

# Use o índice!

O essencial para desenvolvedores

# A culpa não é do Rails

Usando índices para salvar o dia

(no fundo é uma talk sobre banco de dados)

# whoami

Milhouse (Renan Ranelli)

Consultor Independente @ TamingChaos

Senior Software Enginner @ Telnyx LLC

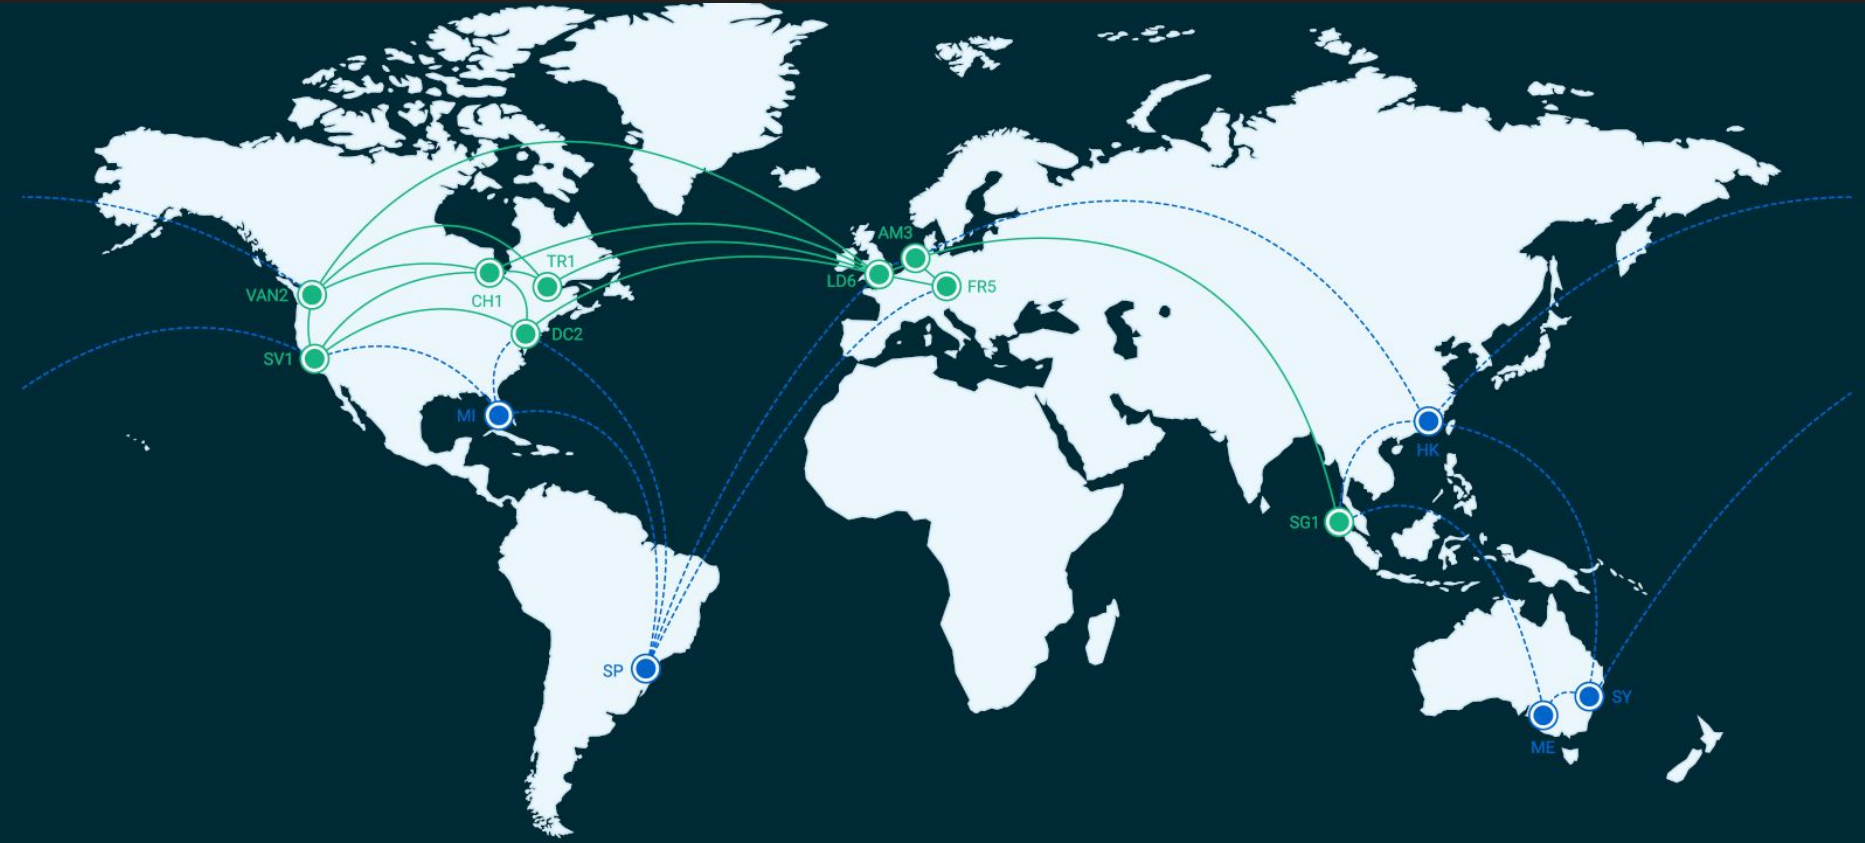
Organizador do meetup de Elixir de São Paulo, o ElugSP

Atualmente: Elixir desde 2015, Python desde 2017 e Ruby antes disso. (E no meio tempo C#, Clojure, JS, Fortran...)

# Telnyx LLC

- A primeira plataforma global de comunicações com rede definida por software

(AWS esta para o datacenter como a Telnyx esta pra telefonia)



POINT OF PRESENCE

FIBER BACKBONE



PLANNED POINT OF PRESENCE

PLANNED FIBER BACKBONE

# Objetivos

- Dar o 1o passo para "subir de nível" na nossa interação com o banco de dados
- Compartilhar a minha experiência a respeito de quais aspectos do banco de dados são mais importantes & menos compreendidos pelos devs
- Tentar transferir um pouco do conhecimento que é meio "lugar comum" na dba-lândia para a developer-lândia
- Não passar vergonha

"X framework é lento"

"X framework é lento"

(Vamos re-escrever em Elixir!)



# Metricas & Observabilidade

É fundamental quando se fala de performance. Conseguir observar seu app rodando é "higiene básica", não é opcional.

Não se começa uma discussão sobre performance sem métricas.

Computadores são notoriamente imprevisíveis.

Mas como estamos na Rails-lândia, a vida é feliz. Estamos a um `gem install` de várias soluções maneiras.

New Relic



Applications

Transactions

Mobile NEW

Servers

Dashboards

F5 LTM NEW

Memcached NEW

MS SQL

MySQL

Nginx NEW

Plugins

Tools

Example S - Master A

Overview

Query Analysis

InnoDB Metrics


Replication

Last 7 days  
Ending now

## Overview

### SQL Volume

This graph shows the total Reads (SELECTs) and Writes (INSERTs,UPDATES,DELETES,REPLACES) executed



PUBLISHED AND SUPPORTED BY



New Relic Inc.

About us | Support site


### Key Utilizations

Percent

InnoDB Buffer Pool Hit Ratio	99%
Query Cache Memory In Use	77.5%
Tmp Tables Written To Disk	41.5%
Query Cache Hit Utilization	36%
Connection Utilization	16.6%


### Database Connections

This graphs shows current database connection information



### Network Traffic

This graph shows the bytes received and set to/from the MySQL instar



### Recent events

All

TODAY

13:00 Alert on Example S - Slave 1  
Details →

Expand all

Collapse all

Duration (ms)	Duration (%)	Segment	Drilldown	Timestamp
218,126		▼ applications/_current_main_charts.html.haml Partial		0.350 s
218,125		▼ applications/_current_main_charts.html.haml Template		0.351 s
218,124		▼ PageHeadersHelper.show_solr_tab?		0.351 s
198,539		▼ find_and_aggregate		0.351 s
160		ApplicationMetricAgent#find_by_sql		0.352 s
18,834		▶ MetricAgent#find_by_sql		1.300 s
353		GC Execution		20.408 s
177,653		▼ .do_find_and_aggregate		20.761 s
0.0		TimesliceTableInfo#find_by_sql		20.762 s
177,602		▼ .fast_time_aggregating_query		20.794 s
170,041		▼ .fast_time_aggregating_query_for_shard		20.795 s
<b>169,231</b>		Timeslice#find		21.060 s
214,102		▼ layouts/application/_main_nav.html.haml Partial		242.198 s
214,102		▼ layouts/application/_main_nav.html.haml Template		242.198 s
192,832		▼ find_and_aggregate		242.199 s
577		▶ ApplicationMetricAgent#find_by_sql		242.199 s
19,460		▶ MetricAgent#find_by_sql		243.546 s
170,565		▼ .do_find_and_aggregate		263.270 s
3.0		TimesliceTableInfo#find_by_sql		263.271 s
170,509		▼ .fast_time_aggregating_query		263.307 s
162,869		▼ .fast_time_aggregating_query_for_shard		263.307 s
543		GC Execution		263.598 s
<b>161,412</b>		Timeslice#find		264.141 s
7,640		▶ .fast_time_aggregating_query_for_shard		426.176 s

New

# New Relic

New Relic. **APM** BROWSER SYNTHETICS MOBILE SERVERS PLUGINS INSIGHTS Alerts<sup>New</sup> Tools Help NewRelic Administration Pro Ann. / Mobile Ent. / L 3

Applications Service maps Key transactions

APPS DataNerd Testbed TIME PICKER Last 30 minutes ending now SERVERS All servers

MONITORING

- Overview
- Service maps
- App map
- Transactions
- Databases**
- External services
- Ruby VMs

EVENTS

- Error analytics<sup>NEW</sup>
- Errors
- Violations
- Deployments
- Thread profiler

REPORTS

- SLA
- Availability
- Capacity
- Scalability
- Web transactions
- Database
- Background jobs

SETTINGS

- Application
- Alert conditions
- Environment
- Class browser

DATABASE

All Memcached MySQL Redis

SORT BY

Most time consuming

Memcached get	863 s
MySQL Account find	378 s
MySQL ClusterAgent find	220 s
MySQL Subscription find	159 s
MySQL Agent find	149 s
MySQL AccountView find	142 s
MySQL AccountProductLevel find	129 s
MySQL ShardSpecification find	126 s
MySQL Alerts::Event find	86.9 s
MySQL SubscriptionProduct find	86.7 s
MySQL ServerAgent find	75.9 s
MySQL AgentSummary find	73.6 s
MySQL RealAgent find	69.3 s
MySQL User find	68.5 s
MySQL Partnership find	55 s
MySQL SamlSsoIntegration find	49 s
MySQL ApplicationSettings find	48 s
MySQL ClusterAgent select	34 s
MySQL Alerts::AccountAlertMigration find	29.7 s
MySQL RealAgent select	26.3 s

Show all database operations table...

Delete all traces

### All databases overview

Top database operations by time consumed

Database query time

Database throughput

Add this chart to the new Dashboards More...

End user 3.08 s 2.46k ppm App server 168 ms 106k rpm 0.0129 err%

# New R

NewRelic APM BROWSER SYNTHETICS MOBILE SERVERS PLUGINS INSIGHTS Alerts New Tools Help NewRelic Administration Pro Ann. / Mobile Ent. / L.

Applications Service maps Key transactions

APPS DataNerd Testbed TIME PICKER Last 30 minutes ending now SERVERS All servers

MONITORING

- Overview
- Service maps
- App map
- Transactions
- Databases**
- External services
- Ruby VMs

EVENTS

- Error analytics NEW
- Errors
- Violations
- Deployments
- Thread profiler

REPORTS

- SLA
- Availability
- Capacity
- Scalability
- Web transactions
- Database
- Background jobs

SETTINGS

- Application
- Alert conditions

DATABASE

All Memcached MySQL Redis

SORT BY Most time consuming

Memcached get	883 s
MySQL Account find	357 s
<b>MySQL ClusterAgent find</b>	<b>227 s</b>
MySQL Agent find	158 s
MySQL Subscription find	150 s
MySQL AccountView find	130 s
MySQL AccountProductLevel find	123 s
MySQL ShardSpecification find	121 s
MySQL Alerts::Event find	84.6 s
MySQL SubscriptionProduct find	82.4 s
MySQL ServerAgent find	75.7 s
MySQL AgentSummary find	75.1 s
MySQL RealAgent find	72.3 s
MySQL User find	65.3 s
MySQL Partnership find	55.7 s
MySQL ApplicationSettings find	48.1 s
MySQL SamlSsoIntegration find	43.1 s
MySQL ClusterAgent select	34.5 s
MySQL Alerts::AccountAlertMigration find	30.5 s
MySQL Deployment find	26.6 s

Show all database operations table...

Delete all traces

### MySQL ClusterAgent find

#### Slow query trace 1,910 ms v2: /applications (GET)

10:19 AM MAX TIME ACTION

**Query**

```
SELECT `agents`.*, true AS use_alert_policies, alertable_violation_summaries.severity AS violation_severity, agent_summaries.problem_ids AS summary_problem_ids, IF(ISNULL(agent_summaries.updated_at), NULL, true) AS summary_reporting, true AS with_health_status FROM `agents` LEFT JOIN cluster_agents_real_agents ON cluster_agents_real_agents.cluster_agent_id = a...
```

**Query analysis (show details)**

Table	Hint
agents	<ul style="list-style-type: none"><li>The table was retrieved with this index: <code>unique_index_for_identity</code></li><li>A temporary table was created to access this part of the query, which can cause poor performance. This typically happens if the query contains GROUP BY and ORDER BY clauses that list columns differently.</li><li>MySQL had to do an extra pass to retrieve the rows in sorted order, which is a cause of poor performance but sometimes unavoidable.</li><li>You can speed up this query by querying only fields that are within the index. Or you can create an index that includes every field in your query, including the primary key.</li><li>Approximately 119 rows of this table were scanned.</li></ul>
cluster_agents_real_agents	<ul style="list-style-type: none"><li>The table was retrieved with this index: <code>PRIMARY</code></li><li>Approximately 13 rows of this table were scanned.</li></ul>
real_agents	<ul style="list-style-type: none"><li>The table was retrieved with this index: <code>PRIMARY</code></li><li>You can speed up this query by querying only fields that are within the index. Or you can create an index that includes every field in your query, including the primary key.</li><li>Approximately 1 row of this table was scanned.</li></ul>
agent_summaries	<ul style="list-style-type: none"><li>The table was retrieved with this index: <code>PRIMARY</code></li><li>You can speed up this query by querying only fields that are within the index. Or you can create an index that includes every field in your query, including the primary key.</li><li>Approximately 1 row of this table was scanned.</li></ul>
alertable_violation_summaries	<ul style="list-style-type: none"><li>The table was retrieved with this index: <code>unique_index_on_alertable_violation_summaries</code></li><li>You can speed up this query by querying only fields that are within the index. Or you can create an index that includes every field in your query, including the primary key.</li></ul>

Environment

End user 2.97 s 2.6k rpm | App server 173 ms 109k rpm 0.0138 errk



# Transaction trace

Track as Key Transaction



Delete this trace

Dec 9, '16 8:38 pm    4,890 ms    301 ms (6.17%)    77.3 ms (1.58%)  
 TRACE TIME            RESP. TIME            USER CPU BURN            SYSTEM CPU BURN

Summary    Trace details    **Database queries**

Show fast queries (< 5ms)

Total duration	Call count	Query
4,340 ms	7	SELECT fc.*FROM field_config fcWHERE (fc.field_name = :db_condition_placeholder?) AND (fc.active = :db_condition_placeholder?) AND (fc.storage_active = :db_condition_placeholder?) AND (fc.deleted = :db_condition_placeholder?)
98 ms	3	DELETE FROM variable WHERE (name = :db_condition_placeholder?)
10 ms	1	INSERT INTO semaphore (name, value, expire) VALUES (:db_insert_placeholder?, :db_insert_placeholder?, :db_insert_placeholder?)

### Query details



Duration 4,344 ms

### Stack trace

```

in PDOStatement::execute called at /mnt/www/html/newrelic/docroot/includes/database/database.inc (2171)
in DatabaseStatementBase::execute called at /mnt/www/html/newrelic/docroot/includes/database/database.inc (683)
in DatabaseTransaction::query called at /mnt/www/html/newrelic/docroot/includes/database/database.inc (26)

```

# Transaction trace

Track as Key Transaction



Delete this trace

Dec 9, '16 8:38 pm    4,890 ms    301 ms (6.17%)    77.3 ms (1.58%)  
TRACE TIME            RESP. TIME            USER CPU BURN            SYSTEM CPU BURN

Summary    Trace details    **Database queries**

Show fast queries (< 5ms)

Total duration	Call count	Query
<b>4,340 ms</b>	7	SELECT fc.*FROM field_config fcWHERE (fc.field_name = :db_condition_placeholder_?) AND (fc.active = :db_condition_placeholder_?) AND (fc.storage_active = :db_condition_placeholder_?) AND (fc.deleted = :db_condition_placeholder_?)
98 ms	3	DELETE FROM variable WHERE (name = :db_condition_placeholder_?)
10 ms	1	INSERT INTO semaphore (name, value, expire) VALUES (:db_insert_placeholder_?, :db_insert_placeholder_?, :db_insert_placeholder_?)

### Query details

Duration: 4,344 ms

### Stack trace

```
in PDOStatement::execute called at /mnt/www/html/newrelic/docroot/includes/database/database.inc (2171)
in DatabaseStatementBase::execute called at /mnt/www/html/newrelic/docroot/includes/database/database.inc (683)
in DatabaseTransaction::query called at /mnt/www/html/newrelic/docroot/includes/database/database.inc (26)
```



# (Lembrando que existem alternativas ao NR)

- AppSignal
- AppOptics
- Pingdom Server Monitor (formerly Scout App)
- Stackify Retrace
- DynaTrace (formerly Ruxit)
- Atatus
- Datadog
- LogicMonitor
- AppDynamics
- Cloudwatch & todas as parafernalias AWS
- .... etc etc etc etc

# Ok. Mas agora como eu resolvo a query lenta?

Pra conseguir raciocinar sobre o que vimos, precisamos entender sobre como o banco de dados funciona.

Vamos focar em apenas duas coisas no banco de dados: **Índices** e **Tabelas**.

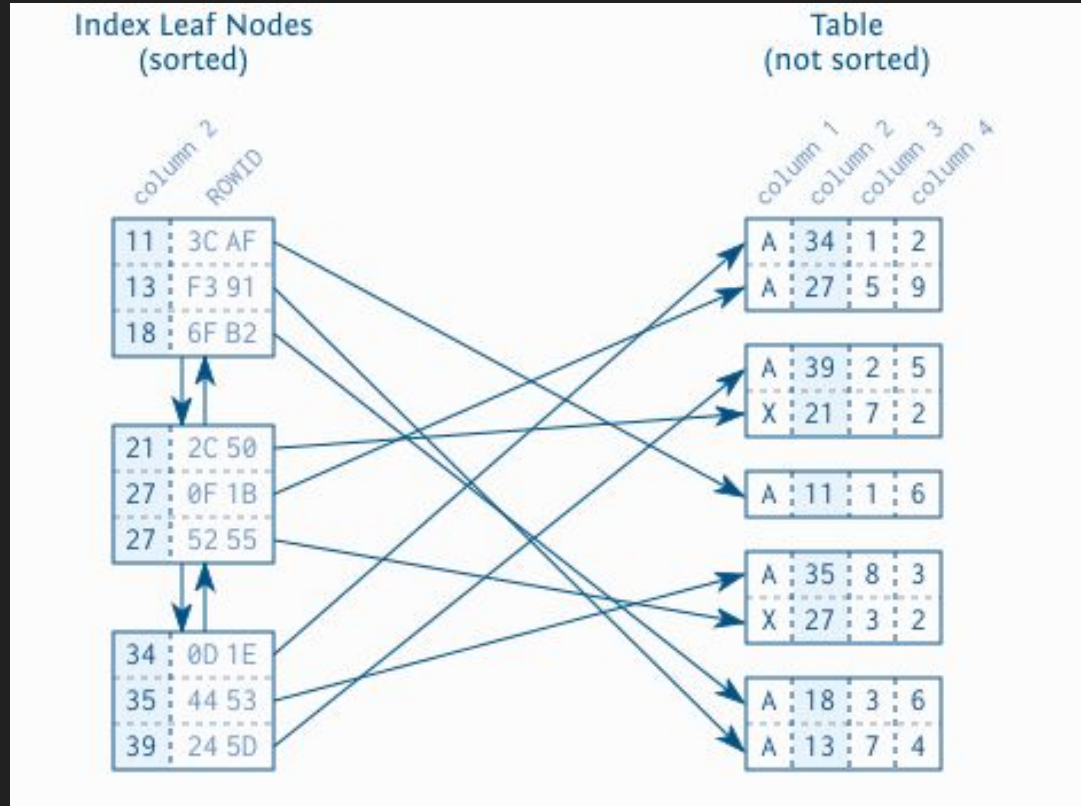
- Vamos imaginar a "**tabela**" como um arquivo em que sempre escrevemos os dados sequencialmente. (pensa no `>>` do shell)
- Índices são um pouco mais envolvidos e vamos falar em seguida

Use o indice, luke!

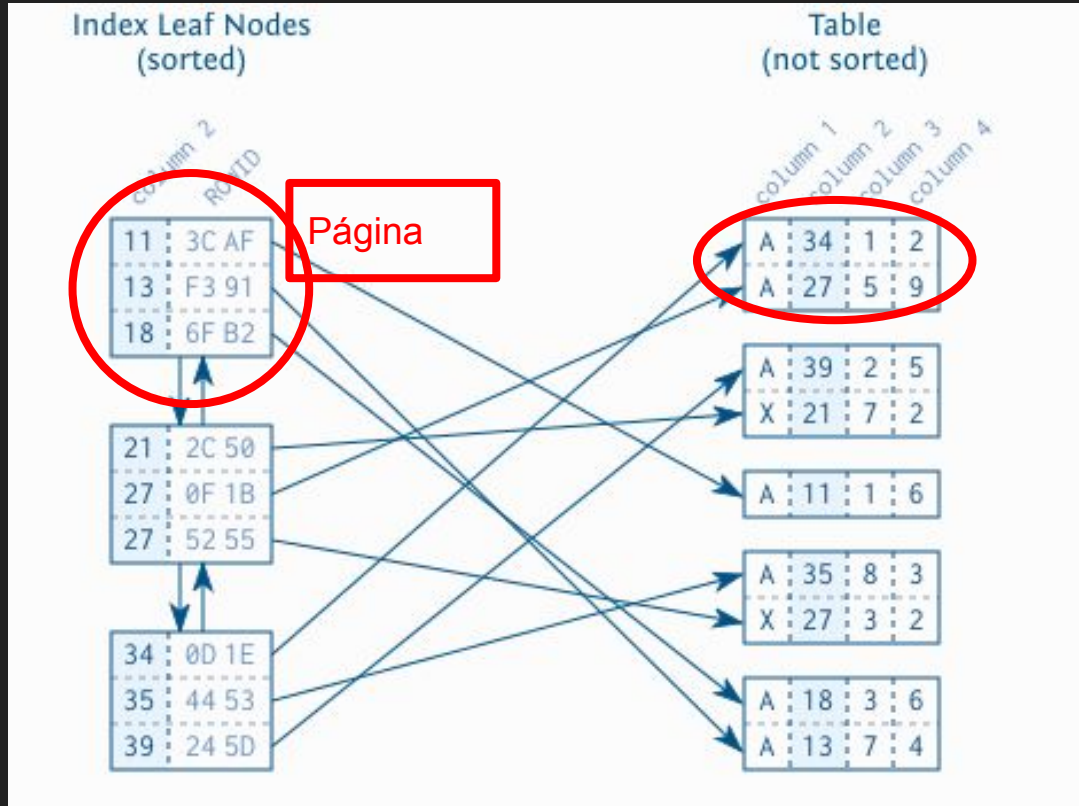
<https://use-the-index-luke.com/>



# A anatomia do índice



# A anatomia do índice



# A principal coisa mal-compreendida

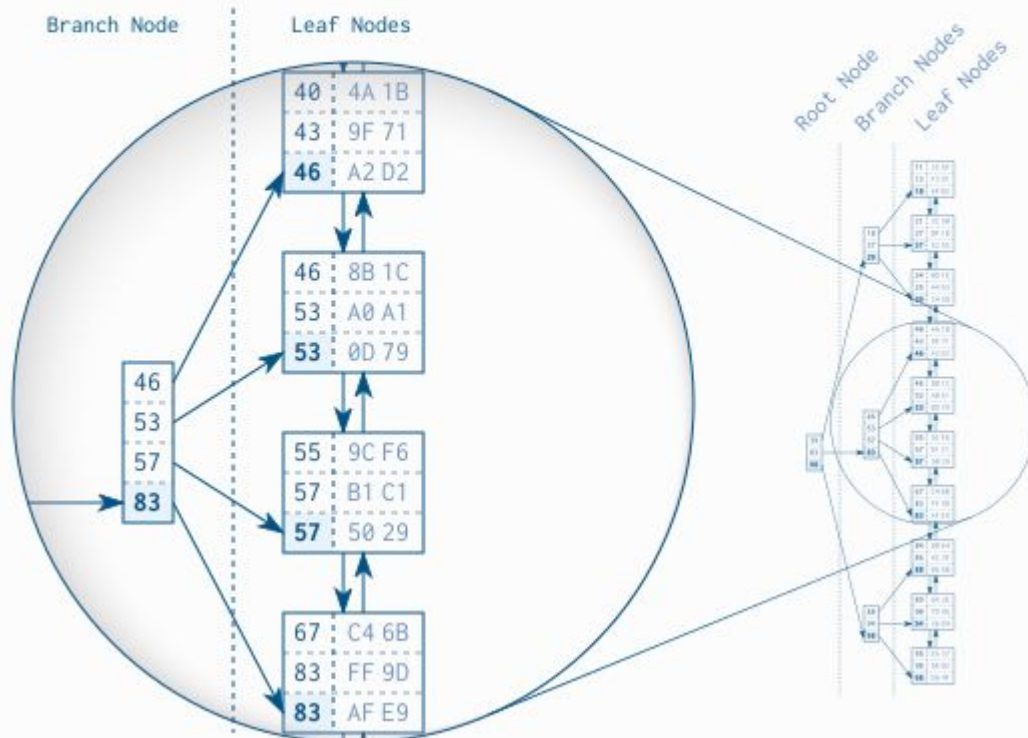
Quando vamos recuperar dados do banco, só podemos recuperar \*páginas\* inteiras, e não registros.

Se você precisa de um registro, você precisa ler toda a \*página\* que o contém. Não há escapatória.

**PÁGINAS LIDAS >>> NUMERO DE REGISTROS**

A anatomy

Figure 1.2 B-tree Structure



Ta, mas como usa

Temos essa estrutura doida de ponteiros. Como que isso ajuda?



Ta, mas como usa

Temos essa estrutura doida de p que isso ajuda?

Table  
(not sorted)

	column 1	column 2	column 3	column 4
A	34	1	2	
A	27	5	9	
A	39	2	5	
X	21	7	2	
A	11	1	6	
A	35	8	3	
X	27	3	2	
A	18	3	6	
A	13	7	4	

Ta, mas como usa

Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2


A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4

# A anatomia do índice

Buscando: id==57



# A anatomia do índice

Buscando: id==55

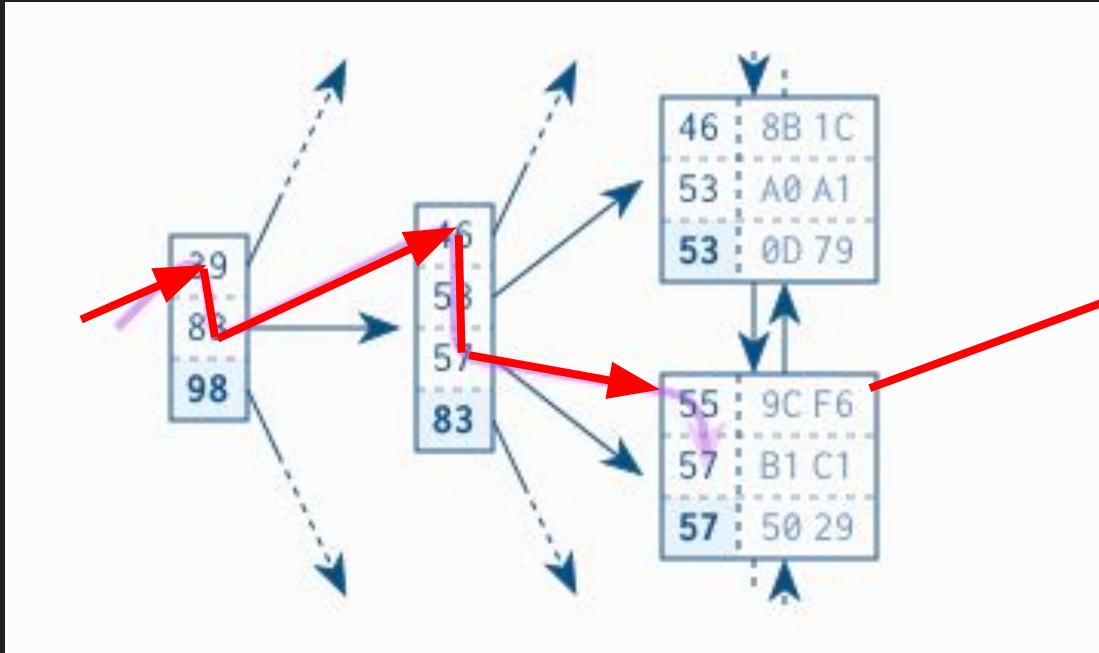


Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2

A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4

Animações pra quem gosta:

<https://www.cs.usfca.edu/~galles/visualization/BPlusTree.html>

# Range Scan

```
#+begin_src sql :results code  
SELECT first_name, last_name, subsidiary_id, phone_number  
  FROM employees  
  WHERE subsidiary_id >= 40 and subsidiary_id <= 55  
#+end_src
```

# Range Scan

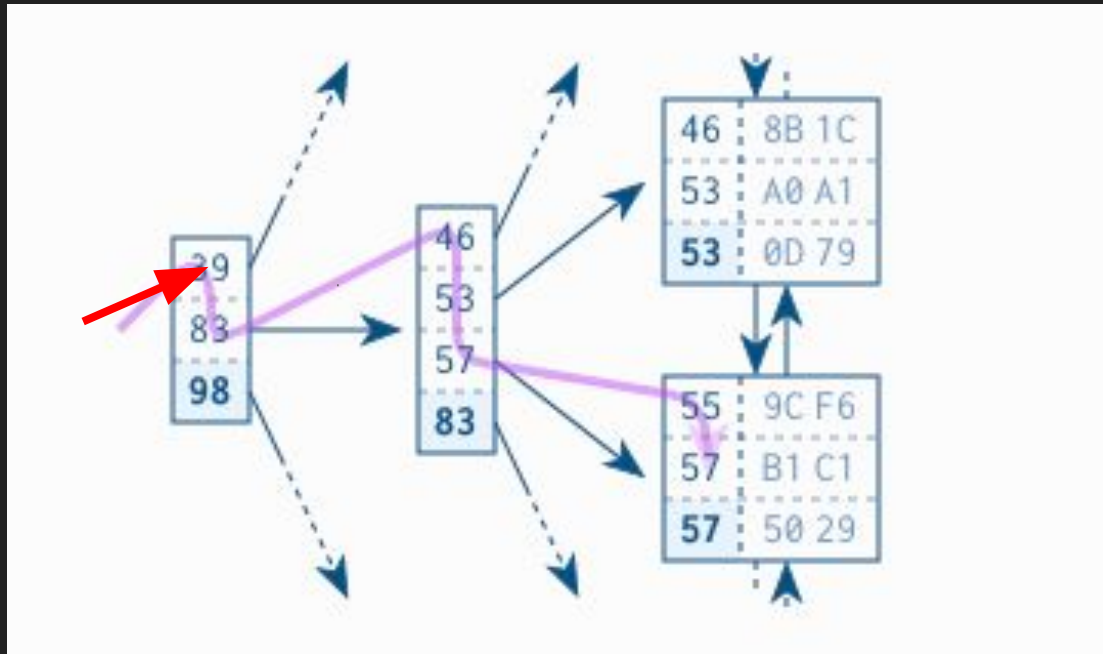


Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2

A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4

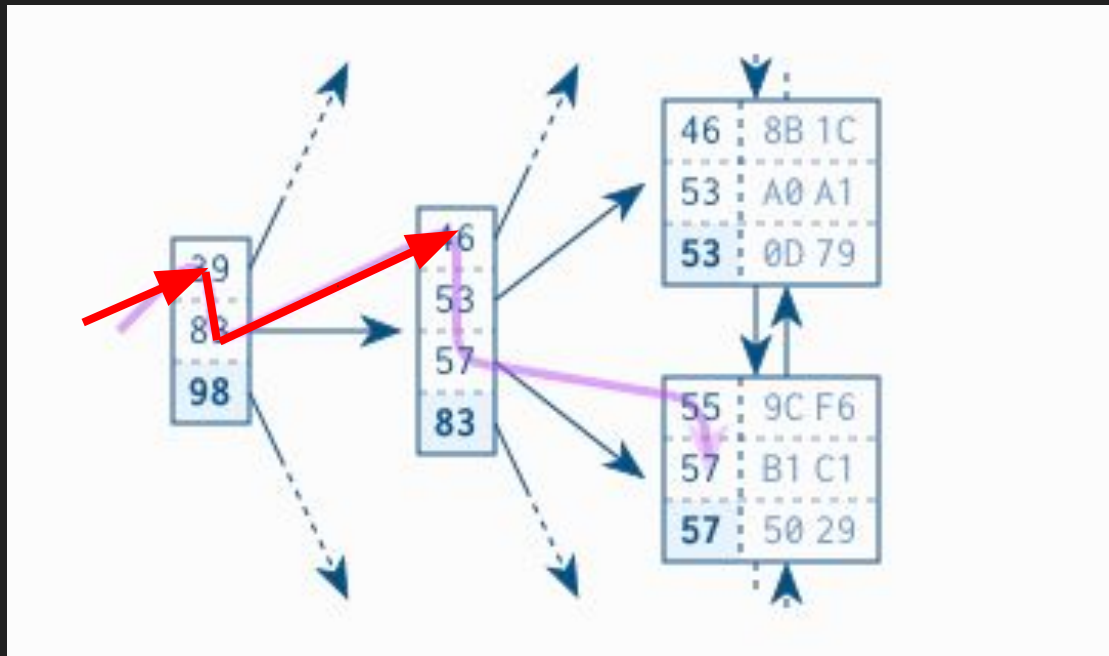


Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2

A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4



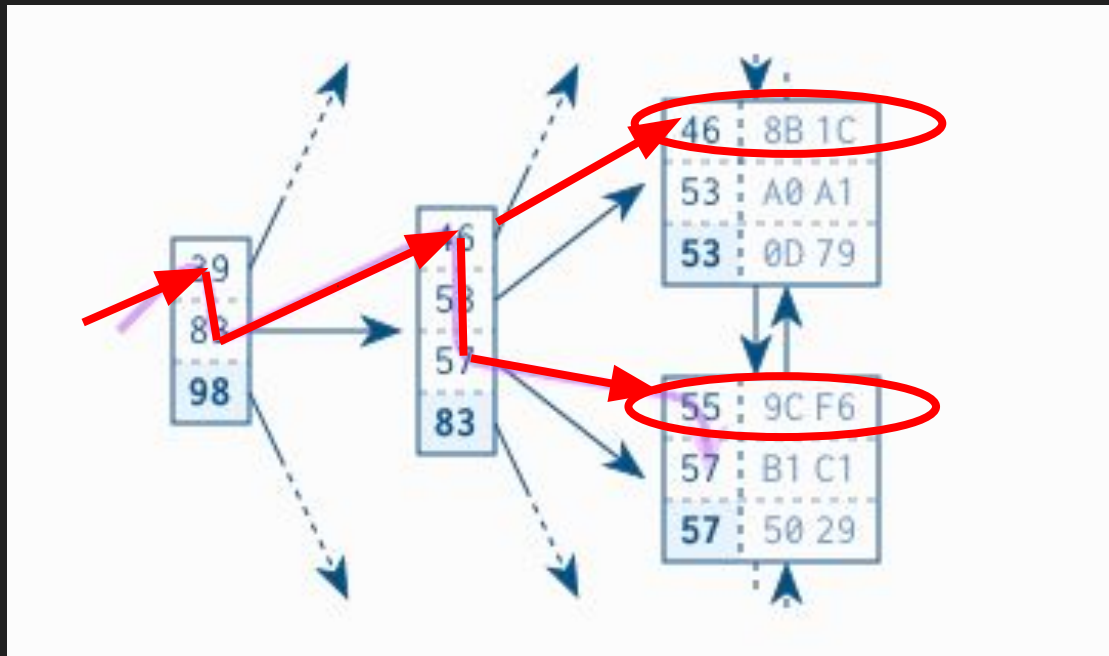


Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2

A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4

# Range Scan




Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2

A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4

# Range Scan




Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

A	34	1	2
A	27	5	9

A	39	2	5
X	21	7	2

A	11	1	6
---	----	---	---

A	35	8	3
X	27	3	2

A	18	3	6
A	13	7	4

```
#+begin_src sql :results code
```

```
SELECT first_name, last_name, subsidiary_id, phone_number  
FROM employees  
WHERE last_name = 'WINAND'  
AND subsidiary_id = 30
```

```
#+end_src
```

```
#+begin_src
```

```
-----  
|Id |Operation          | Name          | Rows | Cost |  
-----  
| 0 |SELECT STATEMENT  |               |    1 |   30 |  
|*1 | TABLE ACCESS BY INDEX ROWID | EMPLOYEES     |    1 |   30 |  
|*2 | INDEX RANGE SCAN  | EMPLOYEES_PK  |   40 |    2 |  
-----
```

Predicate Information (identified by operation id):

- ```
-----  
1 - filter("LAST_NAME"='WINAND')  
2 - access("SUBSIDIARY_ID"=30)
```

```
#+end_src
```

Top

Se você tem um banco chave-valor, isso já é o suficiente.

Então é só jogar índice em tudo e já era?

Não é bem assim...

Caso degenerado: tabela esparça

# Caso degenerado: tabela esparça

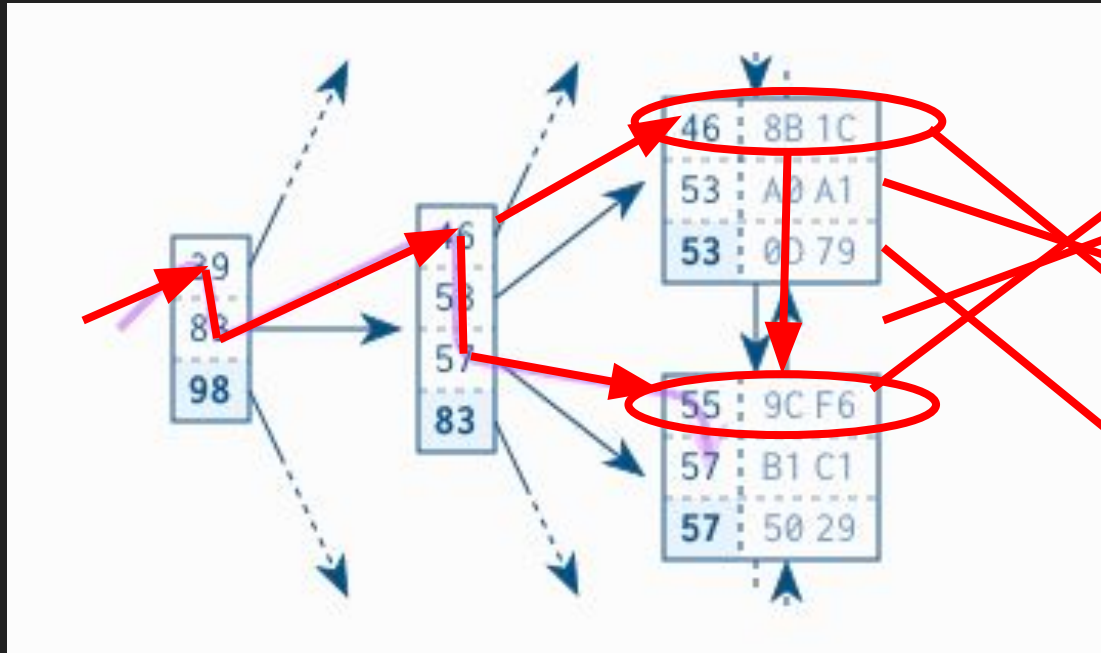


Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

|   |    |   |   |
|---|----|---|---|
| A | 34 | 1 | 2 |
| A | 27 | 5 | 9 |
| A | 39 | 2 | 5 |
| X | 21 | 7 | 2 |
| A | 11 | 1 | 6 |
| A | 35 | 8 | 3 |
| X | 27 | 3 | 2 |
| A | 18 | 3 | 6 |
| A | 13 | 7 | 4 |

# Caso degenerado: tabela esparsa

```
#+begin_src
```

```
Total rows: 100000
```

```
-----  
| Id | Operation          | Name          | Rows  | Cost |  
-----  
|  0 | SELECT STATEMENT   |               | 100000 | 477 |  
|* 1 | TABLE ACCESS FULL| EMPLOYEES     |    100 | 477 |  
-----
```

```
Predicate Information (identified by operation id):
```

```
-----  
  1 - filter("LAST_NAME"='WINAND' AND "SUBSIDIARY_ID"=30)
```

```
#+end_src
```



## Bottomline

O índice não faz milagre

Se a engine de consulta do banco não tiver certeza que o uso do índice vai reduzir o número de leituras, ele não será usado. (ESTATÍSTICAS!)

# Como resolve?

1. "Clusterizar" a tabela
2. Evitar tocar a tabela por completo (covering index)
3. Diminua o result-set
4. Diminua o número de registros por página  
(desperdiça espaço)
5. Pede ajuda pro DBA

# Clustered-table




Table  
(not sorted)

|   | column 1 | column 2 | column 3 | column 4 |
|---|----------|----------|----------|----------|
| A | 34       | 1        | 2        |          |
| A | 27       | 5        | 9        |          |
| A | 39       | 2        | 5        |          |
| X | 21       | 7        | 2        |          |
| A | 11       | 1        | 6        |          |
| A | 35       | 8        | 3        |          |
| X | 27       | 3        | 2        |          |
| A | 18       | 3        | 6        |          |
| A | 13       | 7        | 4        |          |

# Clustered-table




Table  
(not sorted)

|   | column 1 | column 2 | column 3 | column 4 |
|---|----------|----------|----------|----------|
| A | 34       | 1        | 2        |          |
| A | 27       | 5        | 9        |          |
| A | 39       | 2        | 5        |          |
| X | 21       | 7        | 2        |          |
| A | 11       | 1        | 6        |          |
| A | 35       | 8        | 3        |          |
| X | 27       | 3        | 2        |          |
| A | 18       | 3        | 6        |          |
| A | 13       | 7        | 4        |          |

# Index-only scan

```
#+begin_src sql :results code  
CREATE INDEX sales_sub_eur  
  ON sales ( subsidiary_id, eur_value )  
SELECT SUM(eur_value) FROM sales  
  WHERE subsidiary_id = ?  
#+end_src
```

```
#+begin_src
```

| ----- |                  |               |       |      |
|-------|------------------|---------------|-------|------|
| Id    | Operation        | Name          | Rows  | Cost |
| ----- |                  |               |       |      |
| 0     | SELECT STATEMENT |               | 1     | 104  |
| 1     | SORT AGGREGATE   |               | 1     |      |
| * 2   | INDEX RANGE SCAN | SALES_SUB_EUR | 40388 | 104  |
| ----- |                  |               |       |      |

```
#+end_src
```

# Clustered-table




Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

|   |    |   |   |
|---|----|---|---|
| A | 34 | 1 | 2 |
| A | 27 | 5 | 9 |

|   |    |   |   |
|---|----|---|---|
| A | 39 | 2 | 5 |
| X | 21 | 7 | 2 |

|   |    |   |   |
|---|----|---|---|
| A | 11 | 1 | 6 |
|---|----|---|---|

|   |    |   |   |
|---|----|---|---|
| A | 35 | 8 | 3 |
| X | 27 | 3 | 2 |

|   |    |   |   |
|---|----|---|---|
| A | 18 | 3 | 6 |
| A | 13 | 7 | 4 |

# Clustered-table




Table  
(not sorted)

column 1  
column 2  
column 3  
column 4

|   |    |   |   |
|---|----|---|---|
| A | 34 | 1 | 2 |
| A | 27 | 5 | 9 |

|   |    |   |   |
|---|----|---|---|
| A | 39 | 2 | 5 |
| X | 21 | 7 | 2 |

|   |    |   |   |
|---|----|---|---|
| A | 11 | 1 | 6 |
|---|----|---|---|

|   |    |   |   |
|---|----|---|---|
| A | 35 | 8 | 3 |
| X | 27 | 3 | 2 |

|   |    |   |   |
|---|----|---|---|
| A | 18 | 3 | 6 |
| A | 13 | 7 | 4 |

A performance depende não só da sua consulta, mas também da \*disposição dos dados no disco\*, i.e. a ordem em que foram escritos na tabela (!!)

(ou seja, tem coisas que código apenas não resolve)



E não é só isso...

Tudo o que falamos também é condicionado às "estatísticas" de uso, e é radicalmente influenciado por caches em vários níveis e afins.

A engine do banco é \*muito\* esperta.

Tudo o que hipotetizamos precisa ser \*verificado\*

Tudo o que falamos também é influenciado por:

- Qual banco vc ta usando (postgresql, mysql, oracle, sqlserver...)
- Configuração do banco
- Hardware
- Sistema de arquivos e separação de volumes
- Hardware
- Rede
- etc

# Se vc quiser ficar manjeiro



<https://www.amazon.com.br/PostgreSQL-High-Performance-Ibrar-Ahmed/dp/1784392979>

# Discussão

Tá. Falamos isso tudo que provavelmente a maior parte da plateia já sabia.

Mas por que isso é importante?????

# Conclusão

Falar de performance é difícil.

Existem diversas dimensões que precisamos avaliar sobre o que \*performance\* significa. Latência? Throughput? Consumo de memória? CPU? Variância de cada um desses? Dependência desses com escalabilidade horizontal/vertical, etc, etc.

Até um sistema web simples (e.g. `backend` + `banco` + `workers` + `jquery`), daqueles que cabe no heroku já possuem várias peças que influenciam a experiência total do usuário.

Dizer que "O framework " é lento e partir pra uma reescrita sem entender essas coisas e sem métricas, além de não resolver o problema, é um jeito muito eficiente de queimar dinheiro e neurônio

Obrigado

(dscip qualquer coisa)

# Contatos

Milhouse (Renan Ranelli)

[milhouseonsoftware.com](http://milhouseonsoftware.com)

[@renanranelli](https://twitter.com/renanranelli) @ Twitter

[/rranelli](https://github.com/rranelli) @ Github