

The Startup Hangover: Supporting 15 mil users

Phil Calçado – SoundCloud

@pcalcado

<http://philcalcado.com>

ohai!



i'm phil.

i work here:



SOUNDCLOUD



define:analogy



Search

About 35,500,000 results (0.09 seconds)

Everything

Images

Maps

Videos

a·nal·o·gy /əˈnæləjē/

Noun:

1. A comparison between two things, typically on the basis of their structure and for the purpose of explanation or clarification.
2. A correspondence or partial similarity.

Synonyms: similarity - resemblance - likeness - parity - parallel

[More info »](#) [Wikipedia](#) - [Dictionary.com](#) - [Answers.com](#) - [Merriam-Webster](#)

--

law we make up



Customs (Prohibited Imports) Regulations 1956

Statutory Rules 1956 No. 90 as amended

made under the

Customs Act 1901

4W Importation of cat or dog fur

(1) In this regulation:

authorised person means a person authorised in writing by the Minister for this regulation.

cat fur means the pelt or hair of an animal of the species *Felis catus*.

cat or dog fur product means a product or other thing that consists, wholly or partly, of cat fur or dog fur.

dog fur means the pelt or hair of an animal of the species *Canis familiaris*.

(2) The importation into Australia of cat fur, dog fur or a cat or dog fur product is prohibited unless:

(a) permission to import the goods has been granted under subregulation (3) and is in force; and

(b) the permission is produced to the Collector.

(3) The Minister, or an authorised person, may, on application, grant a permission, in writing, for the importation of cat fur, dog fur or a cat or dog fur product.

law we make up

Object-Oriented Programming: An Objective Sense of Style

K. Lieberherr, I. Holland, A. Riel

161 Cullinane Hall, College of Computer Science
Northeastern University, 360 Huntington Ave., Boston MA 02115

We introduce a simple, programming language independent rule (known in-house as the Law of DemeterTM) which encodes the ideas of encapsulation and modularity in an easy to follow form for the object-oriented programmer. You tend to get the

law of nature



law of nature



Greenspuns Tenth Rule Of Programming

This is a humorous observation once made by [Philip Greenspun](#):

Any sufficiently complicated C or Fortran program contains an ad-hoc, informally-specified, bug-ridden, slow implementation of half of [CommonLisp](#).

<http://philip.greenspun.com/research/>

Robert Morris's addendum: "Including Common Lisp."



SOUNDCLOUD

Eric Wahlforss

CTO OF SOUND CLOUD

ABOUT

BOOTSTRAPPING

4 STORIES ABOUT

25 PEOPLE

1 MIO. REGISTERED USERS

SOUND CLOUD

FLICKR FOR AUDIO

Moved from BERLIN GERMANY TO STHLM SWEDEN

MEETING THE SCANDINAVIAN VC GUY

DOWN AT RICKETY OFFICE

NICKING THE FINNISH DESIGN CHAIRS FROM HIS GARAGE

ERIC: "I STILL HAVE THEM!"

BOOTSTRAPPING YOUR OFFICE SPACE

2 THROW A PARTY!

WE HAD A GREAT ROOFTOP

FLEW IN THE MOST INTERESTING SOUND CLOUD USERS

THEY PLAYED SETS + INVITED THEIR FRIEND

PARTY LASTED FOR **2 DAYS!**

APPLICATIONS ARRIVED

OFFICE GOT TRASHED

BOOTSTRAP YOUR COMMUNITY

3 REMEMBER THIS!

CO-FOUNDER ALEX HAD THIS LEATHER JACKET

WE KNEW THIS LASER CUTTING GUY

PUT SOUND CLOUD LOGO ON THE JACKET

WORE THE JACKET EVERY DAY

PEACOCKING IS POWERFUL!

BOOTSTRAP YOUR FUNDING

SoundCloud: share your rocking tunes in a super effective way (invites)

23RD JUNE 2008 by [ERNST-JAN PFAUTH](#)

When you've read the [interview](#) with Media Futurist [Gerd Leonhard](#) a couple of

SoundCloud
super e

23RD JUN

When you've

SoundCloud hits 3 Million users, grows by 50% in 100 days

10TH FEBRUARY 2011 by MARTIN BRYANT

Just the other day we [declared](#) 2011 "The year of social audio", and new figures from [SoundCloud](#) today seem to bear that out.

The audio sharing platform has announced that it has reached 3 million



SoundCloud
super e

23RD JUN

When you've

SoundCloud hits 3 Million users, grows by 50% in 100 days

10TH FEBRUARY 2

Just the other day w
"The year of social a
figures from [SoundC](#)
bear that out.

The audio sharing pl
announced that it ha



[SoundCloud](#), the social audio platform has been gaining steadily in popularity since its launch in 2008. Until this year. In the past few months, the company founded by

SoundCloud
super e

23RD JUN

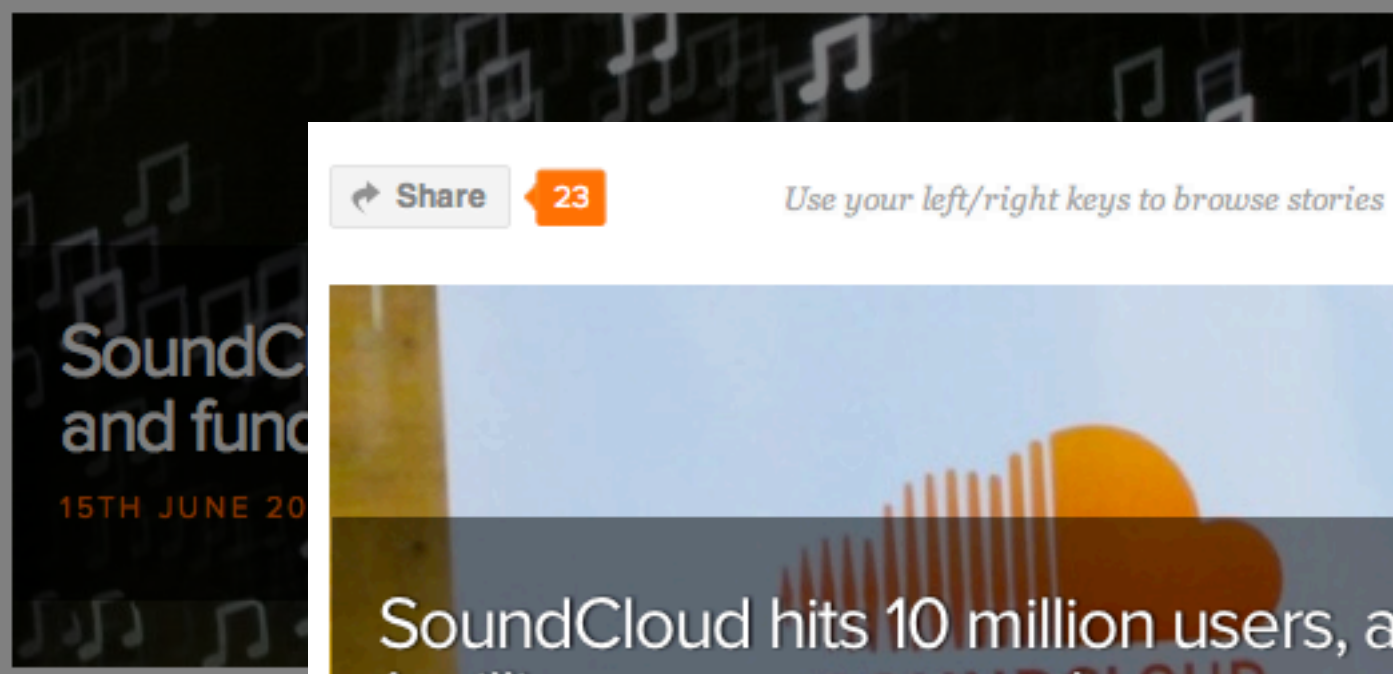
When you've

SoundCloud hits 3 Million users, grows by 50% in 100 days

10TH FEBRUARY 2011

Just the other day we
"The year of social at
figures from [SoundC](#)
bear that out.

The audio sharing pl
announced that it ha



[SoundCloud](#), the s
its launch in 2008



Audio sharing platform [SoundCloud](#) has announced that it has signed up a total of 10 million users, rapid acceleration in takeup for the startup which launched in 2008 and had just 3 million users [in February 2011](#)

SoundCloud
super e

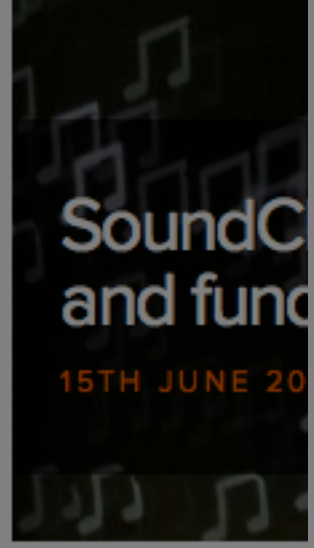

23RD JUN

When you've


SoundCloud hits 3 Million users, grows by 50% in 100 days

Just the other day we
"The year of social a
figures from [SoundC](#)
bear that out.

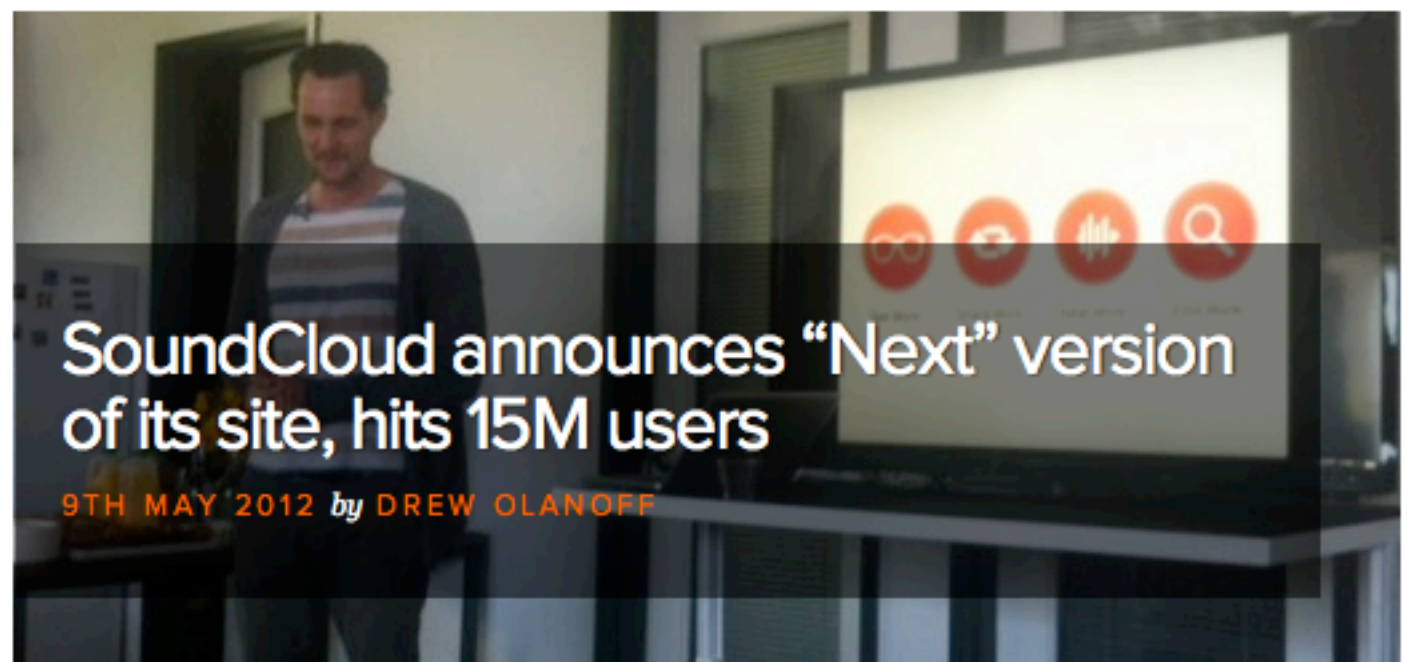
The audio sharing pl
announced that it ha



[SoundCloud](#), the s
its launch in 2008




Audio sharing plat
10 million users, ra
and had just 3 mill




Social audio sharing service [SoundCloud](#) is announcing its "next" version of the

how to survive?


Sunday, 07 January, 2001 08:40:52



Wednesday, 29 October, 2008 17:33:54



Friday, 23 January, 2009 13:49:43




~2008


App



~2011



App



law of nature

How Do Committees Invent?

Melvin E. Conway

Copyright 1968, F. D. Thompson Publications, Inc.
Reprinted by permission of
Datamation magazine,
where it appeared April, 1968.

Conclusion

The basic thesis of this article is that organizations which design systems (in the broad sense used here) are constrained to produce designs which are copies of the communication structures of these organizations. We have seen that this fact has

no point in
fighting nature



Service A



Service B



Service C



teams are built
around goals,
not systems

making layers explicit



Dashboard

You

Tracks

People

Groups

Upload & Share

Search SoundCloud

Dashboard



Your activity **10+** | Incoming tracks **10+** | Exclusive tracks

David Noël uploaded 2 new tracks

Test test

David Noël 1 minute ago

Share Save to Favorites Download

Test

David Noël 4 minutes ago

Share Save to Favorites

David Noël shared this 4 minutes ago: "Test"

Mini Update | More

	This Week
Plays	16
Comments	-
Reposts	-
Tracks	10
Views	68

Write a review

- Go to track
- Save to Favorites
- Download
- Stop following
- Mute David Noël

maryanne favorited a track



Itch by Eyes On 6 minutes ago

a set is a collection of your sounds

Following 443 Followers 307 **10+**

Smashing Pumpkins 12 | 84408

kakkmaddafakka 9 | 662

Netwerk Music Gro 566 | 5711



Phil Calçado

Phil Calçado
Berlin, Germany

17 | 307

Developer @ SoundCloud

- philcalcado.com
- Twitter
- Last.fm
- Songkick

Conifers ↔ Phil Calçado 7 days
Legs And Arms

3.41

▶ 1817 | 13 | 2 | 34

RockstarGames ↔ Phil Calçado 7 days
Emicida - 9 Círculos (Max Payne 3 Original Soundtrack)

3.16

▶ 12506 | 11 | 1 | 39

Phil Calçado 3 7 days
Meio Alternativo

13.19

Find sounds f
New search algo
mean faster sear
people and sound
place.

443 following

- jaycieh**
59 | 1977
- marcohamersma**
29 | 118
- Angelo Pappas**
23 | 77

add other users'
sounds to a set

122 likes

Conifers
Legs And Arms

application layer

vs.

domain layer

law we make up

teams are end-to-end

great for
“lateral hires”
to a team

law we make up


you built it?

you support it

automate yourself
out of support

law we make up

async by default



law we make up

pull requests or pairing

as

code review


github Explore Gist Blog Help pcalcado

PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network **Pull Requests** 12 Issues 18 Wiki Graphs

Open georgi wants someone to merge 3 commits into **master** from **api-user-sounds-query** #774

Discussion Commits 3 Diff 5

 georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned 👤 · No milestone 📅 ·

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by created_at descending and id ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).

Tests:


1. **SUCCESS** 2. **SUCCESS** 3. **SUCCESS** 4. **SUCCESS** 5. **SUCCESS** 6. **SUCCESS**

Open


+ 177 additions

- 200 deletions

[All Pull Requests](#)

3 participants 

 ↔ georgi added some commits 7 days ago

[acb1517](#)  API: Order user sounds by date and id

 ohm started a discussion in the diff 7 days ago

`app/models/stream_item.rb` [View full changes](#)

```

... .. @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34     find_by_sql(sql)
35 35     end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

```

3  ohm [repo collab](#) 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

github Search... Explore Gist Blog Help pcalcado

PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network Pull Requests 12 Issues 18

Open georgi wants someone to merge 3 commits into master from api-user-sounds-query #774

Discussion Commits 3 Diff 5

georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned No milestone

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by created_at descending and id ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).

Tests:
 1. SUCCESS 2. SUCCESS 3. SUCCESS 4. SUCCESS 5. SUCCESS 6. SUCCESS

3 participants

georgi added some commits 7 days ago
 acb1517 API: Order user sounds by date and id

ohm started a discussion in the diff 7 days ago

app/models/stream_item.rb View full changes

```

... @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34     find_by_sql(sql)
35 35     end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

```

3 ohm repo collab 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

short-lived branch

github Explore Gist Blog Help pcalcado


PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network **Pull Requests** 12 Issues 18 Wiki Graphs

Open georgi wants someone to merge 3 commits into **master** from **api-user-sounds-query** #774

Discussion Commits 3 Diff 5

 georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned  No milestone 

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by created_at descending and id ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).


Tests:

1. **SUCCESS** 2. **SUCCESS** 3. **SUCCESS** 4. **SUCCESS** 5. **SUCCESS** 6. **SUCCESS**

Open

+ 177 additions
- 200 deletions

[All Pull Requests](#)

3 participants 

 georgi added some commits 7 days ago

[acb1517](#)  API: Order user sounds by date and id


 ohm started a discussion in the diff 7 days ago

`app/models/stream_item.rb` [View full changes](#)

```

... .. @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34     find_by_sql(sql)
35 35     end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

```

 ohm [repo collab](#) 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

github Search... Explore Gist Blog Help pcalcado

PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network Pull Requests 12 Issues 18 Wiki Graphs

Open georgi wants someone to merge 3 commits into master from api-user-sounds-query #774

Discussion Commits 3 Diff 5

georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned No milestone

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by created_at descending and id ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).

Tests:
 1. SUCCESS 2. SUCCESS 3. SUCCESS 4. SUCCESS 5. SUCCESS 6. SUCCESS

3 participants

georgi added some commits 7 days ago
 acb1517 API: Order user sounds by date and id

ohm started a discussion in the diff 7 days ago

app/models/stream_item.rb View full changes

```


... @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34     find_by_sql(sql)
35 35     end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

```

3 ohm repo collab 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

branch goes green



github Explore Gist Blog Help pcalcado

PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network **Pull Requests** 12 Issues 18 Wiki Graphs

Open georgi wants someone to merge 3 commits into `master` from `api-user-sounds-query` #774

Discussion Commits 3 Diff 5

 georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned 👤 · No milestone 📅 ·

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by `created_at` descending and `id` ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).

Tests:


1. **SUCCESS** 2. **SUCCESS** 3. **SUCCESS** 4. **SUCCESS** 5. **SUCCESS** 6. **SUCCESS**

Open

+ 177 additions

- 200 deletions

[All Pull Requests](#)

3 participants 

 ↔ georgi added some commits 7 days ago

`acb1517`  API: Order user sounds by date and id

 ohm started a discussion in the diff 7 days ago

```

app/models/stream_item.rb View full changes
... .. @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34     find_by_sql(sql)
35 35     end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

```

 ohm **repo collab** 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

github Search... Explore Gist Blog Help pcalcado

PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network Pull Requests 12 Issues 18 Wiki Graphs

Open georgi wants someone to merge 3 commits into master from api-user-sounds-query #774

Discussion Commits 3 Diff 5

georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned No milestone

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by created_at descending and id ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).

Tests:
 1. SUCCESS 2. SUCCESS 3. SUCCESS 4. SUCCESS 5. SUCCESS 6. SUCCESS

3 participants

georgi added some commits 7 days ago
 acb1517 API: Order user sounds by date and id

ohm started a discussion in the diff 7 days ago

app/models/stream_item.rb View full changes

```

... @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34   find_by_sql(sql)
35 35 end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

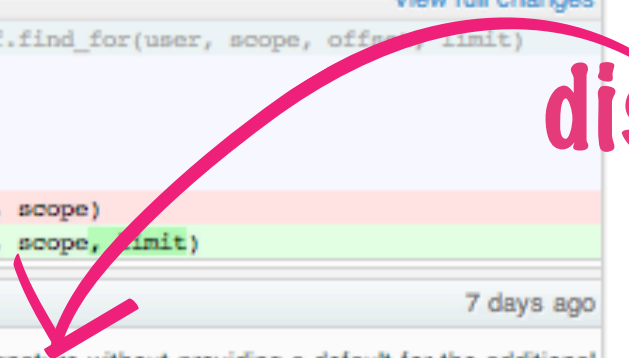
```

3

ohm repo collab 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

discussion




github Explore Gist Blog Help pcalcado

PRIVATE soundcloud / soundcloud Admin Pull Request Unwatch 8 Your Fork 2

Code Network **Pull Requests** 12 Issues 18 Wiki Graphs

Open georgi wants someone to merge 3 commits into `master` from `api-user-sounds-query` #774

Discussion Commits 3 Diff 5

 georgi opened this pull request 7 days ago
API: Order user sounds by date and id

No one is assigned 👤 · No milestone 📅 ·

This changeset fixes the ordering for the user sounds endpoint. It's possible that two items have the same date, so that during pagination the union query will order them differently. To make this stable we order by `created_at` descending and `id` ascending. Furthermore it is probably more efficient to add limits to the sub queries.

@groble: I checked this query with our social ringtones friend and querying the last page wasn't really slower (1500 tracks).

Tests:


1. **SUCCESS** 2. **SUCCESS** 3. **SUCCESS** 4. **SUCCESS** 5. **SUCCESS** 6. **SUCCESS**


Open


+ 177 additions


- 200 deletions

[All Pull Requests](#)

3 participants 

 ↔ georgi added some commits 7 days ago

[acb1517](#)  API: Order user sounds by date and id


 ohm started a discussion in the diff 7 days ago

`app/models/stream_item.rb` [View full changes](#)

```

... .. @@ -26,15 +26,15 @@ def self.find_for(user, scope, offset, limit)
33 33
34 34     find_by_sql(sql)
35 35     end
36 36
37 - def self.tracks_sql(user, scope)
37 + def self.tracks_sql(user, scope, limit)

```

3  ohm [repo collab](#) 7 days ago

You're changing the method signature without providing a default for the additional argument, without changing any of the calls to this method. Am I missing something?

law we make up

use the best tool for the
job, but practice empathy

no_brainer choices:
ruby, go, scala, java, R

our fav outliers:
node.js, clojure,
haskell, python

law we make up

everyone is an architect

bad calls are ok,
gather feedback fast and
make sure we can revert it
when needed

law we make up

end_to_end testing
before checking in code

impersonators / stubs /
mocks
for everything we can

and we often can


```
$ ./precheckin.sh
Buildfile: /Users/pcalcado/code/foss/bananiser/build.xml
```

```
init:
```

```
[delete] Deleting directory /Users/pcalcado/code/foss/bananiser/target
[mkdir] Created dir: /Users/pcalcado/code/foss/bananiser/target
[mkdir] Created dir: /Users/pcalcado/code/foss/bananiser/target/bin
[mkdir] Created dir: /Users/pcalcado/code/foss/bananiser/target/test
```

```
resolve:
```

```
[ivy:retrieve] :: Ivy 2.2.0 - 20100923230623 :: http://ant.apache.org/ivy/ ::
[ivy:retrieve] :: loading settings :: file = /Users/pcalcado/code/foss/bananiser/ivysettings.xml
[ivy:retrieve] :: resolving dependencies :: org.apache#bananiser;working@28cfdaf2bcaa.netpoint.com
```

```
test:
```

```
[javac] Compiling 8 source files to /Users/pcalcado/code/foss/bananiser/target/test
[junit] Running com.soundcloud.bananiser.BananaJobsTest
2012-05-23 09:51:28.132 java[94430:1903] Unable to load realm info from SCDynamicStore
[junit] Tests run: 2, Failures: 0, Errors: 0, Time elapsed: 0.424 sec
[junit] Running com.soundcloud.bananiser.BananiserTest
[junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 0.003 sec
[junit] Running com.soundcloud.bananiser.NoOpReducerTest
[junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 0.265 sec
[junit] Running com.soundcloud.bananiser.SeqFileCompatibleMapperTest
[junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 0.022 sec
[junit] Running com.soundcloud.bananiser.utilities.BananaUtilityTest
[junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 0.09 sec
[junit] Running com.soundcloud.bananiser.utilities.CatUtilityTest
[junit] Tests run: 0, Failures: 0, Errors: 0, Time elapsed: 0.001 sec
[junit] Running com.soundcloud.bananiser.utilities.cat.CatMapperTest
[junit] Tests run: 2, Failures: 0, Errors: 0, Time elapsed: 0.003 sec
```

```
pack:
```

```
[jar] Building jar: /Users/pcalcado/code/foss/bananiser/target/bananiser.jar
```

```
functional-test:
```

```
[echo] ===== Test: Text->Text
[exec] 2012-05-23 09:51:33.999 java[94438:1903] Unable to load realm info from SCDynamicStore
[exec] 12/05/23 09:51:34 INFO security.UserGroupInformation: JAAS Configuration already set up for Hadoop, not re-installing
[exec] 12/05/23 09:51:34 INFO jvm.JvmMetrics: Initializing JVM Metrics with processName=JobTracker, sessionId=
[exec] 12/05/23 09:51:34 WARN util.NativeCodeLoader: Unable to load native-hadoop library for your platform... using built-in
```

law we make up

product manager
is a benevolent dictator

and they often can code

law we make up

20% time is sacred

just don't have the whole
team going away at once

law we make up

no multi-tasking



everyone is a singleton

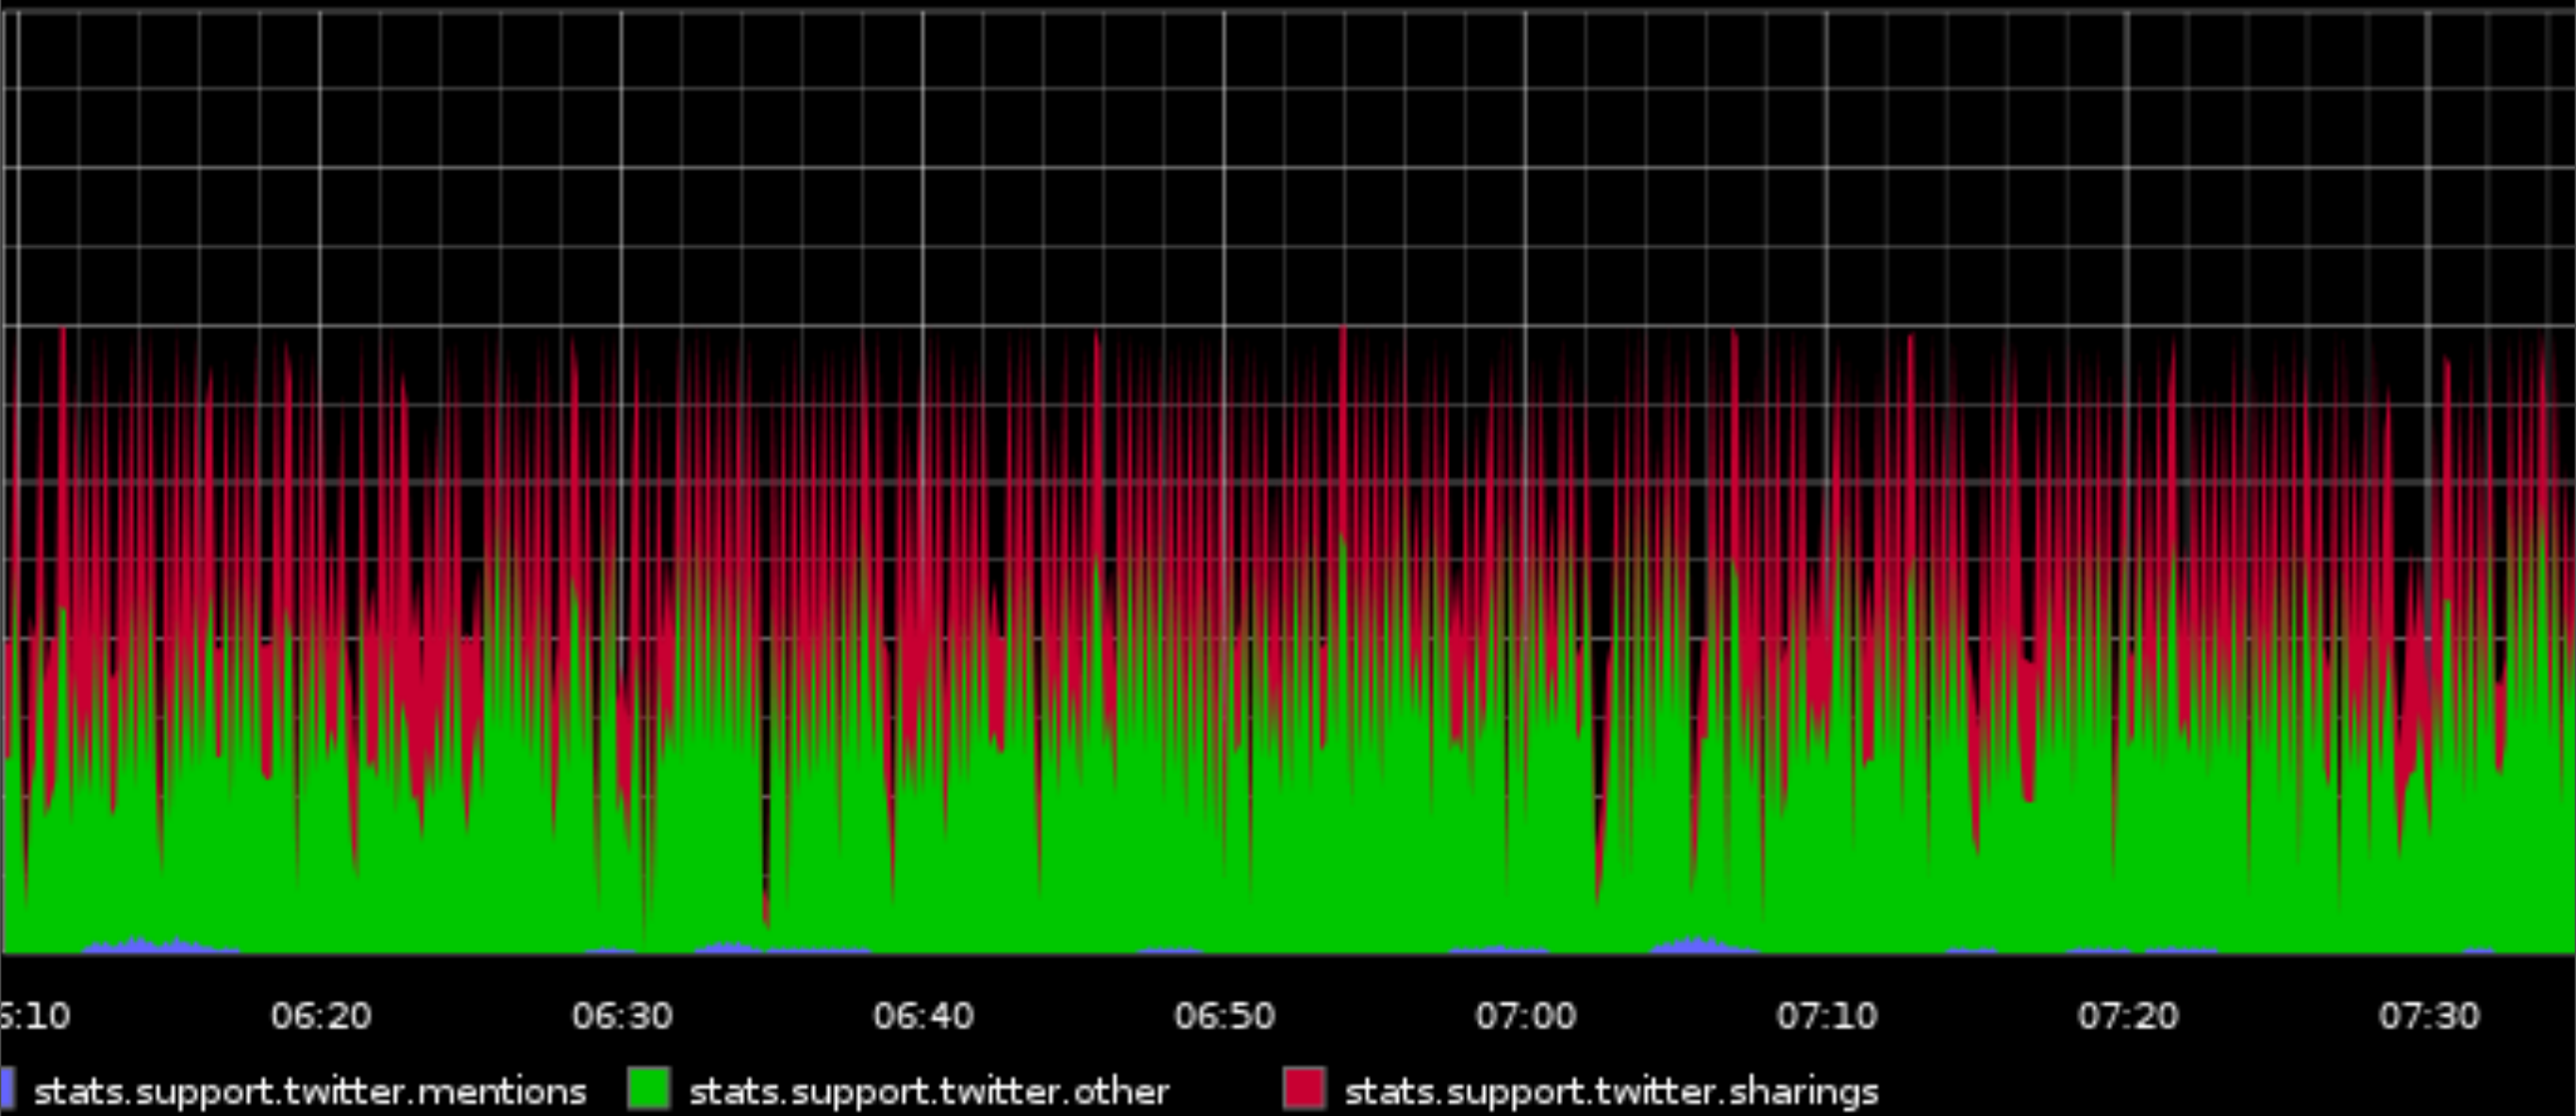
law we make up

**“done” means
someone is using it
(a. k. a. continuous
deployment)**

cycle time: ~ 3 days

law we make up

measure all the things



Qs?

Phil Calçado – SoundCloud

@pcalcado

<http://philcalcado.com>



SOUNDCLOUD

<http://bit.ly/work-at-soundcloud>