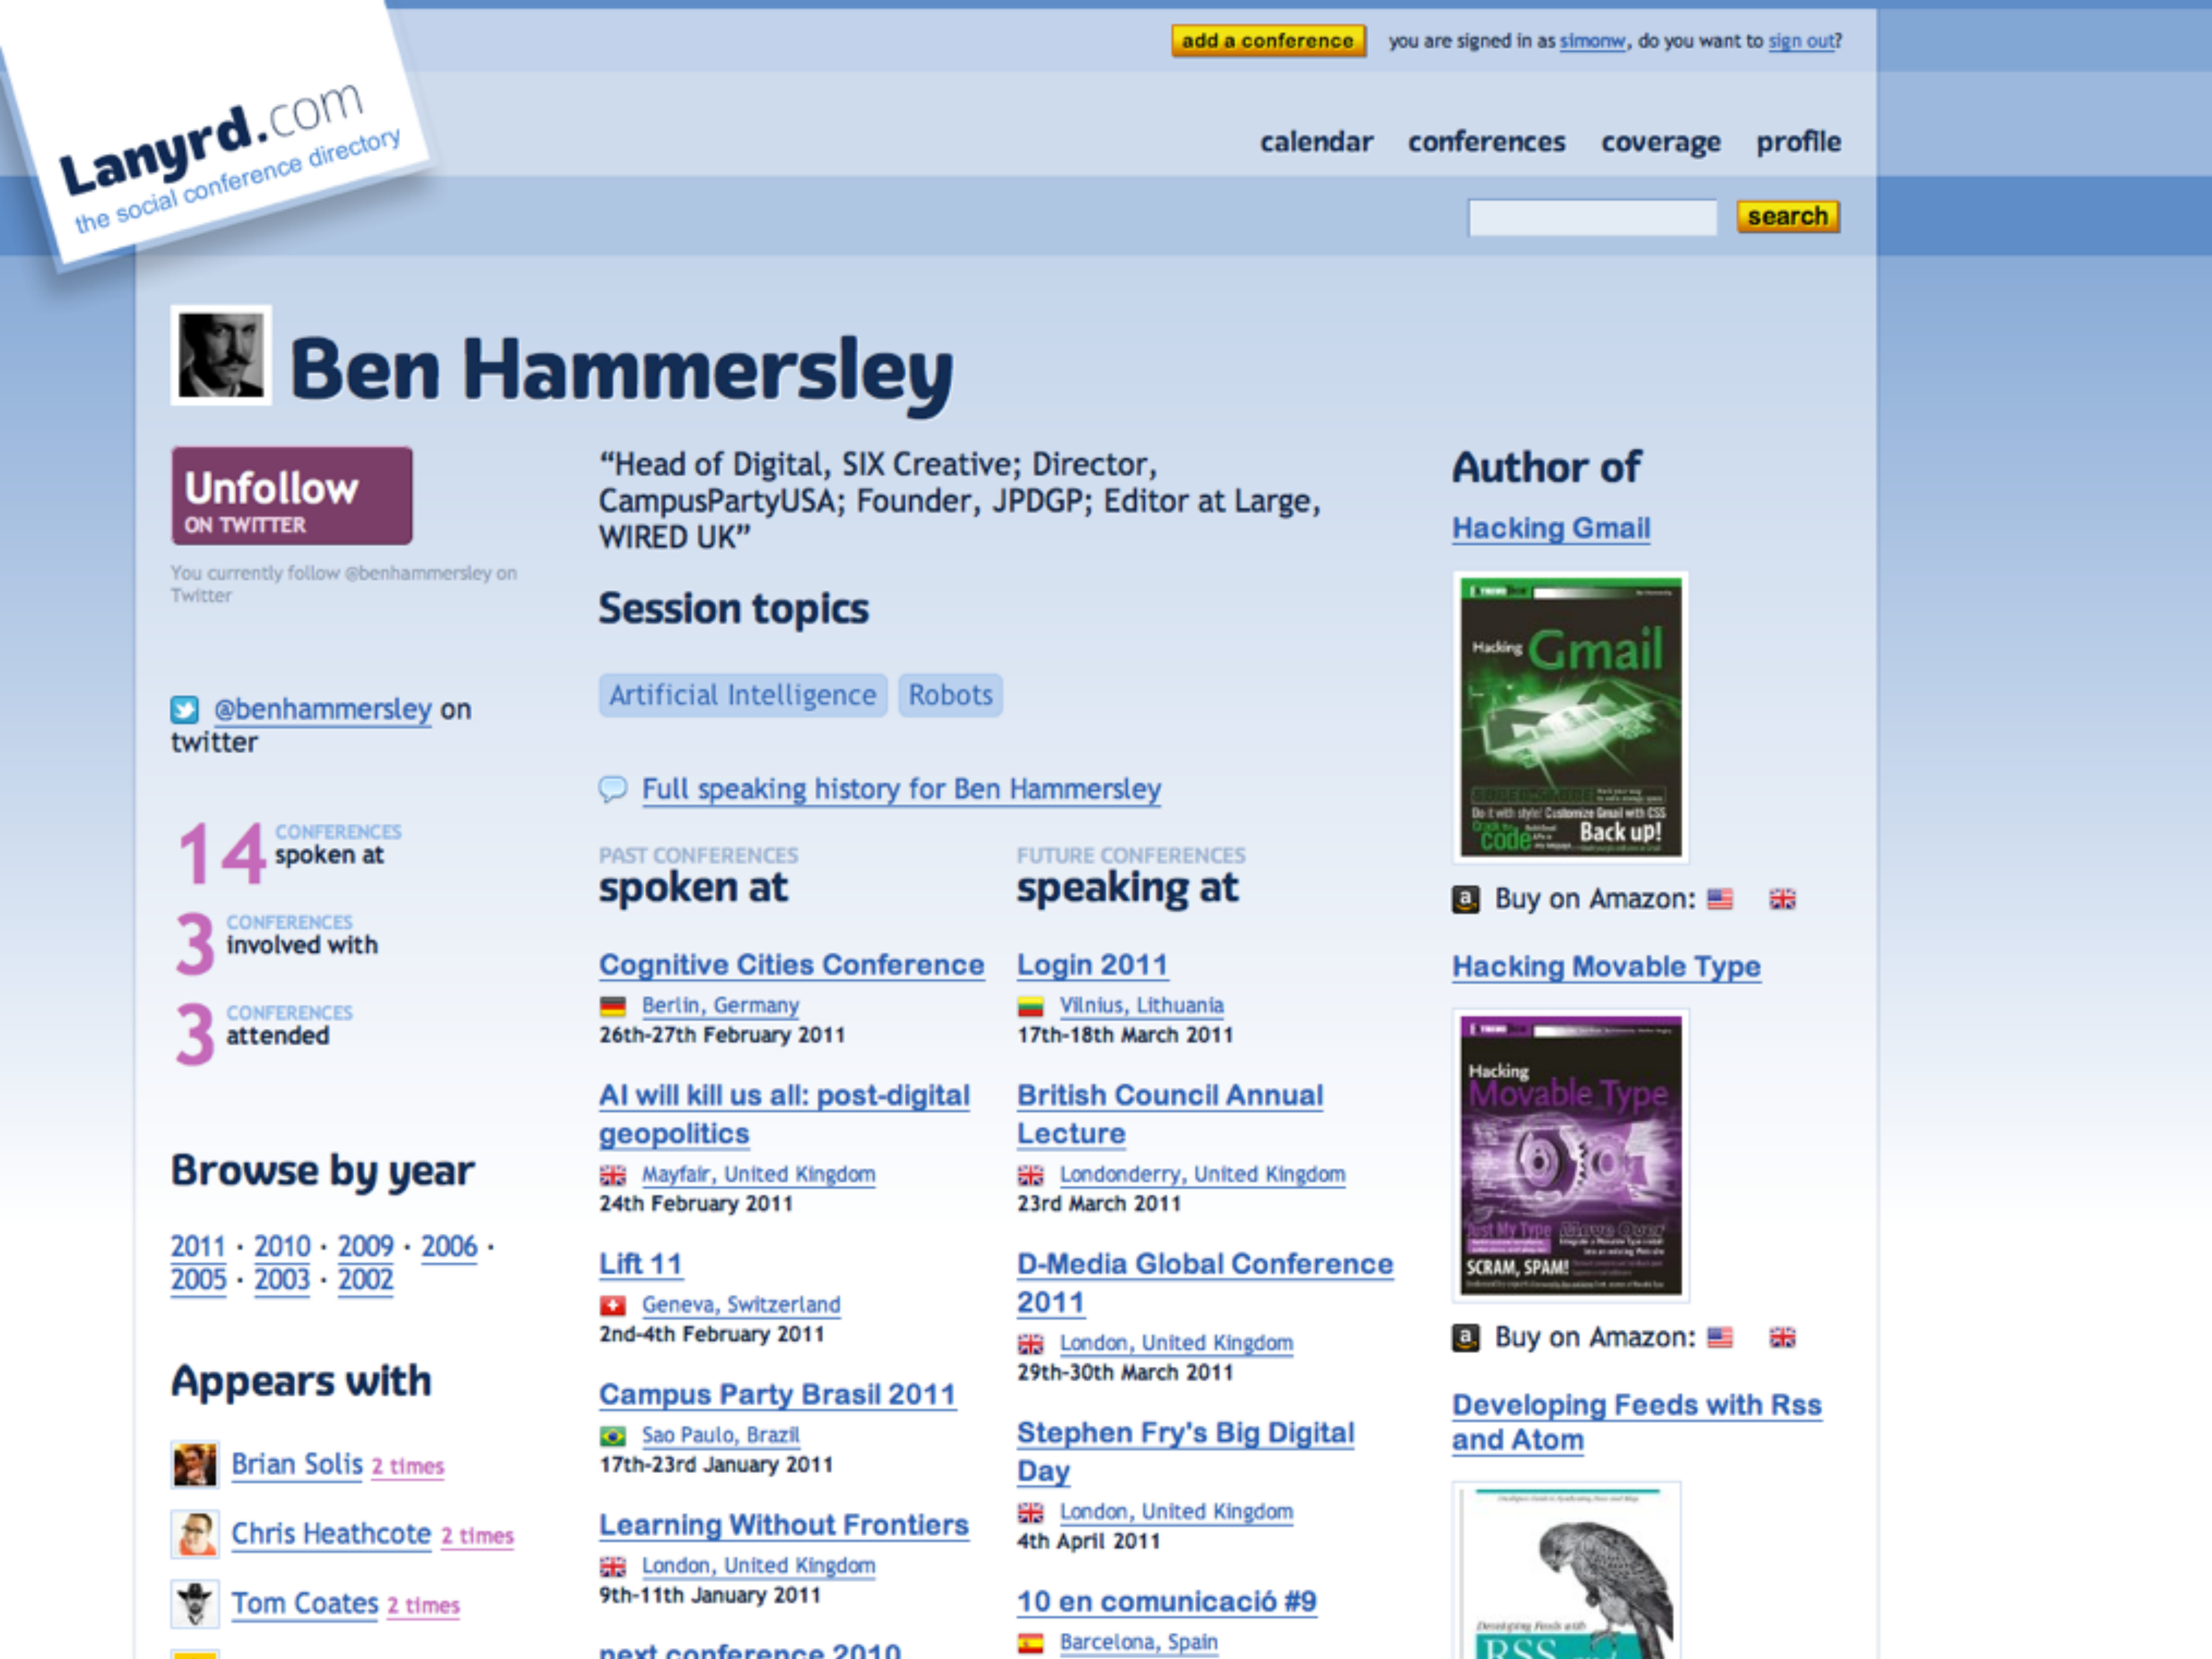

Lanyrd.com
the social conference directory

add a conference you are signed in as simonw, do you want to sign out?

calendar conferences coverage profile

search

Ben Hammersley

Unfollow
ON TWITTER

You currently follow @benhammersley on Twitter

@benhammersley on twitter

14 CONFERENCES spoken at

3 CONFERENCES involved with

3 CONFERENCES attended

Browse by year

2011 · 2010 · 2009 · 2006 · 2005 · 2003 · 2002

Appears with

Brian Solis 2 times

Chris Heathcote 2 times

Tom Coates 2 times

"Head of Digital, SIX Creative; Director, CampusPartyUSA; Founder, JPDGP; Editor at Large, WIRED UK"

Session topics

Artificial Intelligence Robots

Full speaking history for Ben Hammersley

PAST CONFERENCES
spoken at

Cognitive Cities Conference
Berlin, Germany
26th-27th February 2011

AI will kill us all: post-digital geopolitics
Mayfair, United Kingdom
24th February 2011

Lift 11
Geneva, Switzerland
2nd-4th February 2011

Campus Party Brasil 2011
Sao Paulo, Brazil
17th-23rd January 2011

Learning Without Frontiers
London, United Kingdom
9th-11th January 2011

next conference 2010

FUTURE CONFERENCES
speaking at

Login 2011
Vilnius, Lithuania
17th-18th March 2011

British Council Annual Lecture
Londonderry, United Kingdom
23rd March 2011

D-Media Global Conference 2011
London, United Kingdom
29th-30th March 2011

Stephen Fry's Big Digital Day
London, United Kingdom
4th April 2011

10 en comunicació #9
Barcelona, Spain

Author of

Hacking Gmail

Buy on Amazon:

Hacking Movable Type

Buy on Amazon:

Developing Feeds with Rss and Atom

Haystack

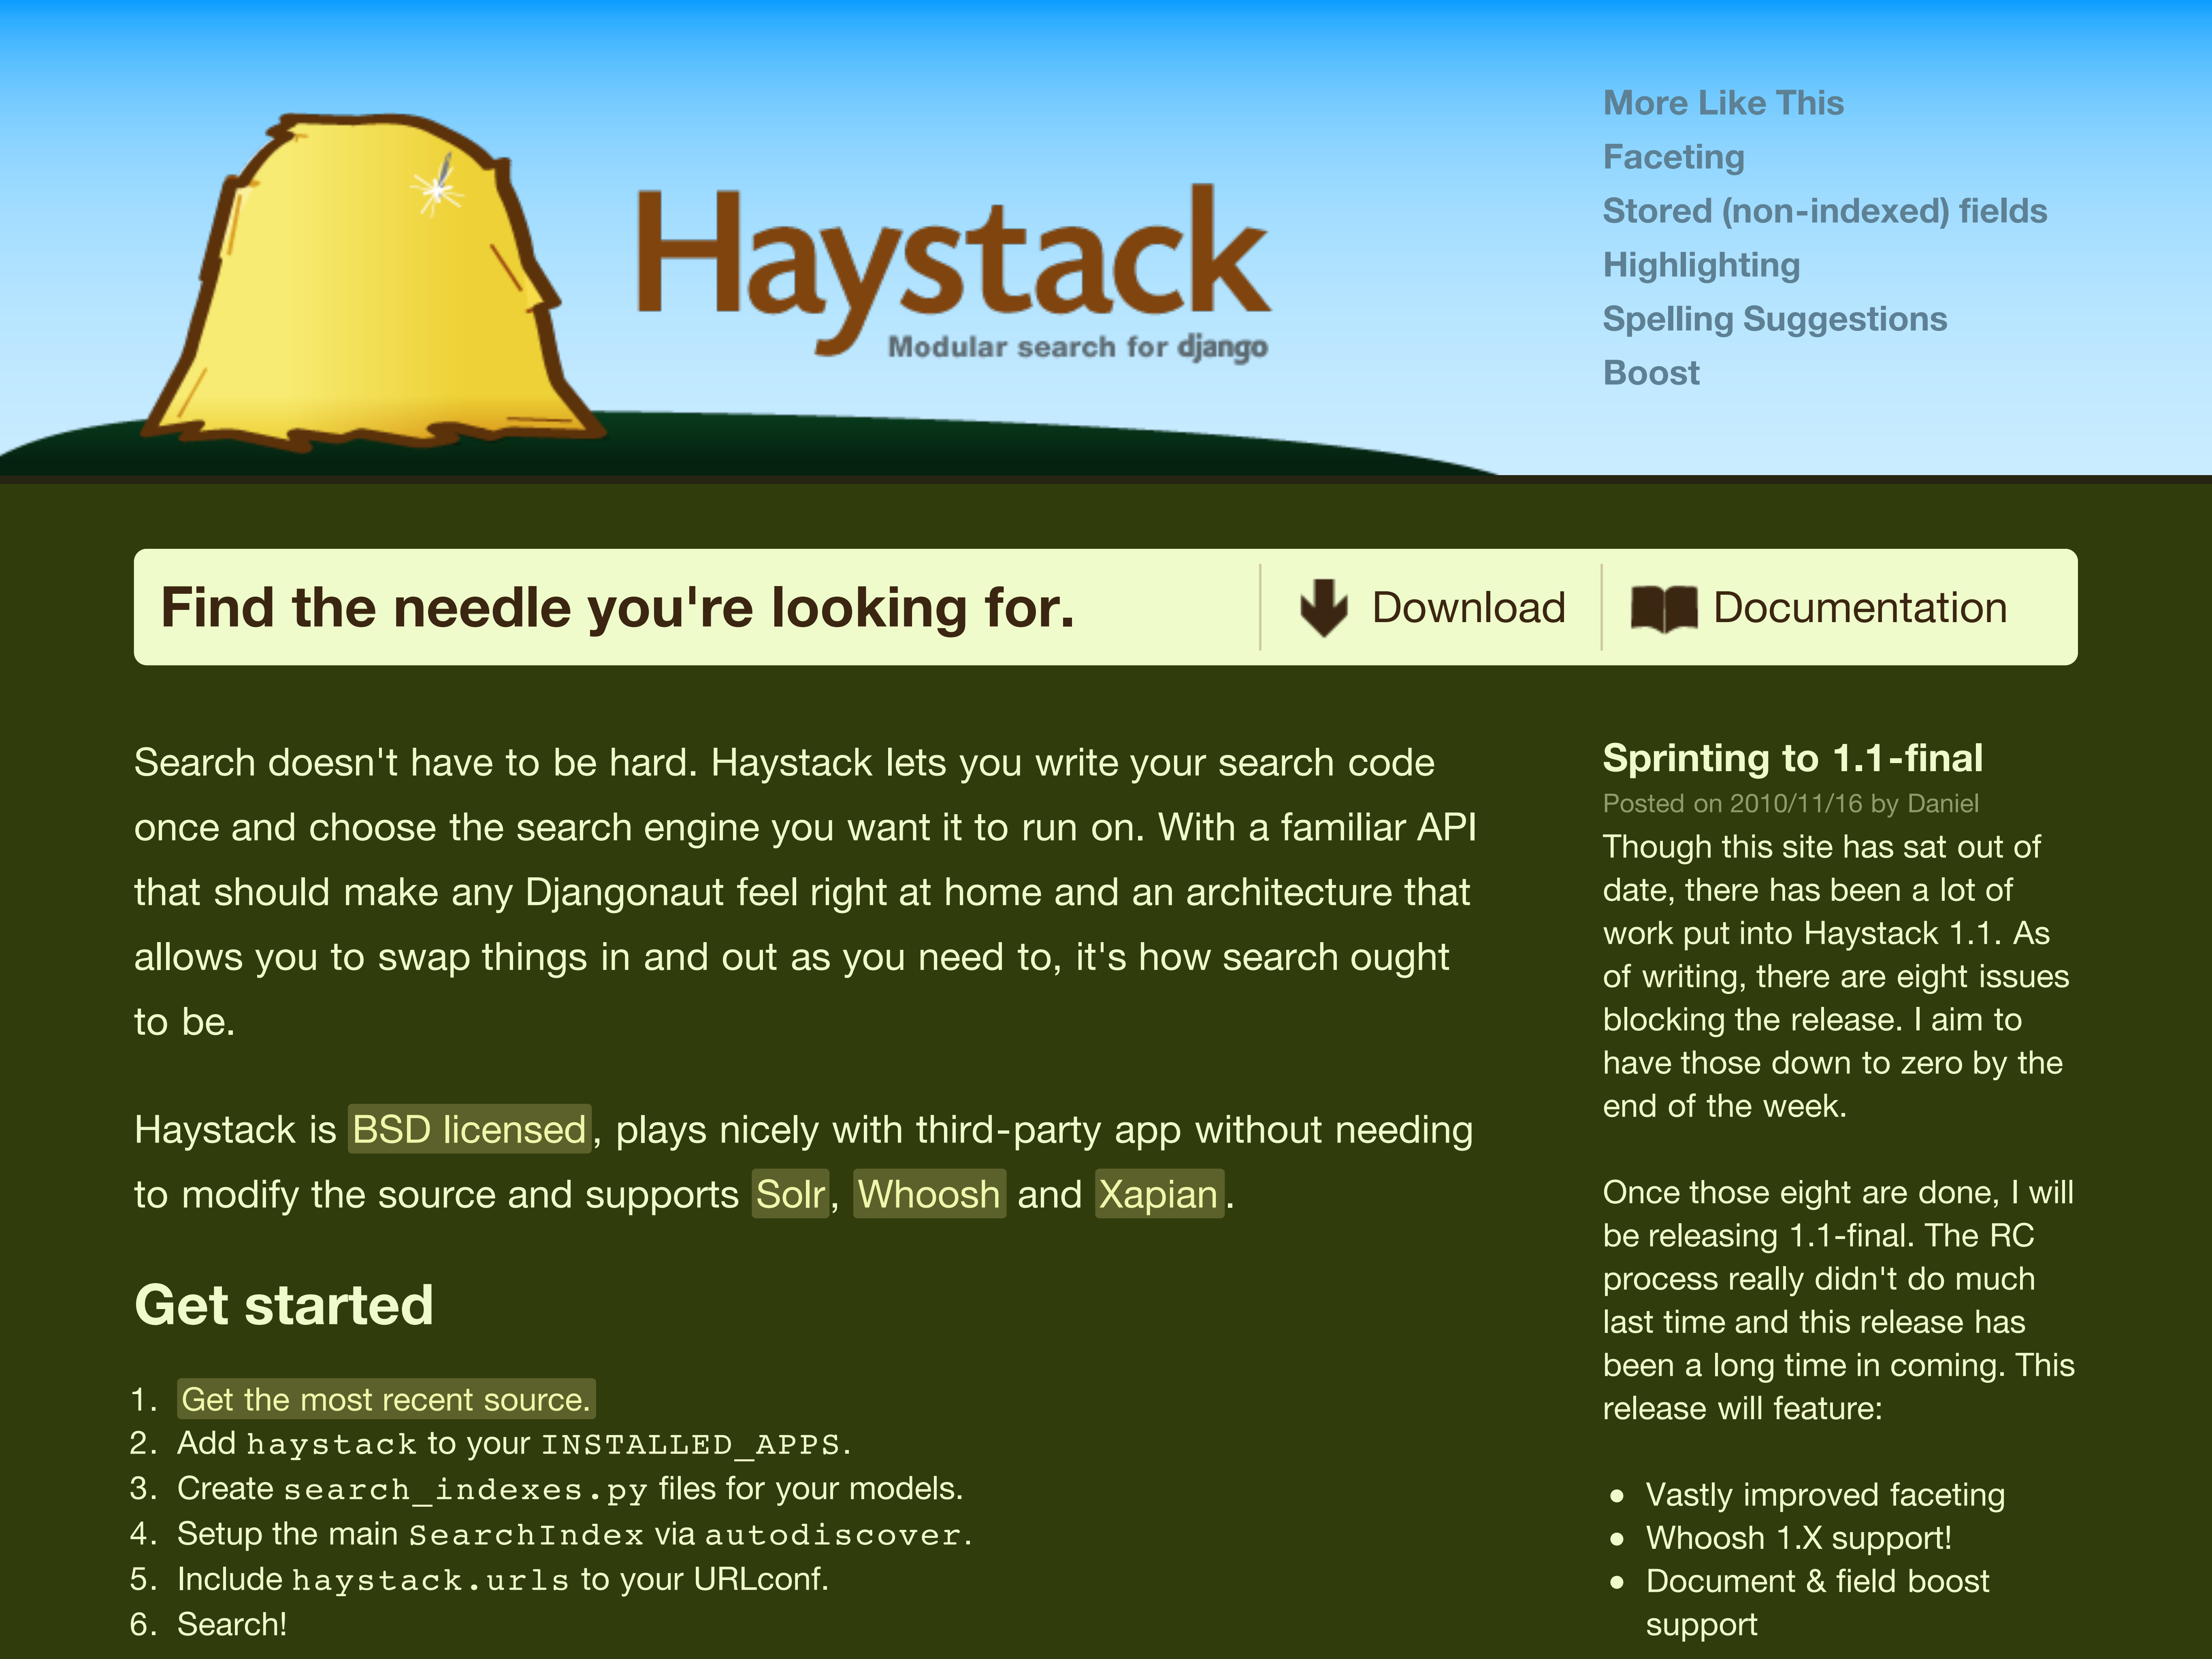

- More Like This
- Faceting
- Stored (non-indexed) fields
- Highlighting
- Spelling Suggestions
- Boost

Find the needle you're looking for.

Search doesn't have to be hard. Haystack lets you write your search code once and choose the search engine you want it to run on. With a familiar API that should make any Djangonaut feel right at home and an architecture that allows you to swap things in and out as you need to, it's how search ought to be.

Haystack is BSD licensed, plays nicely with third-party app without needing to modify the source and supports Solr, Whoosh and Xapian.

Get started

1. Get the most recent source.
2. Add `haystack` to your `INSTALLED_APPS`.
3. Create `search_indexes.py` files for your models.
4. Setup the main `SearchIndex` via `autodiscover`.
5. Include `haystack.urls` to your URLconf.
6. Search!

Sprinting to 1.1-final

Posted on 2010/11/16 by Daniel

Though this site has sat out of date, there has been a lot of work put into Haystack 1.1. As of writing, there are eight issues blocking the release. I aim to have those down to zero by the end of the week.

Once those eight are done, I will be releasing 1.1-final. The RC process really didn't do much last time and this release has been a long time in coming. This release will feature:

- Vastly improved faceting
- Whoosh 1.X support!
- Document & field boost support

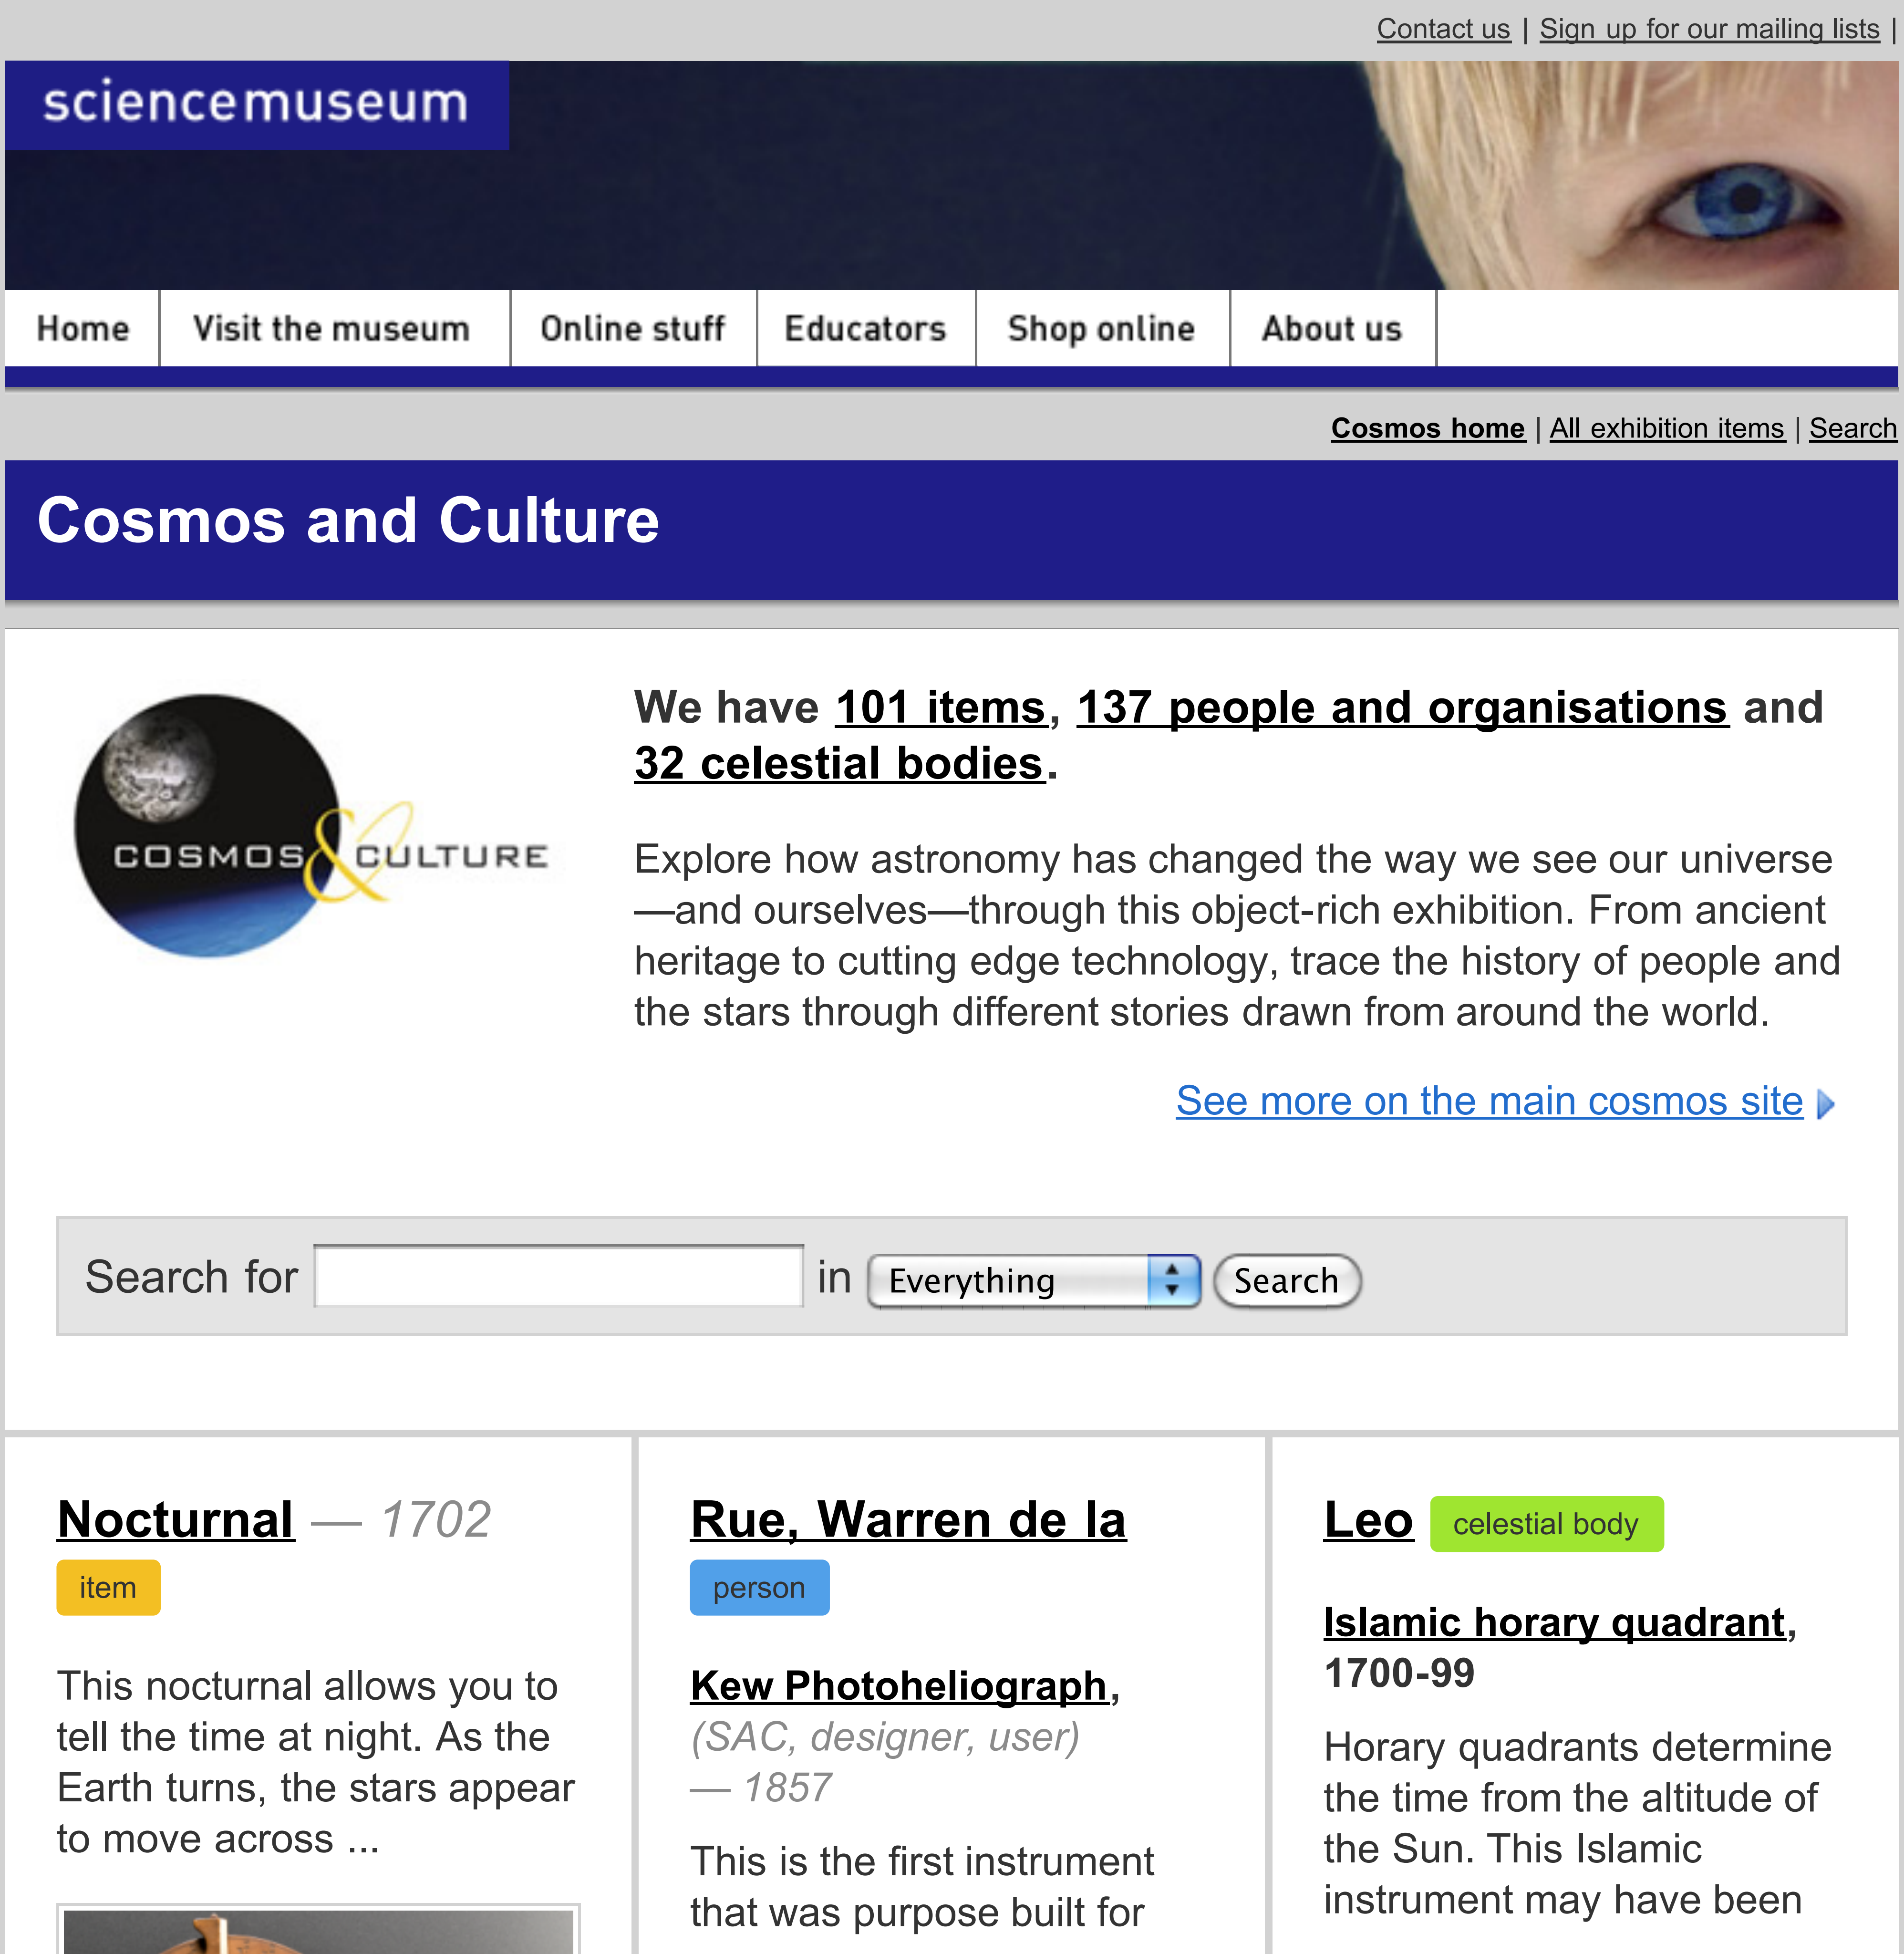

Contact us | Sign up for our mailing lists |

sciencemuseum

Home | Visit the museum | Online stuff | Educators | Shop online | About us

Cosmos home | All exhibition items | Search

Cosmos and Culture

We have 101 items, 137 people and organisations and 32 celestial bodies.

Explore how astronomy has changed the way we see our universe—and ourselves—through this object-rich exhibition. From ancient heritage to cutting edge technology, trace the history of people and the stars through different stories drawn from around the world.

See more on the main cosmos site ▸

Search for [ ] in Everything ▾ Search

Nocturnal — *1702*

item

This nocturnal allows you to tell the time at night. As the Earth turns, the stars appear to move across ...

Rue, Warren de la

person

Kew Photoheliograph, *(SAC, designer, user)* — *1857*

This is the first instrument that was purpose built for

Leo celestial body

Islamic horary quadrant, 1700-99

Horary quadrants determine the time from the altitude of the Sun. This Islamic instrument may have been

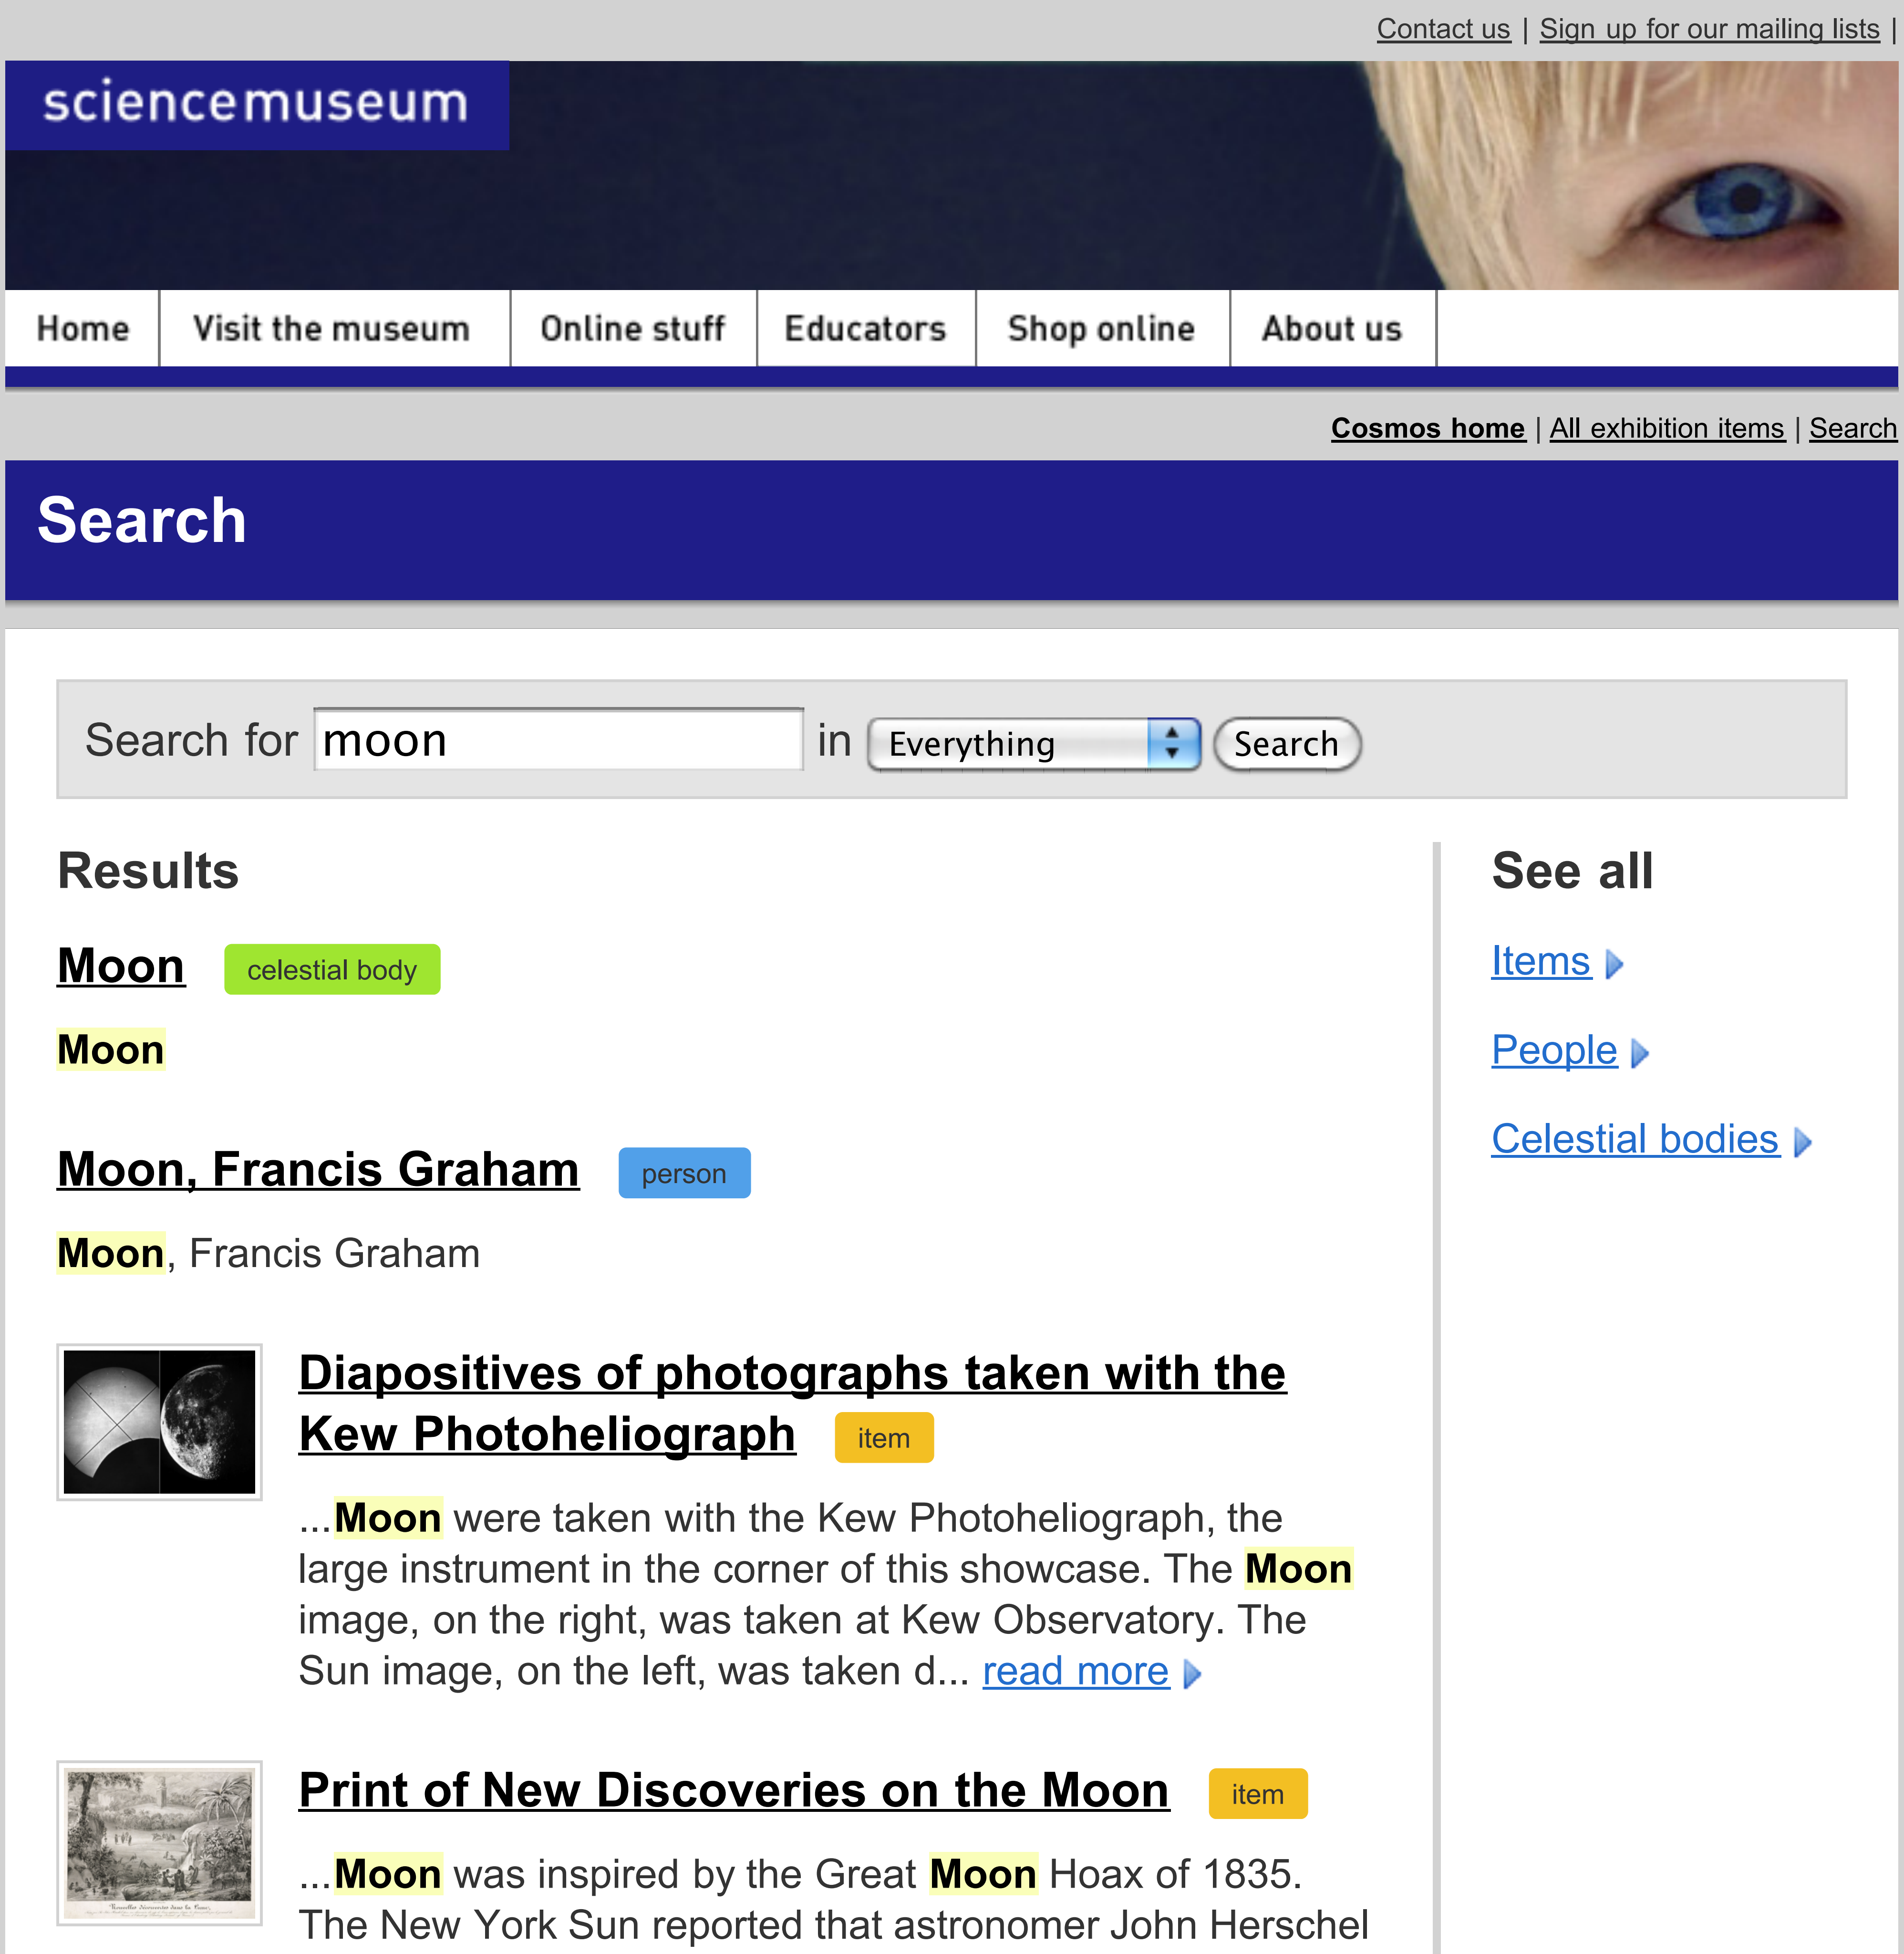
Contact us | Sign up for our mailing lists |
sciencemuseum
Home
Visit the museum
Online stuff
Educators
Shop online
About us
Cosmos home | All exhibition items | Search

Search

Search for moon in Everything

Results

Moon celestial body

Moon

Moon, Francis Graham person

Moon, Francis Graham

Diapositives of photographs taken with the Kew Photoheliograph item

...Moon were taken with the Kew Photoheliograph, the large instrument in the corner of this showcase. The Moon image, on the right, was taken at Kew Observatory. The Sun image, on the left, was taken d... read more

Print of New Discoveries on the Moon item

...Moon was inspired by the Great Moon Hoax of 1835. The New York Sun reported that astronomer John Herschel

See all

Model-oriented search

- Define **search_indexes.py** (like admin.py) for your application
- Hook up default haystack search views
- Write a quick **search.html** template
- Run **./manage.py rebuild_index**

```
from haystack import indexes
from haystack import site
from models import MuseumObject, Person, CelestialBody

class MuseumObjectIndex(indexes.SearchIndex):
    text = indexes.CharField(document=True, model_attr='text')
    def get_queryset(self):
        return MuseumObject.objects.all()

site.register(MuseumObject, MuseumObjectIndex)


class PersonIndex(indexes.SearchIndex):
    text = indexes.CharField(document=True, model_attr='name')
    def get_queryset(self):
        return Person.objects.all()

site.register(Person, PersonIndex)
```

```
(r'^search/', include('haystack.urls')),
<ul class="listing">
{% for result in page.object_list %}
    {% if result.model_name == "museumobject" %}
    <li class="item">
        <img src="{{ result.object.image_inline }}" width="75" alt="">
        <h3><a href="{{ result.object.get_absolute_url }}">{{ result.object.name }}</a>
        <span class="type-indicator">item</span></h3>
        <p>{% highlight result.text with request.GET.q %}</p>
    </li>
    {% endif %}
    {% if result.model_name == "person" %}
    <li class="person">
        <h3><a href="{{ result.object.get_absolute_url }}">{{ result.object.name }}</a>
        <span class="type-indicator">person</span></h3>
        <p>{% highlight result.text with request.GET.q %}</p>
    </li>
    {% endif %}
{% endfor %}
</ul>
```

Pick your backend

- Whoosh - pure Python
  - For sites with no write traffic
- Solr - Java web service application server
  - Best bet for medium-huge sites
- Xapian - embedded C library
  - Haven’t tried this myself

Main Wiki

Search the site with Solr Search
Powered by Lucid Imagination

Last Published: Sat, 04 Jun 2011 12:23:42 GMT

About
- Welcome
- Who We Are

Documentation
Resources
Related Projects

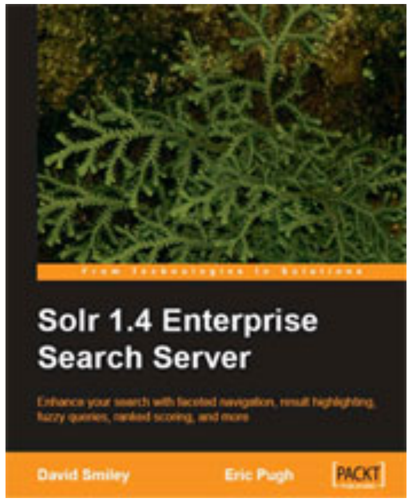

Welcome to Solr

PDF

What Is Solr?

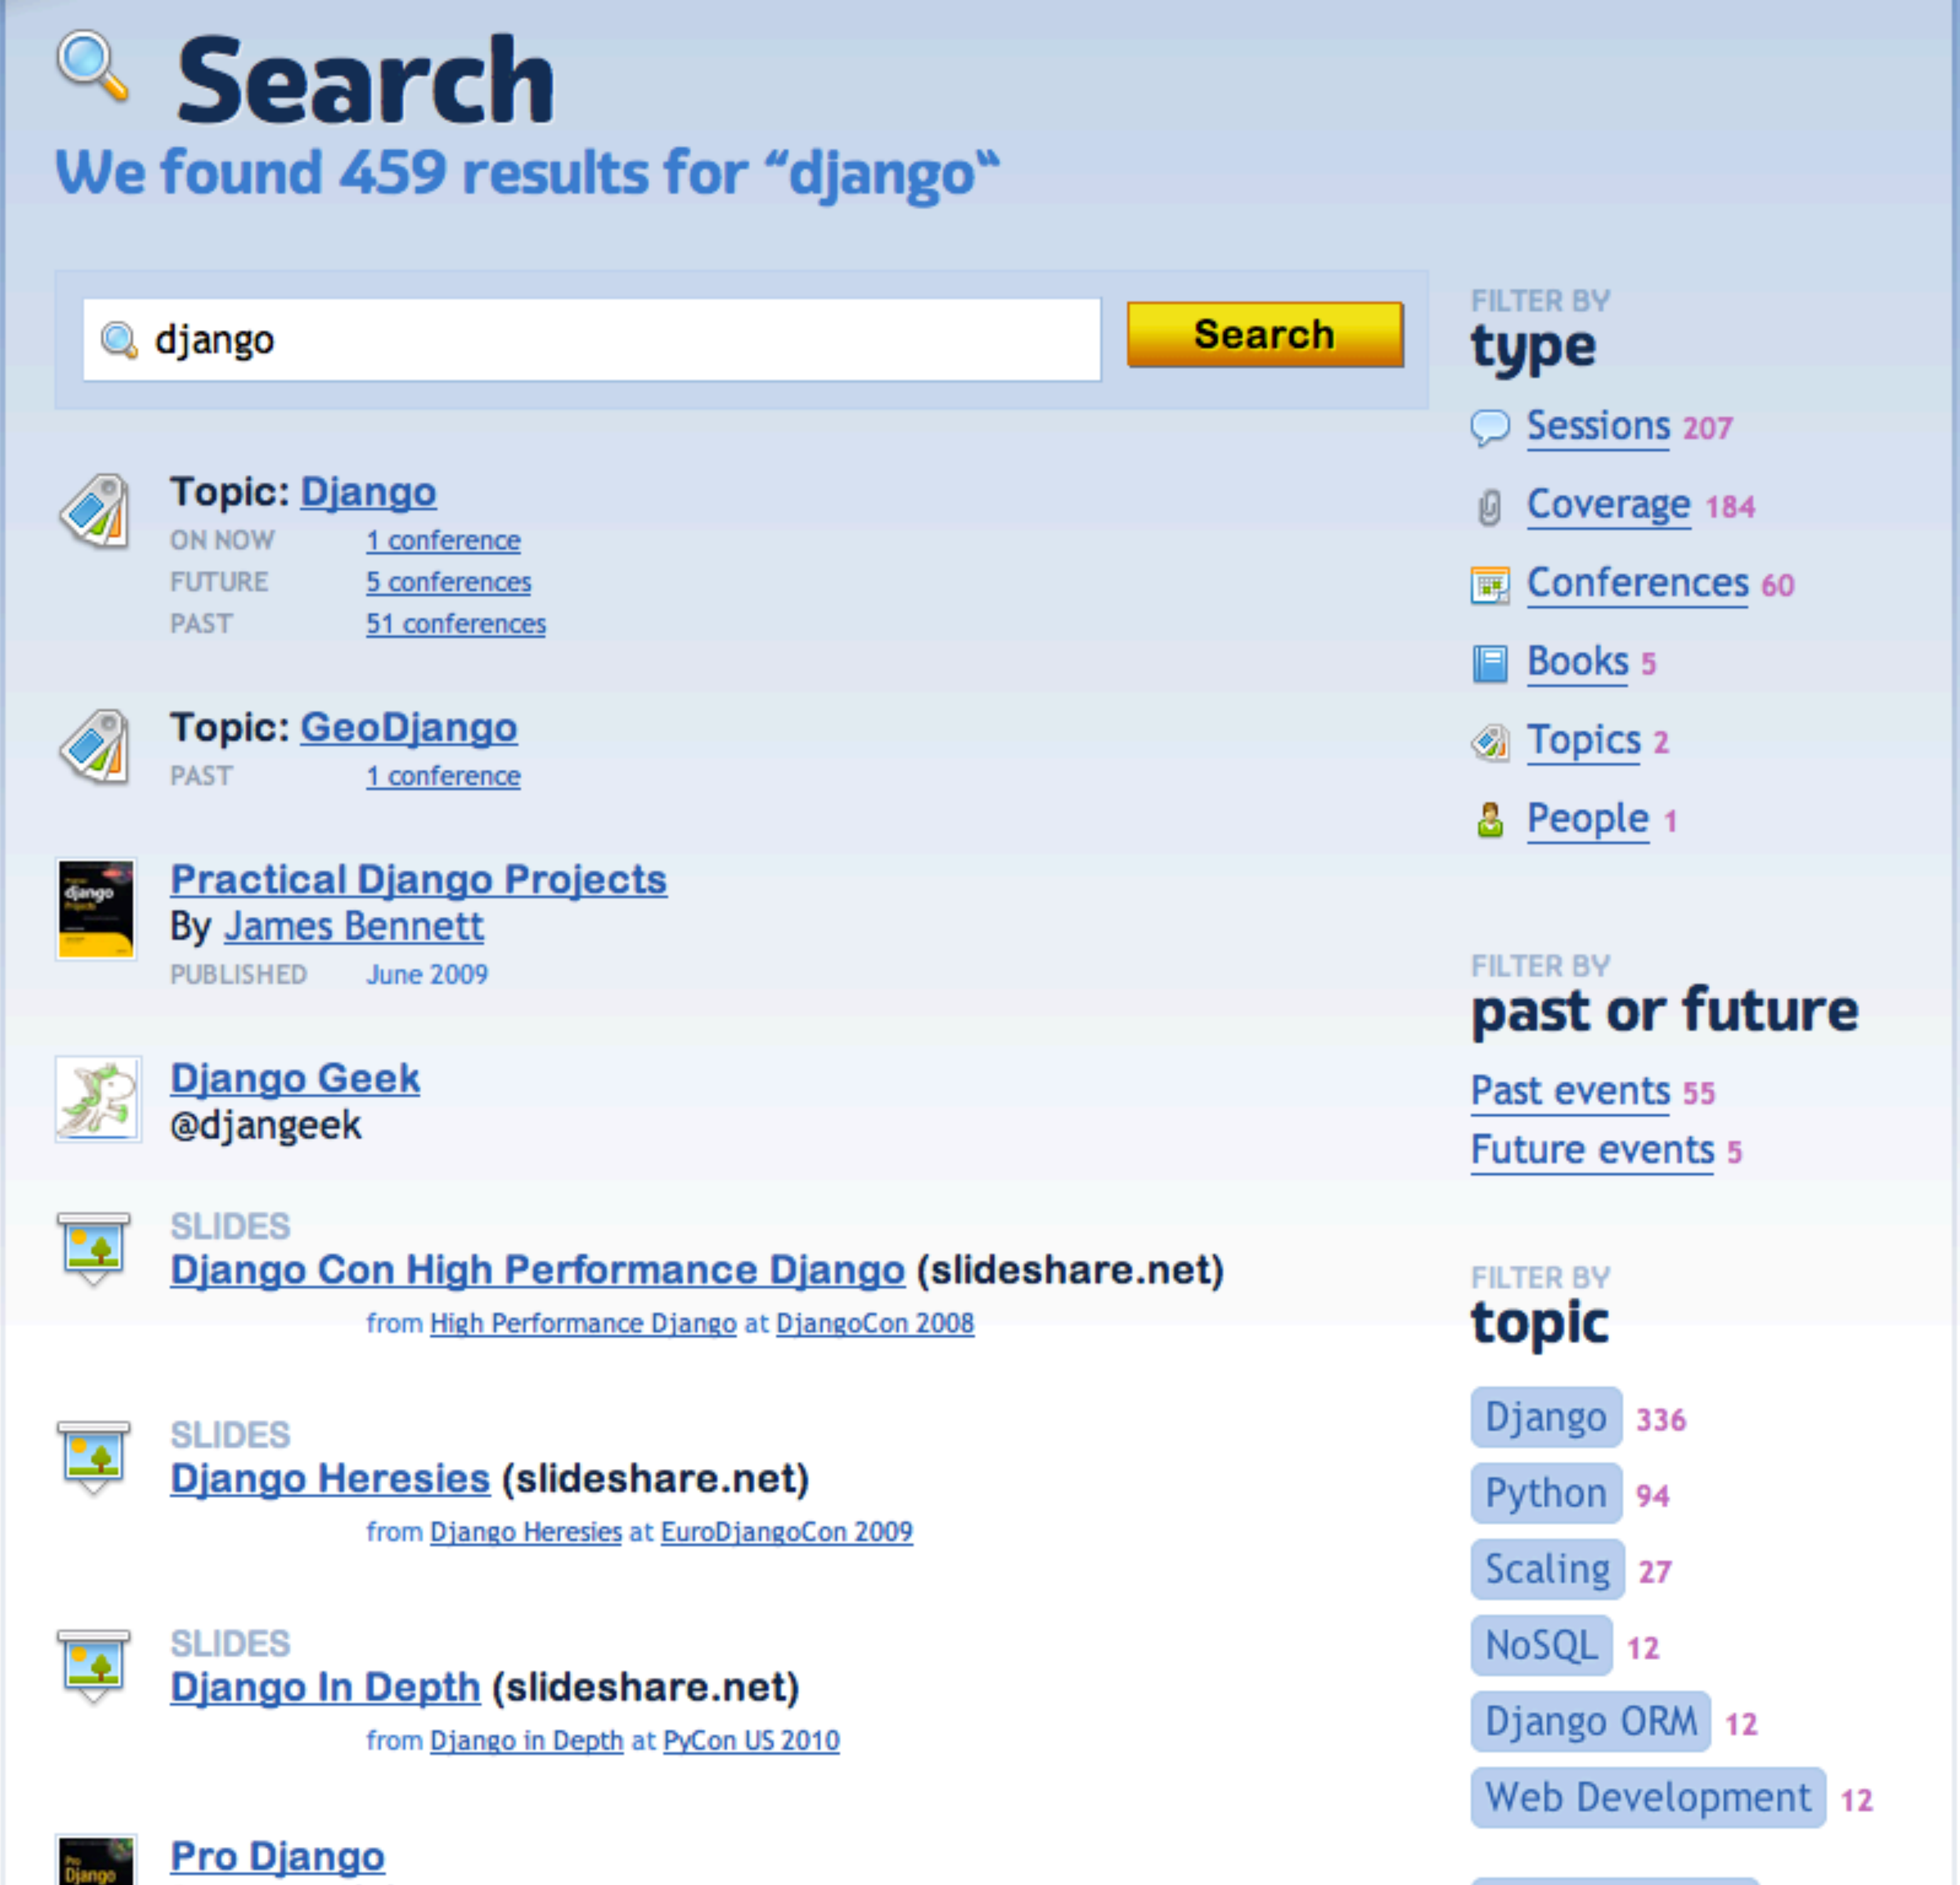

Search

We found 459 results for “django“

django

Topic: Django

ON NOW	1 conference
FUTURE	5 conferences
PAST	51 conferences

Topic: GeoDjango

PAST	1 conference

Practical Django Projects
By James Bennett
PUBLISHED June 2009

Django Geek
@djangeek

SLIDES
Django Con High Performance Django (slideshare.net)
from High Performance Django at DjangoCon 2008

SLIDES
Django Heresies (slideshare.net)
from Django Heresies at EuroDjangoCon 2009

SLIDES
Django In Depth (slideshare.net)
from Django in Depth at PyCon US 2010

Pro Django

FILTER BY
type

Sessions	207
Coverage	184
Conferences	60
Books	5
Topics	2
People	1

FILTER BY
past or future

Past events	55
Future events	5

FILTER BY
topic

Django	336
Python	94
Scaling	27
NoSQL	12
Django ORM	12
Web Development	12

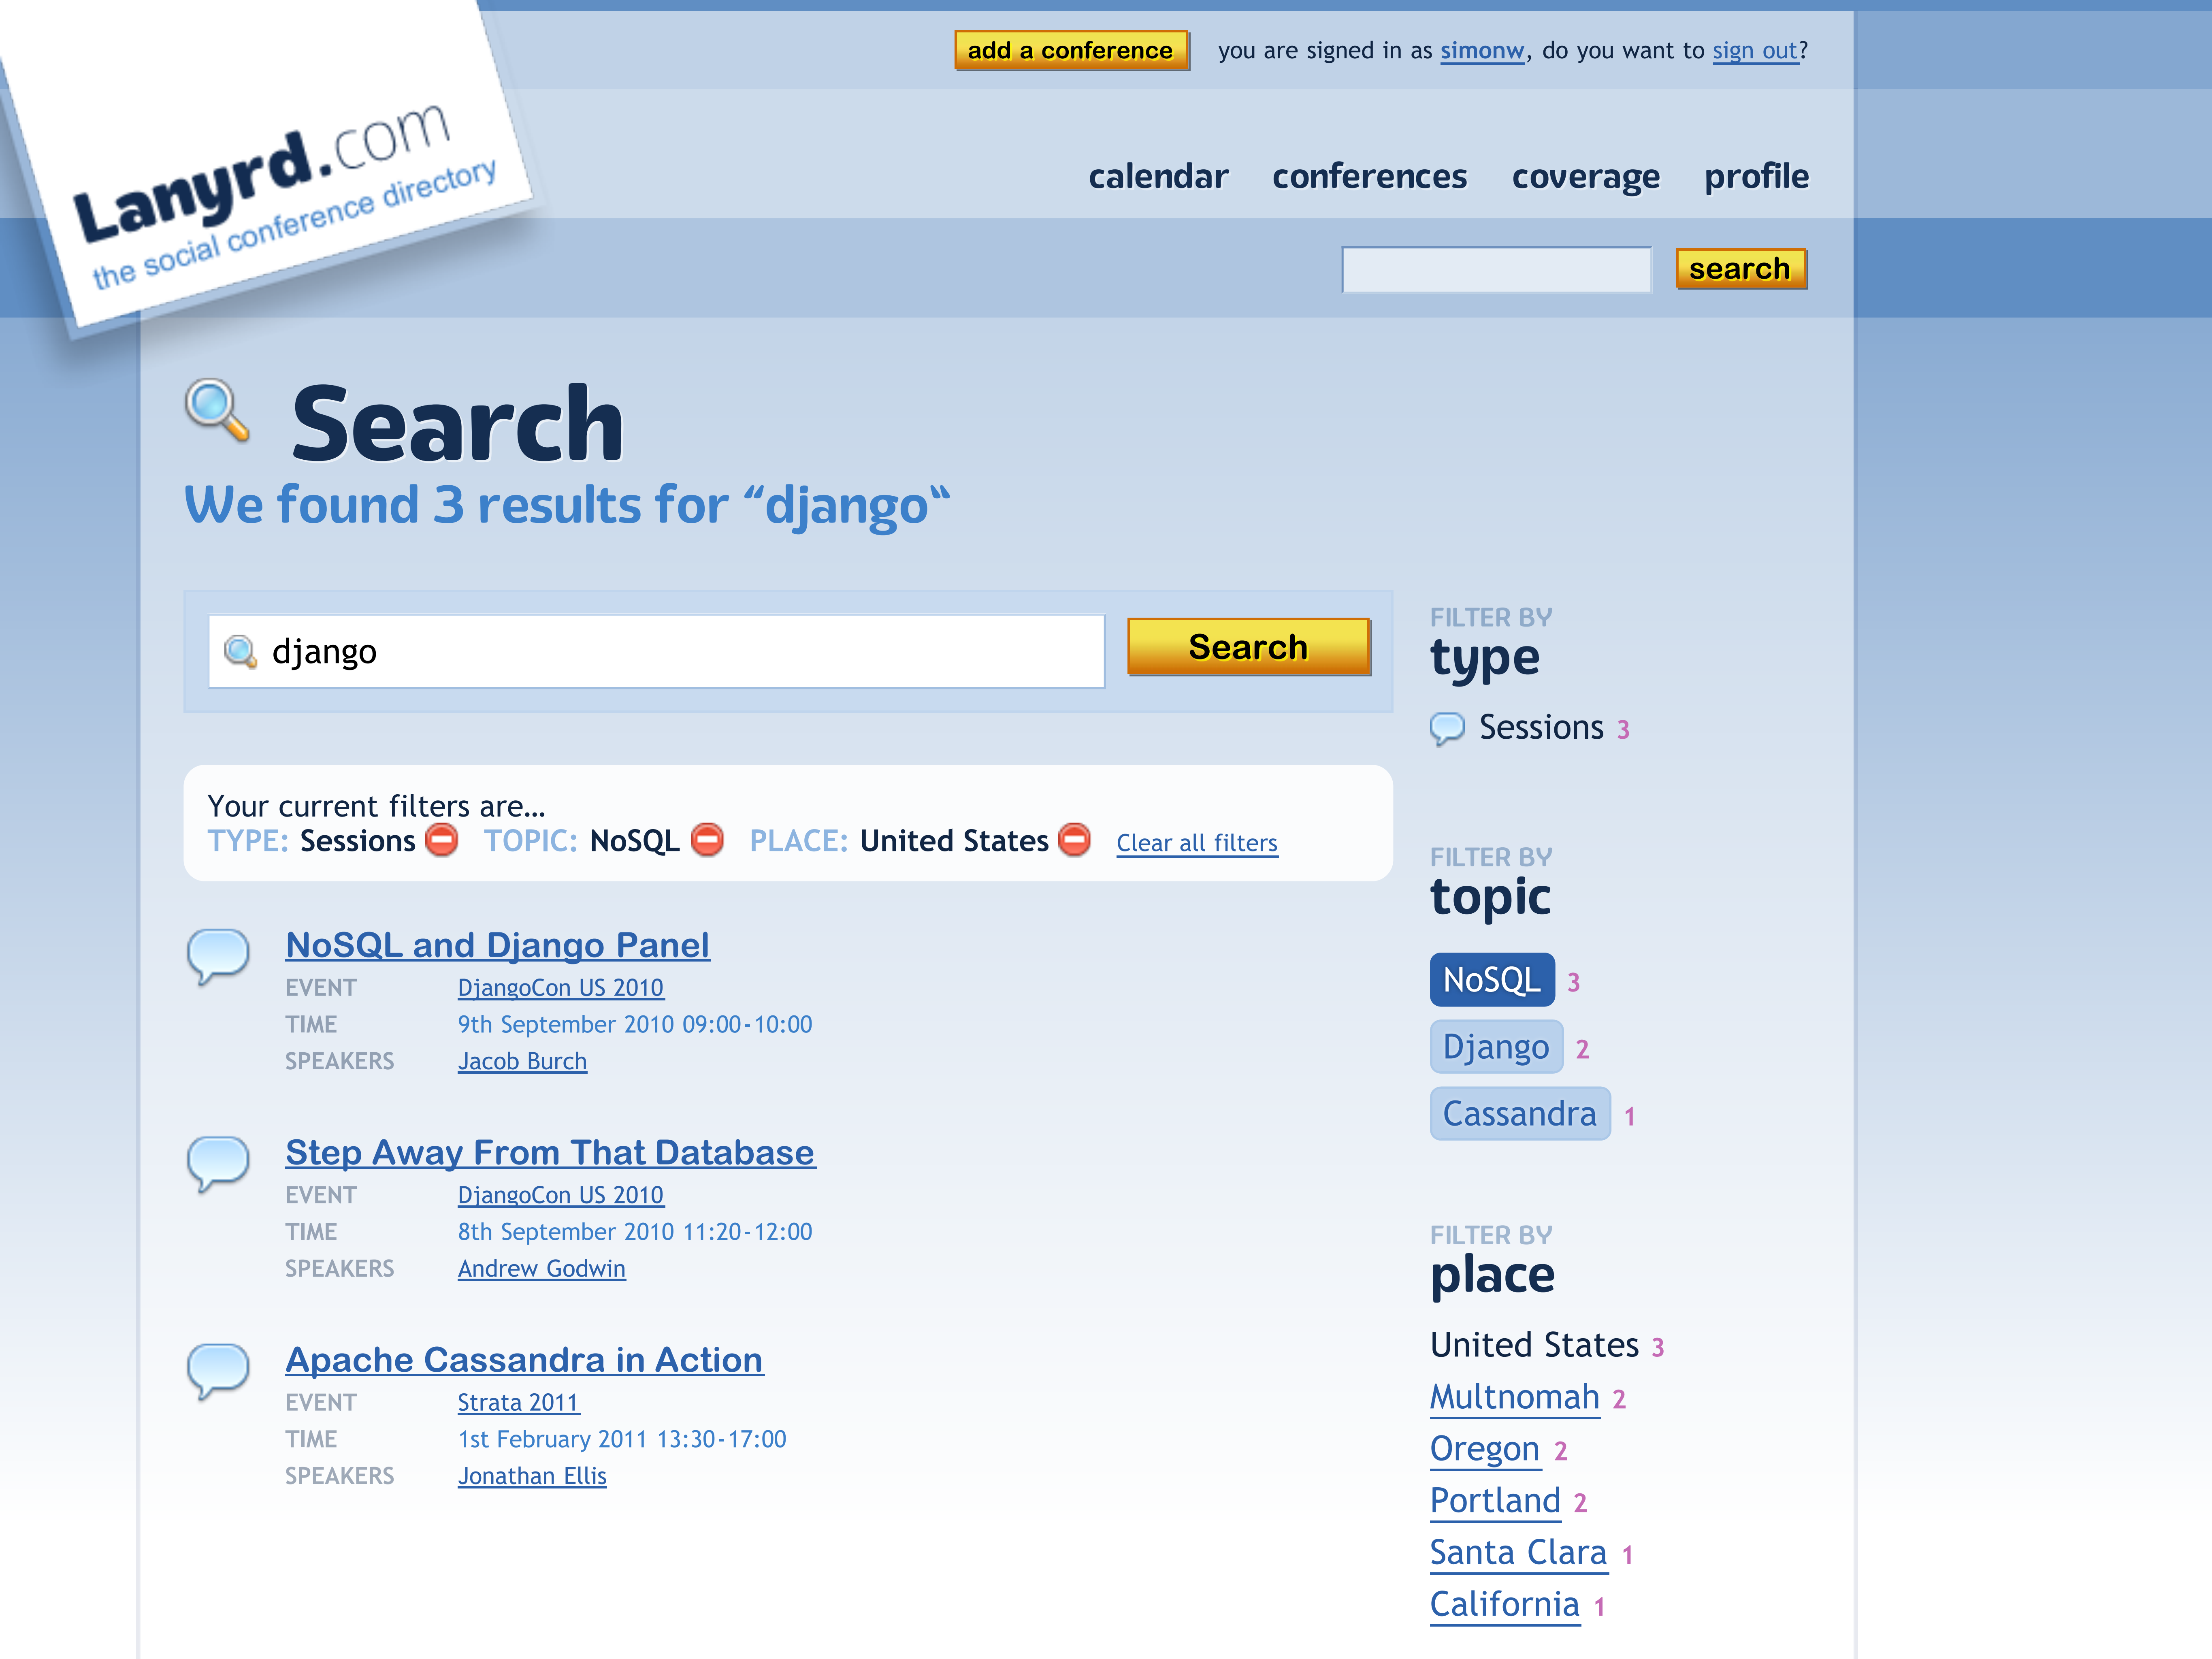

you are signed in as simonw, do you want to sign out?

Search

We found 3 results for “django“

django

Search

Your current filters are...
TYPE: Sessions ⛔ TOPIC: NoSQL ⛔ PLACE: United States ⛔ Clear all filters

NoSQL and Django Panel

EVENT	DjangoCon US 2010
TIME	9th September 2010 09:00-10:00
SPEAKERS	Jacob Burch

Step Away From That Database

EVENT	DjangoCon US 2010
TIME	8th September 2010 11:20-12:00
SPEAKERS	Andrew Godwin

Apache Cassandra in Action

EVENT	Strata 2011
TIME	1st February 2011 13:30-17:00
SPEAKERS	Jonathan Ellis

FILTER BY
type

Sessions	3

FILTER BY
topic

NoSQL	3
Django	2
Cassandra	1

FILTER BY
place

United States	3
Multnomah	2
Oregon	2
Portland	2
Santa Clara	1
California	1

```
class BookIndex(indexes.SearchIndex):
    text = indexes.CharField(document=True, use_template=True)
    speakers = indexes.MultiValueField()
    topics = indexes.MultiValueField()

    def prepare_speakers(self, obj):
        return [a.user.t_id for a in obj.authors.exclude(
            user = None
        ).select_related('user')]

    def prepare_topics(self, obj):
        return list(obj.topics.values_list('pk', flat=True))
```

search/indexes/books/book_text.txt

```
{{ object.title }}
{{ object.tagline }}
{% for author in object.authors.all %}
    {{ author.display_name }}
    {{ author.user.t_screen_name }}
{% endfor %}
{% for topic in object.topics.all %}
    {{ topic.name_en }}
{% endfor %}
```

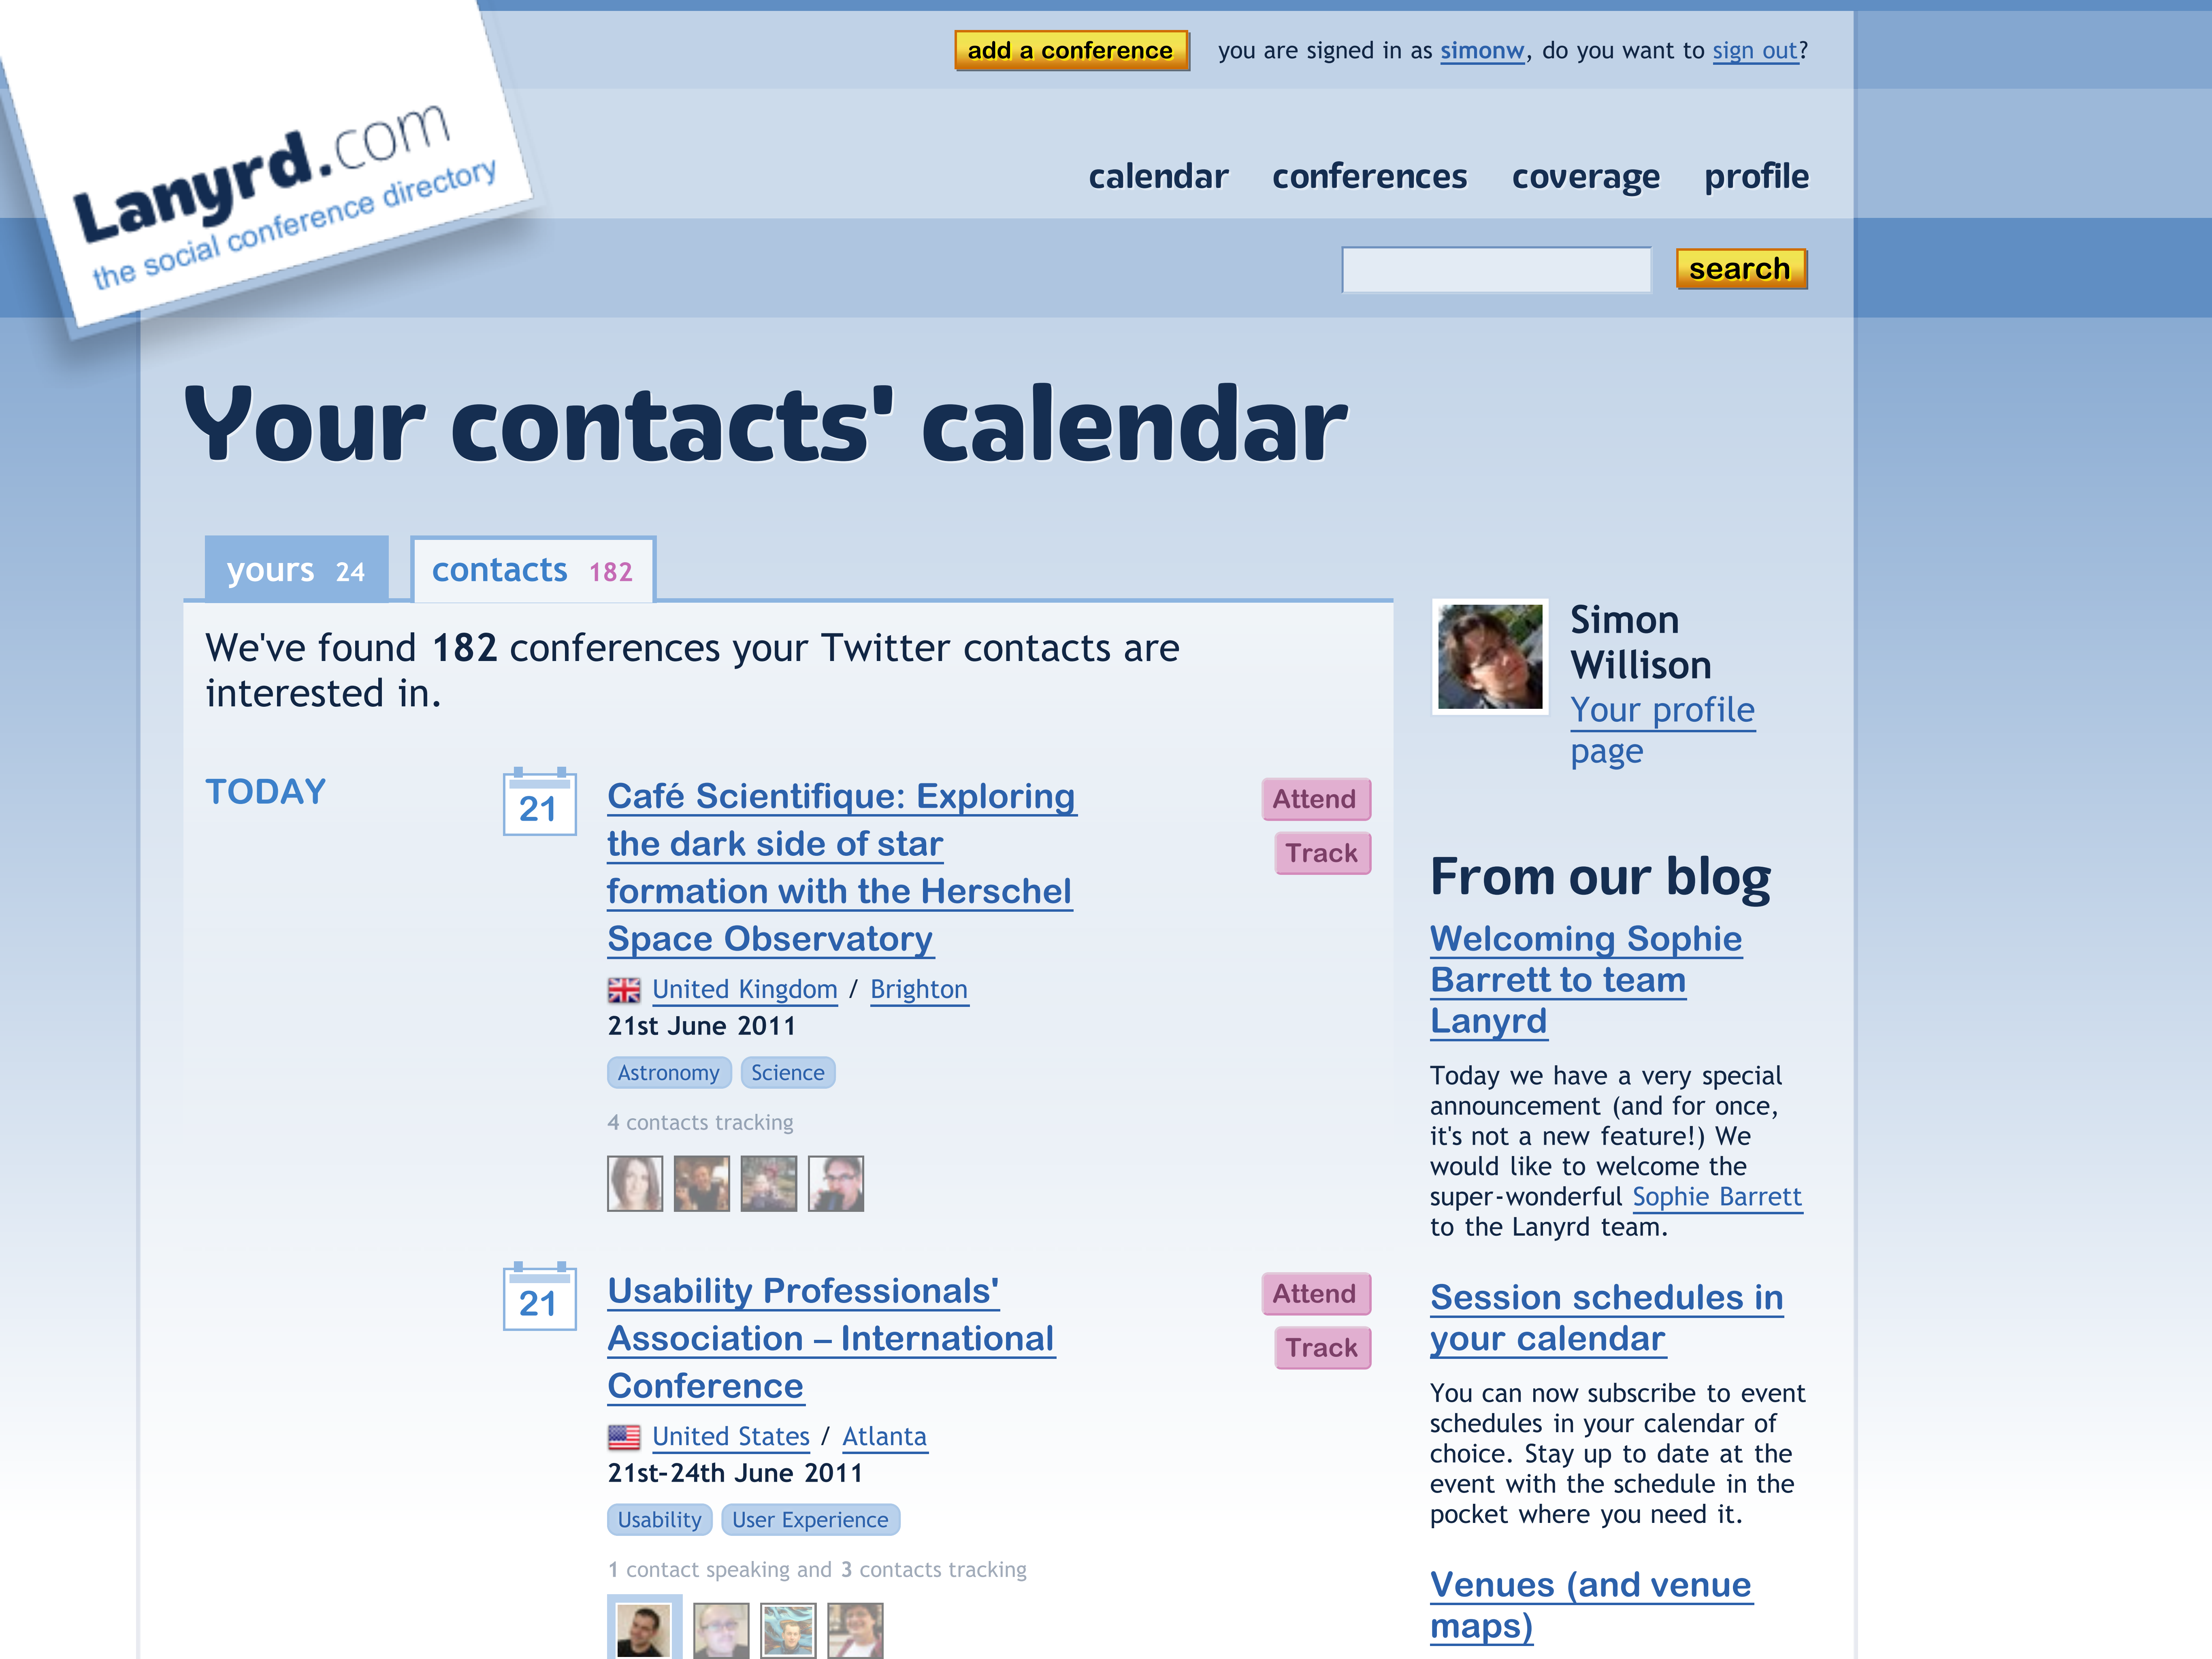

you are signed in as simonw, do you want to sign out?

Your contacts' calendar

yours 24 | contacts 182

We've found **182** conferences your Twitter contacts are interested in.

TODAY

21

Café Scientifique: Exploring the dark side of star formation with the Herschel Space Observatory

United Kingdom / Brighton

21st June 2011

Astronomy | Science

4 contacts tracking

Attend
Track

21

Usability Professionals' Association – International Conference

United States / Atlanta

21st-24th June 2011

Usability | User Experience

1 contact speaking and 3 contacts tracking

Attend
Track

Simon Willison
Your profile page

From our blog

Welcoming Sophie Barrett to team Lanyrd

Today we have a very special announcement (and for once, it's not a new feature!) We would like to welcome the super-wonderful Sophie Barrett to the Lanyrd team.

Session schedules in your calendar

You can now subscribe to event schedules in your calendar of choice. Stay up to date at the event with the schedule in the pocket where you need it.

Venues (and venue maps)

```
sqs = SearchQuerySet()
sqs = sqs.models(Conference)

or_string = ' OR '.join(contact_ids)
sqs = sqs.narrow('attendees:(%s)' % or_string)
```

Staying fresh

- Search engines usually don’t like accepting writes too frequently
  - RealTimeSearchIndex for low traffic sites
- ./manage.py update_index --age=6 (hours)
  - Uses index.get_updated_field()
- Roll your own (message queue or similar...)

Smarter indexing

```
class Article(models.Model):
    needs_indexing = models.BooleanField(
        default = True, db_index = True
    )
    ...
    def save(self, *args, **kwargs):
        self.needs_indexing = True
        super(Article, self).save(*args, **kwargs)
```

```
index = site.get_index(model)
updated_pks = []

objects = index.load_all_queryset().filter(
    needs_indexing=True
)[:100]
if not objects:
    return

for object in objects:
    updated_pks.append(object.pk)
    index.update_object(object)

index.load_all_queryset().filter(
    pk__in = updated_pks
).update(needs_indexing = False)
```

Replication

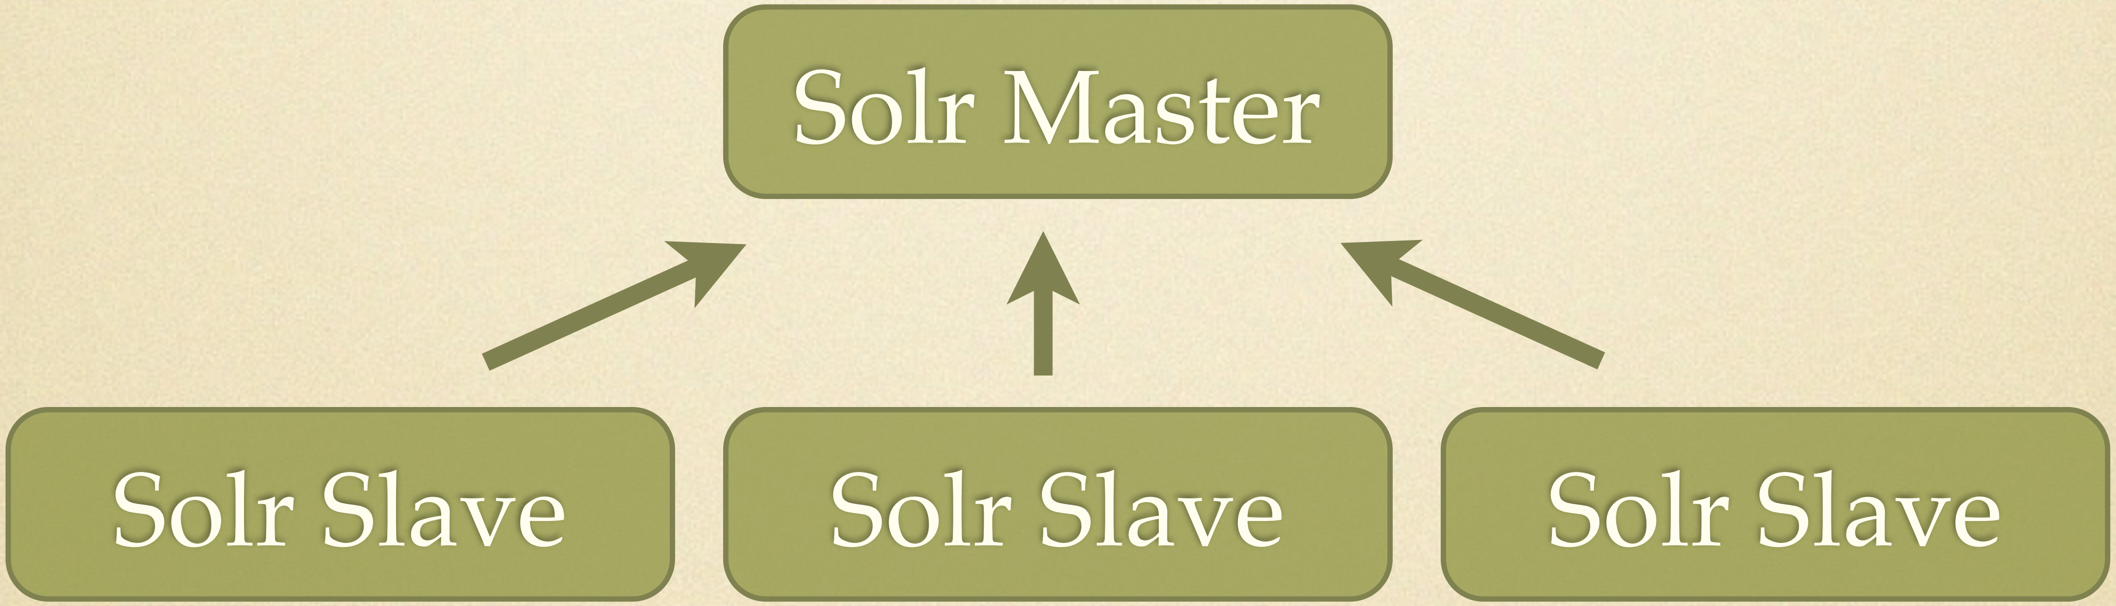

nginx + Solr replication trick

```
upstream solrmaster {
    server 10.68.43.214:8080;
}
upstream solrslaves {
    server 10.68.43.214:8080;
    server 10.193.138.80:8080;
    server 10.204.143.106:8080;
}
```

```
server {
    listen   8983;
    location /solr/update {
        proxy_pass  http://solrmaster;
    }
    location /solr/select {
        proxy_pass  http://solrslaves;
    }
}
```

?

Celery

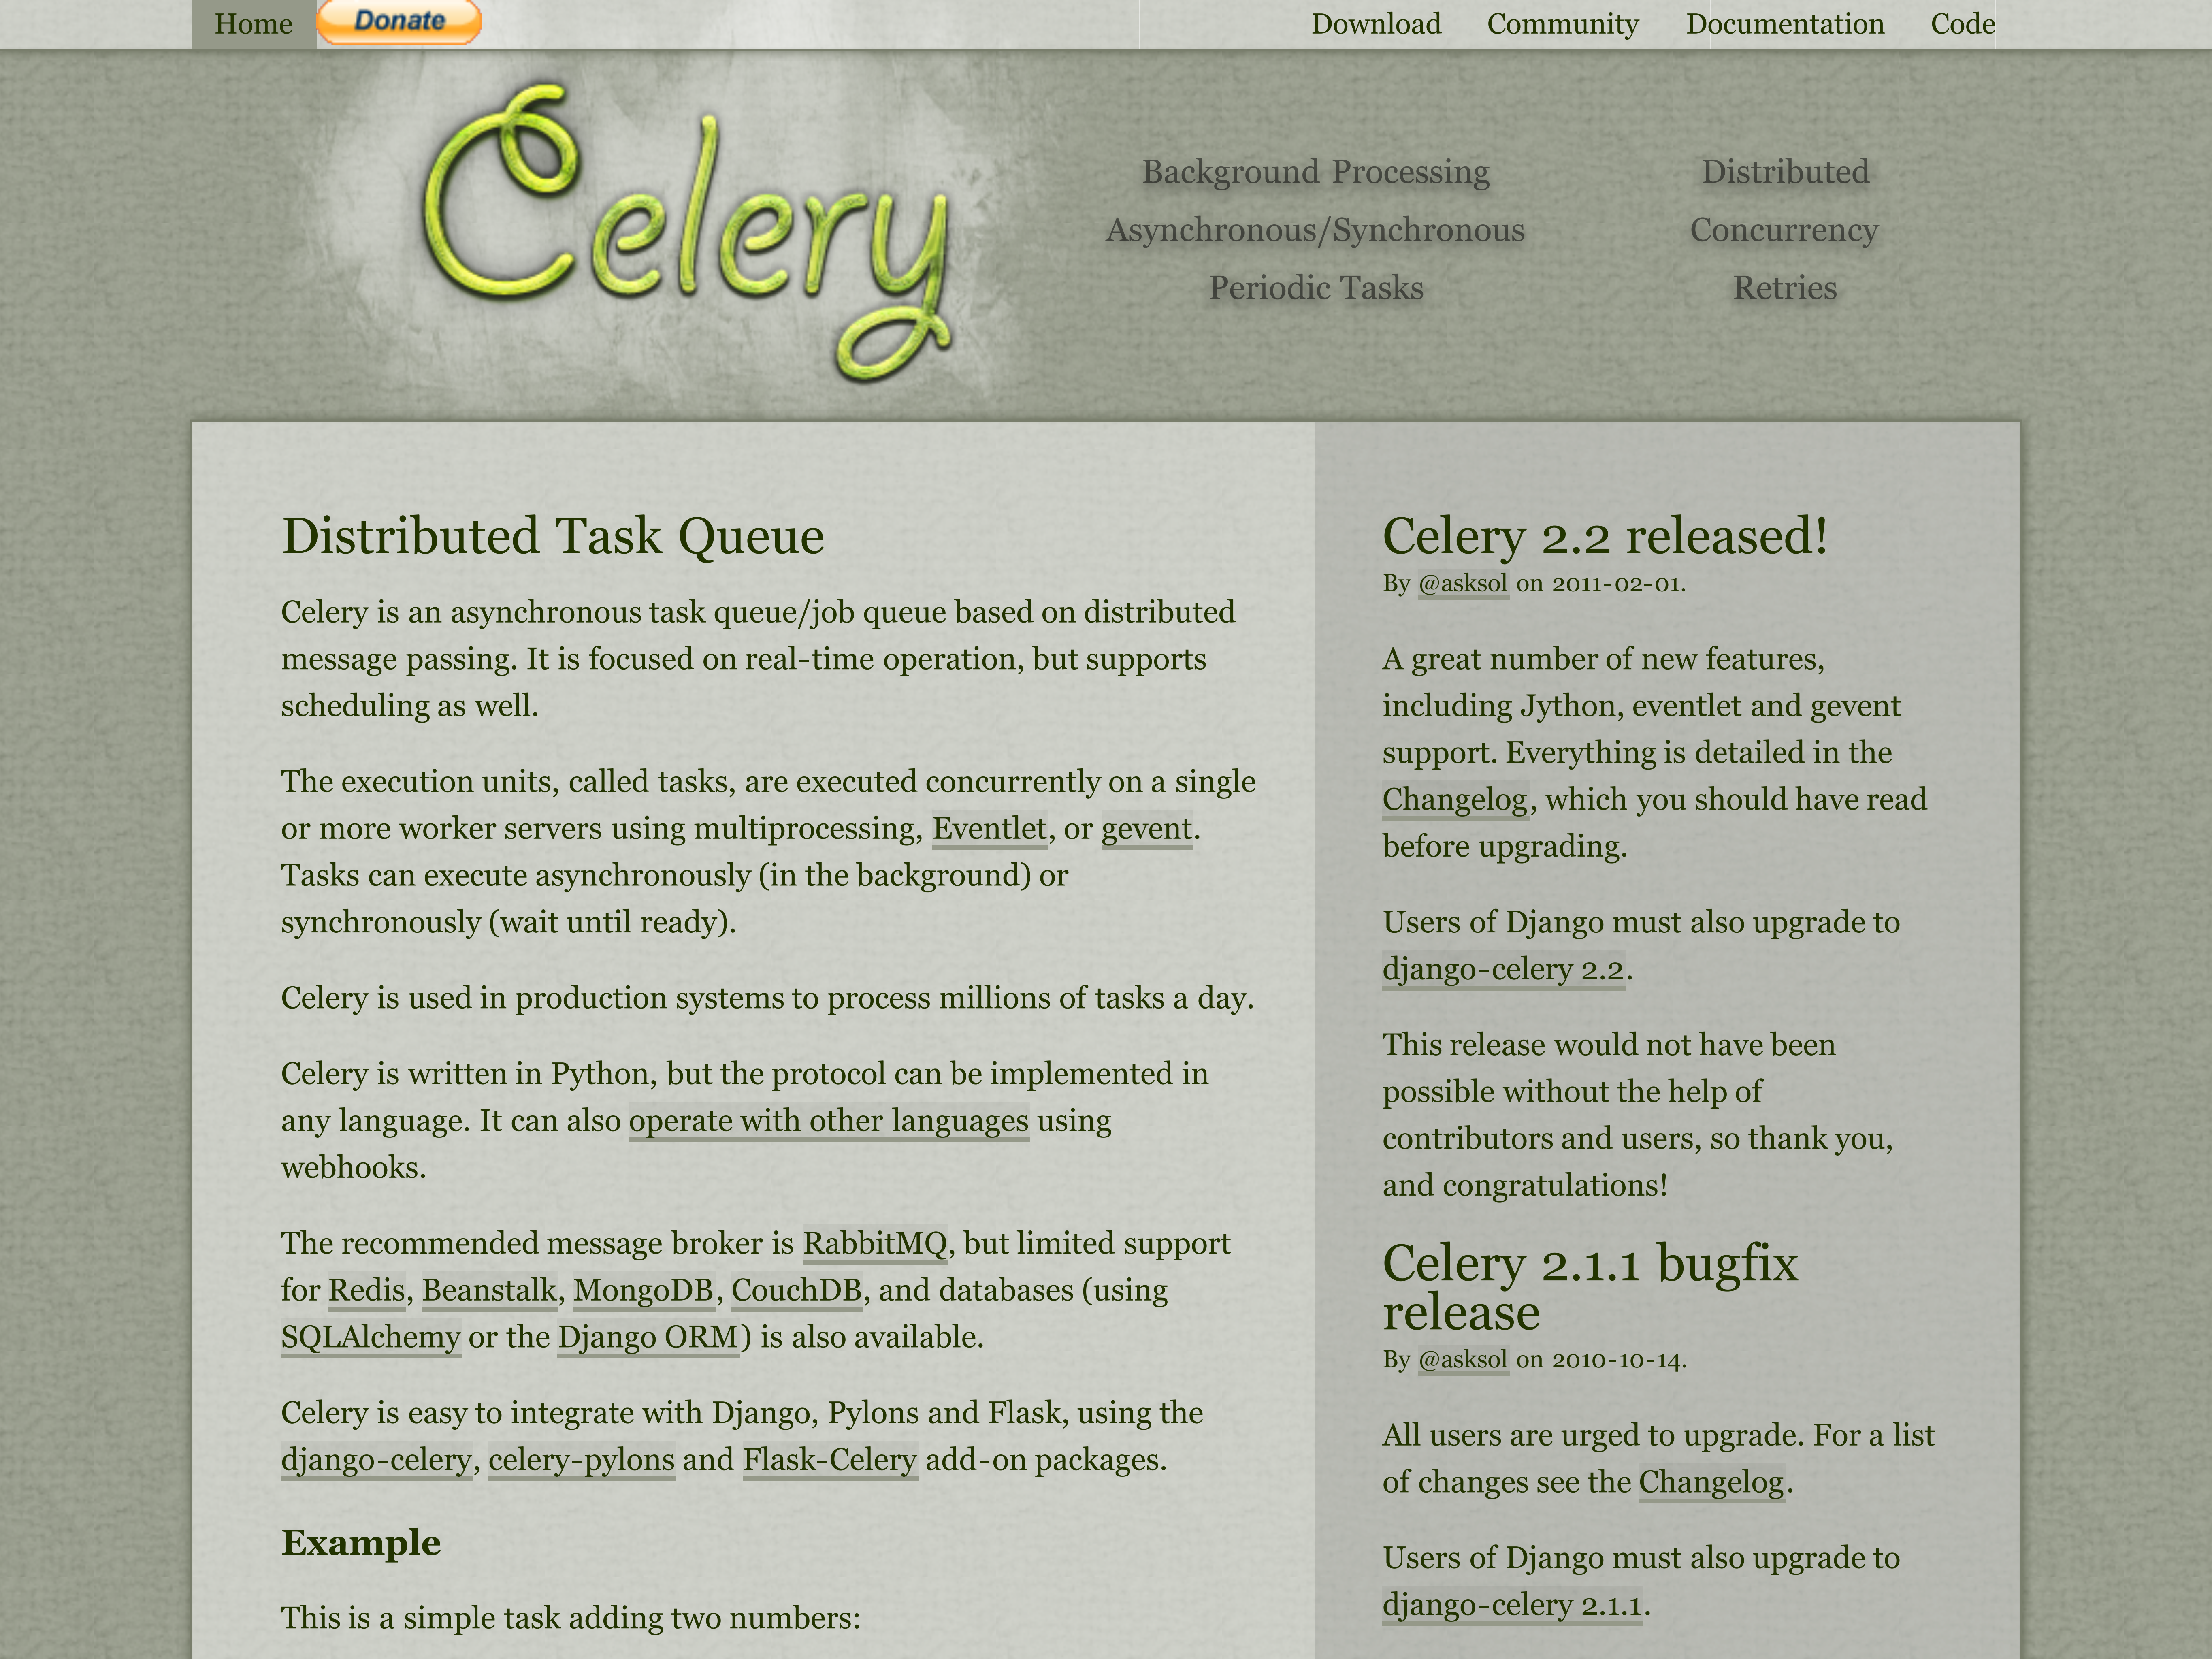

Distributed Task Queue

Celery is an asynchronous task queue/job queue based on distributed message passing. It is focused on real-time operation, but supports scheduling as well.

The execution units, called tasks, are executed concurrently on a single or more worker servers using multiprocessing, Eventlet, or gevent. Tasks can execute asynchronously (in the background) or synchronously (wait until ready).

Celery is used in production systems to process millions of tasks a day.

Celery is written in Python, but the protocol can be implemented in any language. It can also operate with other languages using webhooks.

The recommended message broker is RabbitMQ, but limited support for Redis, Beanstalk, MongoDB, CouchDB, and databases (using SQLAlchemy or the Django ORM) is also available.

Celery is easy to integrate with Django, Pylons and Flask, using the django-celery, celery-pylons and Flask-Celery add-on packages.

Example

This is a simple task adding two numbers:

Celery 2.2 released!

By @asksol on 2011-02-01.

A great number of new features, including Jython, eventlet and gevent support. Everything is detailed in the Changelog, which you should have read before upgrading.

Users of Django must also upgrade to django-celery 2.2.

This release would not have been possible without the help of contributors and users, so thank you, and congratulations!

Celery 2.1.1 bugfix release

By @asksol on 2010-10-14.

All users are urged to upgrade. For a list of changes see the Changelog.

Users of Django must also upgrade to django-celery 2.1.1.

Tasks?

- Anything that takes more than about 200ms
  - Updating a search index
  - Resizing images
  - Hitting external APIs
  - Generating reports

Trivial example

- Fetch the content of a web page

```
from celery.task import task

@task
def fetch_url(url):
    return urllib.urlopen(url).read()

>>> result = fetch_url.delay(‘http://cnn.com/’)
>>> html = result.wait()
```

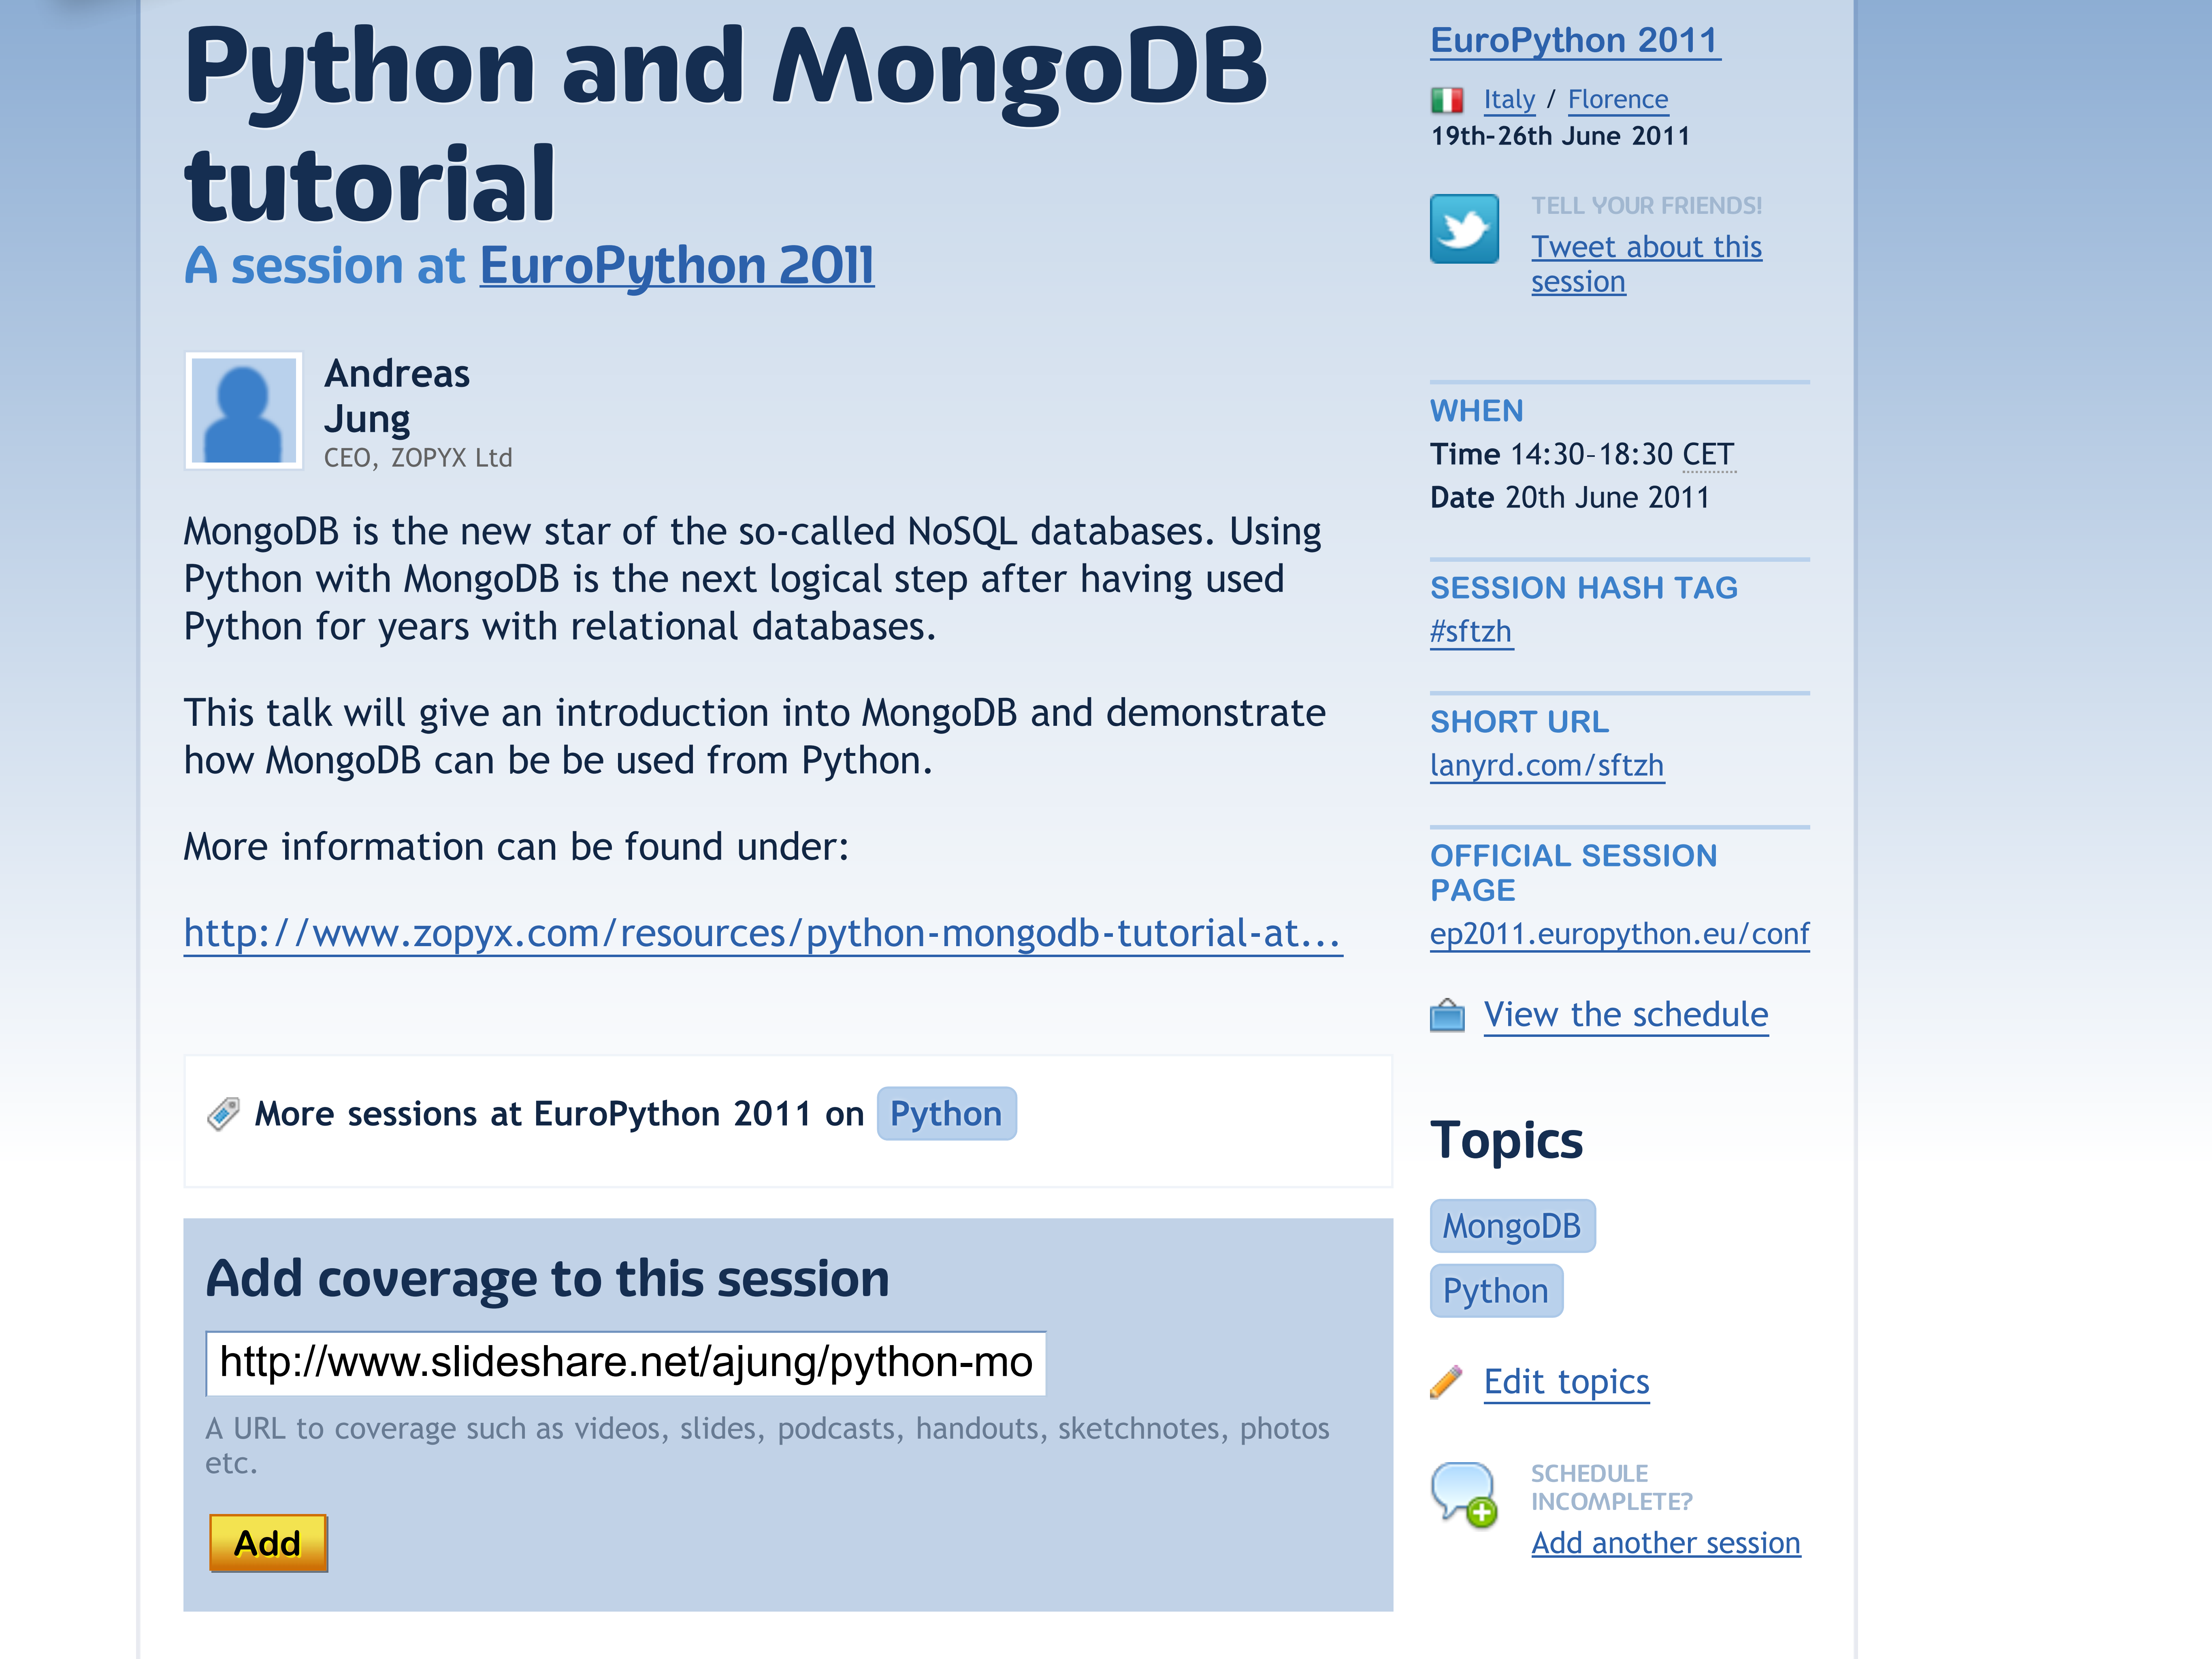

Python and MongoDB tutorial

A session at EuroPython 2011

Andreas Jung
CEO, ZOPYX Ltd

MongoDB is the new star of the so-called NoSQL databases. Using Python with MongoDB is the next logical step after having used Python for years with relational databases.

This talk will give an introduction into MongoDB and demonstrate how MongoDB can be be used from Python.

More information can be found under:

http://www.zopyx.com/resources/python-mongodb-tutorial-at...

More sessions at EuroPython 2011 on Python

Add coverage to this session

http://www.slideshare.net/ajung/python-mo

A URL to coverage such as videos, slides, podcasts, handouts, sketchnotes, photos etc.

Add

EuroPython 2011

Italy / Florence
19th-26th June 2011

TELL YOUR FRIENDS!
Tweet about this session

WHEN
Time 14:30-18:30 CET
Date 20th June 2011

SESSION HASH TAG
#sftzh

SHORT URL
lanyrd.com/sftzh

OFFICIAL SESSION PAGE
ep2011.europython.eu/conf

View the schedule

Topics

MongoDB
Python

Edit topics

SCHEDULE INCOMPLETE?
Add another session

Add coverage

http://www.slideshare.net/ajung/python-mongo-dbtrainingeurop...

Link title

Python mongo db-training-europython-2011

Type of coverage

- ○ Link
- ○ Write-up
- ● Slides
- ○ Video
- ○ Audio
- ○ Sketch notes
- ○ Transcript
- ○ Handout
- ○ Liveblog
- ○ Photos
- ○ Notes

Python and MongoDB tutorial

EuroPython 2011

Italy / Florence
19th-26th June 2011

Coverage preview

From **SlideShare:**

Andreas Jung, www.zopyx.com

PYTHON AND MONGODB
THE PERFECT MATCH

The task itself...

- Tries using http://embed.ly/ to find a preview
- Fetches the HTTP headers and first 2048 bytes
- If HTML, attempts to extract the <title>
- If other, gets the file type and size from headers

Behind the scenes...

```
ar = enhance_link.delay(url)
poll_url = '/working/%s/' % signed.dumps({
    'task_id': ar.task_id,
    'on_done_url': on_done_url,
})
if 'ajax' in request.POST:
    return render_json(request, {
        'ok': True,
        'poll_url': poll_url,
    })
else:
    return HttpResponseRedirect(poll_url)
```

And when it’s done...

```
from celery.backends import default_backend

...
task_id = request.REQUEST.get('id', '')
result = default_backend.get_result(task_id)
```

Configuration

```
# Carrot / Celery: queue uses Redis
CARROT_BACKEND = "ghettoq.taproot.Redis"
BROKER_HOST = " 10.11.11.11" # redis server
BROKER_PORT = 6379
BROKER_VHOST = "6"

# Task results stored in memcached, so they can
# expire automatically
CELERY_RESULT_BACKEND = "cache"
CELERY_CACHE_BACKEND = \
   "memcached://10.11.11.12:11211;..."
```

Advanced Celery

- **celerybeat** for scheduling periodic tasks (a smarter version of cron)
- **celeryev** / **celerymon** for monitoring your worker cluster
- **celerycam** for snapshotting cluster state

The Activity Stream pattern

- How do you implement Twitter?
- Give every user an “inbox” list of message IDs from the people they follow
- Write an ID in to EVERY follower’s inbox when a user tweets
- @timoreilly has 1,473,990 followers
    - redis at 100,000 writes/second = 1.5 seconds

?

Fabric

Fabric

About

Fabric is a Python (2.5 or higher) library and command-line tool for streamlining the use of SSH for application deployment or systems administration tasks.

It provides a basic suite of operations for executing local or remote shell commands (normally or via `sudo`) and uploading/downloading files, as well as auxiliary functionality such as prompting the running user for input, or aborting execution.

Typical use involves creating a Python module containing one or more functions, then executing them via the `fab` command-line tool. Below is a small but complete “fabfile” containing a single task:

```
from fabric.api import run

def host_type():
    run('uname -s')
```

Once a task is defined, it may be run on one or more servers, like so:

```
$ fab -H localhost,linuxbox host_type
[localhost] run: uname -s
[localhost] out: Darwin
[linuxbox] run: uname -s
[linuxbox] out: Linux

Done.
Disconnecting from localhost... done.
Disconnecting from linuxbox... done.
```

In addition to use via the `fab` tool, Fabric’s components may be imported into other Python code,

Table Of Contents

Next topic

Overview and Tutorial

This Page

Show Source

Quick search

Go

Enter search terms or a module, class or function name.

- Turn Python functions in to command line arguments
- High level abstraction over SSH for running commands on remote machines

Automated deployments

- Every project **needs** automated deployments
  - Repeatable, documented, easy to roll back
- Ops people rave about Chef and Puppet
  - Powerful... but hard to set up
- Fabric: take your current procedure, wrap it in a few lines of Python

Simple examples

```
# fabfile.py
from fabric.api import run

def clear_cache():
    run('echo "flush_all" | nc localhost 11211')

$ fab -H memcach1,memcache2 clear_cache
```

Simple examples

```
REDIS_URL = 'http://mirrors.kernel.org/ubuntu/pool/'  + \
    'universe/r/redis/redis-server_2.0.0~rc2-1_amd64.deb'

def bootstrap_redis():
    run('cd /tmp && wget %s' % REDIS_URL)
    sudo('dpkg --force-confnew -i /tmp/redis-server_2.0.0~rc2-1_amd64.deb')
    put('config-files/redis.conf', '/tmp/redis.conf')
    sudo('mv /tmp/redis.conf /etc/redis/redis.conf')
    sudo('/etc/init.d/redis-server restart')
```

Simple examples

```
def git_export():
    env.deploy_date = datetime.datetime.now().strftime(
        '%Y-%m-%d-%H%M%S'
    )
    env.export_path = '/tmp/export/%s' % (env.deploy_date)
    local('mkdir -p %(export_path)s' % env)
    local(
        'cd .. && git archive --prefix=lanyrd/ --format=tar ' +
        'master | tar -x -C %(export_path)s' % env
    )

$ fab git_export
```

Rollback with symlinks

```
def repoint_symlink():
    with settings(warn_only = True):
        run('rm %(deploy_dir)s/previous' % env)
        run('mv %(deploy_dir)s/current %(deploy_dir)s/previous' % env)
    run('ln -s %(deploy_dir)s/%(deploy_date)s %(deploy_dir)s/current' % env)

def rollback():
    run('mv %(deploy_dir)s/current %(deploy_dir)s/_previous' % env)
    run('mv %(deploy_dir)s/previous %(deploy_dir)s/current' % env)
    run('mv %(deploy_dir)s/_previous %(deploy_dir)s/previous' % env)
```

Rollback with symlinks

```
def repoint_symlink():
    with settings(warn_only = True):
        run('rm %(deploy_dir)s/previous' % env)
        run('mv %(deploy_dir)s/current %(deploy_dir)s/previous' % env)
    run('ln -s %(deploy_dir)s/%(deploy_date)s %(deploy_dir)s/current' % env)

def rollback():
    run('mv %(deploy_dir)s/current %(deploy_dir)s/_previous' % env)
    run('mv %(deploy_dir)s/previous %(deploy_dir)s/current' % env)
    run('mv %(deploy_dir)s/_previous %(deploy_dir)s/previous' % env)
```

servers.json

```
{
  "servers": {
    "appserver1": {
      "instance_id": "i-a13432d",
      "public_dns": "ec2-111-11-121-211.compute-1.amazonaws.com",
      "private_dns": "ip-10-195-11-112.ec2.internal",
      "private_ip": "10.195.11.112"
    },
    "appserver2": {
      "instance_id": "i-a34344e",
      "public_dns": "ec2-112-11-121-211.compute-1.amazonaws.com",
      "private_dns": "ip-10-204-111-116.ec2.internal",
      "private_ip": "10.204.111.116"
    },
```

servers.json

```
  ....
  "roles": {
    "appserver": ["appserver1", "appserver2"],
    "solrmaster": ["services3"],
    "solrslave": ["appserver1", "appserver2", "services2"],
    "solrread": ["services3", "appserver1", "appserver2", "services2"],
    "redis": ["services3"],
    "queuebroker": ["services3"],
    "memcached": ["appserver1", "appserver2", "services3"],
  }
}
```

In the fabfile.py

```
_js = json.load(open('servers.json'))
servers = _js['servers']
roles = _js['roles']

def server(name):
    env.hosts = env.hosts or []
    env.hosts.append('ubuntu@%s' % servers[name]['public_dns'])

def role(name):
    for server_name in roles[name]:
        server(server_name)
```

```
$ fab role:memcached clear_cache
```

Dynamic nginx config

```
def deploy_nginx_config():
    def _nginx_backends(role, port):
        return '\n'.join(
            '        server %s:%s;' % (info['private_ip'], port)
            for name, info in servers.items() if name in roles[role]
        )
    content = open('config-files/nginx.conf').read()
    content = content % {
        'backends': _nginx_backends('appserver', 8000),
        'solrmaster': _nginx_backends('solrmaster', 8080),
        'solrslaves': _nginx_backends('solrread', 8080),
    }
    open('/tmp/nginx.conf', 'w').write(content)
    put('/tmp/nginx.conf', '/tmp/nginx.conf')
    sudo('mv /tmp/nginx.conf /etc/nginx/nginx.conf')
```

```
$ fab role:loadbalancer deploy_nginx_config
```

Dream setup

- Web interface to push a git tag to staging
- Big Red Button to push staging to production

?

http://lanyrd.com/sftzq