



The MySQL Ecosystem @ GitHub




Sam Lambert

Director of Technology

 github.com/samlambert

 twitter.com/isamlambert

 samlambert.com



*THE HUMAN BRAIN IS A WONDERFUL
THING. IT STARTS WORKING THE MOMENT
YOU ARE BORN, AND NEVER STOPS UNTIL
YOU STAND UP TO SPEAK IN PUBLIC.*

- GEORGE JESSEL





GitHub

Code hosting and collaboration

What is GitHub?




8
MILLION

DEVELOPERS ARE
WORKING
TOGETHER ON

21
MILLION

PROJECTS ON
GITHUB RIGHT NOW



What is GitHub?




130+ TB
Git data




239
GitHubbers



100
Engineers



Slave



Replica/Secondary



What is MySQL's role at
GitHub?

What is MySQL's role at GitHub?



janky

jenkins mysql redis elasticsearch memcached

pages

nginx

github-fs

git-daemon ernicorn gitmon slumlord

analytics-app

nginx unicorn

gpanel

unicorn resqued timerd

github-babel

babeld gitauth

github-fe

nginx unicorn

alambic-fe

alambic

github-api

nginx unicorn



What is the Ecosystem?

What is the Ecosystem?





Why is the Ecosystem so important?



The GitHub Stack

The GitHub Stack



WEB



Ruby on Rails

Powerful application framework.
Rapid development.



Javascript

UI goodness

BACKEND



C

Powers a number of backend
applications



Puppet

Provisioning



Git

Obviously

DATABASES



MySQL

Main source of truth




Elasticsearch

Indexes all the code, as well as
issues and pull requests.



Redis

Resque, cache data.



```
SELECT DATE_SUB(NOW(),  
INTERVAL 2 YEAR);
```



Small number of MySQL
hosts



Majority of queries served
from one host



Replicas only used for
backups and failover



Old hardware




Unscalable




Contention everywhere



Traffic spikes caused query response times to go up



Need to move data centers



A chance to update
hardware



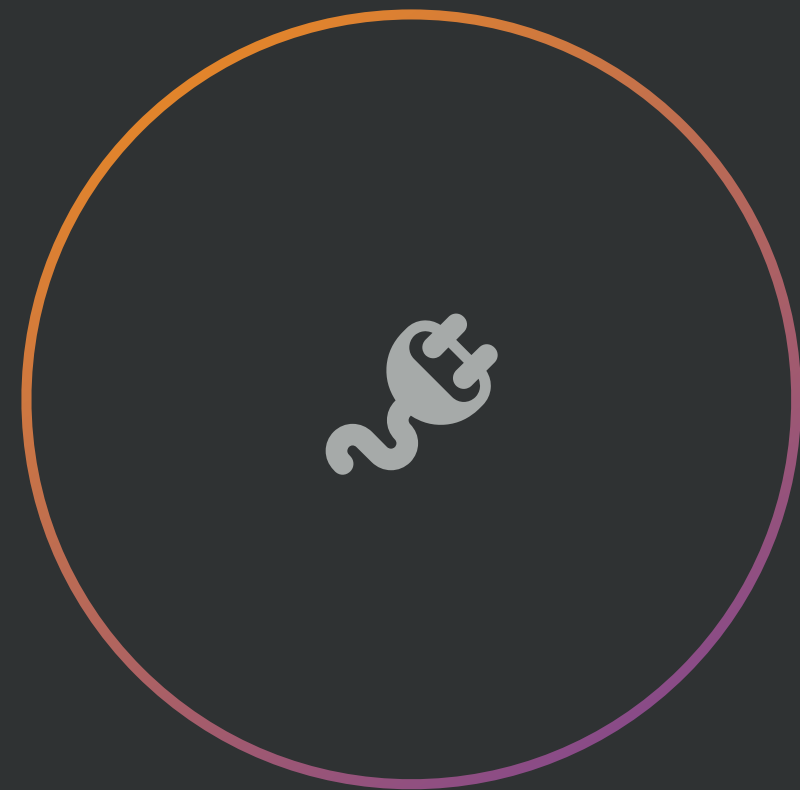
HARDWARE

- More hosts
- Faster CPUs
- SSDs



NETWORKING

- 10gb



A chance to tune config



- Partition by function
- Split features out
- Isolate tables as a whole
- App specific
- Milage may vary



Large volume single table



Constantly growing




No joins - app refactor



Regression testing is
essential



Replaying queries from live




Go Live



Maintenance window

Results!

MySQL app mean time






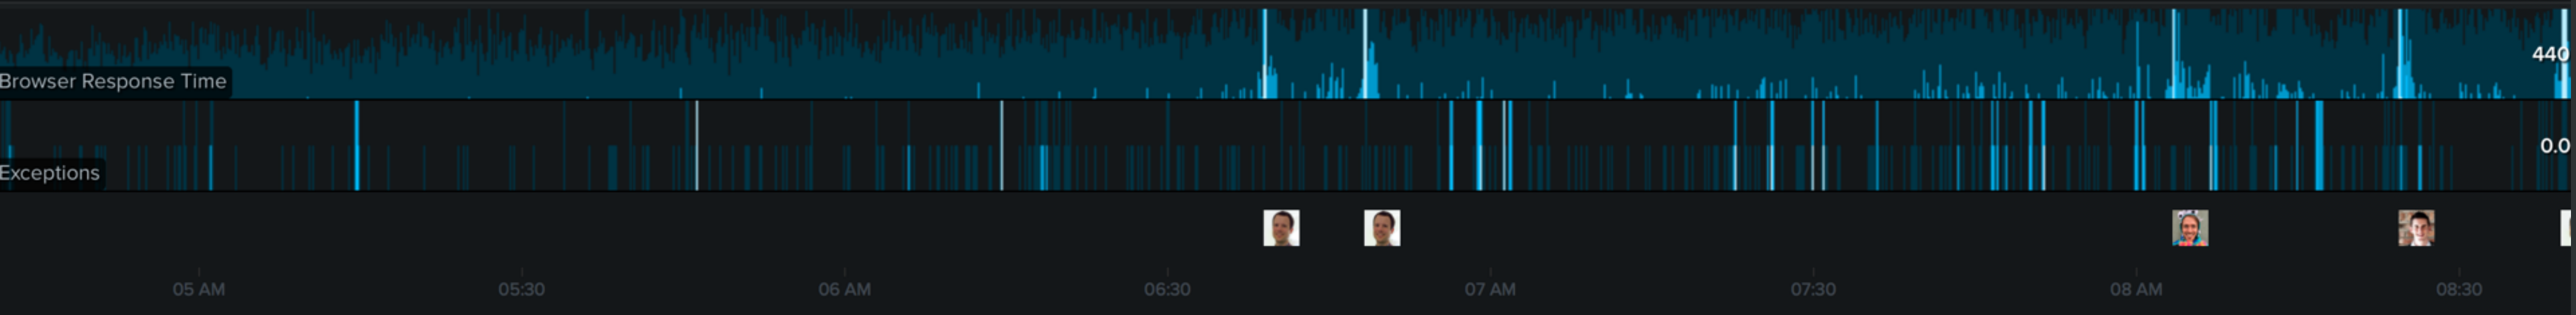
Tooling





Lots to mention



Haystack




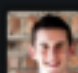
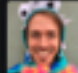


Search...  

Dashboard

Updated: April 13, 2015 08:40:59 **5m** 30m 1h 3h 1d

- Overview
- Dashboard**
- Firehose
- Experienced by you

GitHub::TimeoutMiddleware::RequestTimeout	github/github#30415 github/github#31070 github/github#39550 github/github#31069	3	
last occurred 2 minutes ago — users#show			
NoMethodError		2	
last occurred 3 minutes ago — undefined method `repository' for nil:NilClass			
Braintree::NotFoundError		1	
last occurred 2 minutes ago			
GitRPC::Timeout		1	
last occurred 4 minutes ago — GitRPC::Timeout			
GitHub::TimeoutMiddleware::RequestTimeout		1	
last occurred just now — blob#show			
Timeout::Error		1	
last occurred 2 minutes ago — execution expired			

- Latest Deploys
-  **flipper-stafftools**
deployed by aroben a minute ago
 -  **billing-manager-emails**
deployed by shayfrendt 16 minutes ago
 -  **ifs-data-quotas-refactor**
deployed by tclem 37 minutes ago
 -  **flipper-peek-bar**
deployed by aroben 2 hours ago
 -  **flipper-middleware**
deployed by aroben 2 hours ago

Search...

GitHub / Timeout::Error

[Create issue on github/github](#)

First occurred Monday, December 08 2014 11:11 AM PST

Raw 100 Web flow 100 Git 100

Overview

Dashboard

Firehose

Experienced by you

Latest Deploys

deployed by southgate 3 minutes ago

flipper-stafftools
deployed by aroben 13 minutes ago

billing-manager-emails
deployed by shayfrendt 28 minutes ago

ifs-data-quotas-refactor
deployed by tclem an hour ago

flipper-peek-bar
deployed by aroben 2 hours ago

Searches

Rails 3

404 Errors



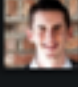
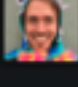

Mobile



Frequent Users	Frequent Repos	Frequent Routes	Frequent Servers
43	228	235	27
8	66	74	26
4	45	52	24
3	18	29	24
3	17	28	24
3	11	17	23
2	8	14	23
2	7	13	22
2	7	13	22
	6		21

April 13, 2015 08:34:58

- Overview
- Dashboard**
- Firehose
- Experienced by you

- Latest Deploys
-  [redacted] deployed by southgate 5 minutes ago
 -  **flipper-stafftools** deployed by aroben 15 minutes ago
 -  **billing-manager-emails** deployed by shayfrendt 30 minutes ago
 -  **lfs-data-quotas-refactor** deployed by tclm an hour ago
 -  **flipper-peek-bar** deployed by aroben 2 hours ago

- Searches
- Rails 3**
 - 404 Errors**
 - Mobile**
 - User Errors**
 - Machine Room**

execution expired

action	BlobController#show	🔍
accept	text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,*/*;q=0.8	🔍
areas_of_responsibility	[:raw, :web_flow, :git]	🔍
current_ref	billing-manager-emails	🔍
dbconn	write	🔍
master_pid	21527	🔍
master_started_at	2015-04-13T15:24:06Z	🔍
method	GET	🔍
newsies_dbconn	notifications_primary	🔍
oauth_access_id	nil	🔍
params	{"controller"=>"blob", "action"=>"show", "user_id"=>"[redacted]"}	🔍
rails	3.2.21.github1	🔍
referrer	https://github.com/[redacted]	🔍
remote_ip	[redacted]	🔍
repo	[redacted]	🔍
request_id	[redacted]	🔍


```
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/xpath_context.rb:8:in `register_namespaces'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node.rb:152:in `block in xpath'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node.rb:150:in `map'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node.rb:150:in `xpath'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node_set.rb:114:in `block in css'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node_set.rb:237:in `block in each'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node_set.rb:236:in `upto'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node_set.rb:236:in `each'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/node_set.rb:103:in `css'  
vendor/gems/2.1.5/ruby/2.1.0/gems/nokogiri-1.6.3.1/lib/nokogiri/xml/document_fragment.rb:79:in `css'  
🐱 lib/github/html/name_prefix_filter.rb:30:in `call'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline/filter.rb:137:in `call'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:121:in `block in perform_filter'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:161:in `block in instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/activesupport-3.2.21.gemset/lib/active_support/notifications.rb:123:in `block in instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/activesupport-3.2.21.gemset/lib/active_support/notifications/instrumenter.rb:20:in `instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/activesupport-3.2.21.gemset/lib/active_support/notifications.rb:123:in `instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:160:in `instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:120:in `perform_filter'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:106:in `block (2 levels) in call'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:105:in `each'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:105:in `inject'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:105:in `block in call'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:161:in `block in instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/activesupport-3.2.21.gemset/lib/active_support/notifications.rb:123:in `block in instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/activesupport-3.2.21.gemset/lib/active_support/notifications/instrumenter.rb:20:in `instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/activesupport-3.2.21.gemset/lib/active_support/notifications.rb:123:in `instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:160:in `instrument'  
🔍 vendor/gems/2.1.5/ruby/2.1.0/gems/html-pipeline-1.11.0/lib/html/pipeline.rb:103:in `call'  
🐱 app/helpers/blob_markup_helper.rb:116:in `block in markup_blob!'  
🔍 vendor/ruby/71c83a957a8b31cf30ecf1e622414c7370815/lib/ruby/2.1.0/timeout.rb:91:in `block in timeout'  
🔍 vendor/ruby/71c83a957a8b31cf30ecf1e622414c7370815/lib/ruby/2.1.0/timeout.rb:35:in `block in catch'  
🔍 vendor/ruby/71c83a957a8b31cf30ecf1e622414c7370815/lib/ruby/2.1.0/timeout.rb:35:in `catch'  
🔍 vendor/ruby/71c83a957a8b31cf30ecf1e622414c7370815/lib/ruby/2.1.0/timeout.rb:35:in `catch'  
🔍 vendor/ruby/71c83a957a8b31cf30ecf1e622414c7370815/lib/ruby/2.1.0/timeout.rb:106:in `timeout'  
🐱 app/helpers/blob_markup_helper.rb:115:in `markup_blob!'  
🐱 app/helpers/blob_markup_helper.rb:97:in `markup_blob'  
🐱 app/helpers/blob_markup_helper.rb:36:in `format_blob'  
🐱 app/views/blob/_blob_content.html.erb:1:in `_app_views_blob__blob_content_html_erb__757533915030191661_69855271779560'
```

```
27 def call
28   return doc unless prefix
29
30   doc.css('*[name], *[id]').each do |element|
31     %w(name id).each do |attribute|
32       if element[attribute]
33         element[attribute] = element[attribute].sub(/\\A(#{prefix})?/, prefix)
34       end
35     end
36   end
```



Slow queries

April 13, 2015 08:35:55

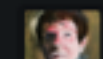

Overview

Dashboard


Firehose


Experienced by you


Latest Deploys

 
deployed by southgate 15 minutes ago

 **flipper-stafftools**
deployed by aroben 25 minutes ago

 **billing-manager-emails**
deployed by shayfrendt 40 minutes ago










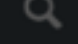




 **ifs-data-quotas-refactor**
deployed by tclm an hour ago

 **flipper-peek-bar**
deployed by aroben 2 hours ago

Searches

Save a search for quick access to it.

```
[1.07 sec] SELECT  
[blurred]
```

accept	application/vnd.github.beta.full+json	
areas_of_responsibility	[:api, :pull_requests]	
connected_host	github-[blurred]	
current_ref	billing-manager-emails	
database_connection	[blurred]rd.iad.github.net./github_production	
master_pid	21343	
master_started_at	2015-04-13T15:24:15Z	
method	GET	
params	{"since"=>"2015-04-13T10:35:25-0500", "repository_id"=>[blurred]}	
rails	3.2.21.github1	
referrer	nil	
remote_ip	[blurred]	
repo	[blurred]	
request_id	[blurred]	

queries

6.234	BEGIN	6.219ms
8.819	SELECT	
comments#post, reque		
12.827	SELECT	
i, route: api/issue_c		
16.846	SELECT	
ments#post, request_		
19.979	SELECT	
tion: github, categor		
29.213	INSERT IN	
[REDACTED]		
49.386	SELECT	
ments#post, request_		
83.505	SELECT N	
ed_users`. `user_id`		
[REDACTED]		
88.441	SELECT N	
users`. `ignored_id`		
t_id: 4A7DB799:09F9:		
93.896	SELECT CO	
spam.id WHERE `issu		
ithub, category: api,		
95.107	UPDATE `i	
107.945	SELECT `a	
e` = 'S' /*applicat		
ms		
108.744	UPDATE `i	
119.699	SELECT `a	
e` = 'S' /*applicat		
ms		
120.875	COMMIT	0.874ms



Haystack is awesome





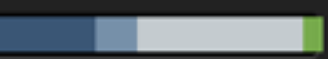
Failbotd









- C daemon
- Takes exceptions from the app
- Queues them
- Posts to Haystack

Staff toolbar




spicycode/gist-hook... Switch to GitHub:lab Rails 3.2.21.github1  160  3705 1.7s 

 This repository Search Explore Gist Blog Help  samlambert +    

32 queries 16.38ms query time 23 rows

START ▲	DURATION	RESULTS	QUERY
45.472ms	0.527ms		lib/github/sql.rb:192 ✖ github_ro_2@github-... /github_production (github-...)
79.021ms	0.309ms		lib/github/sql.rb:192 ✖ github_ro_2@github-... /github_production (github-...)
80.941ms	0.884ms	1	lib/github/sql.rb:264 ✖ github_ro_2@github-... /github_production (github-...)
84.958ms	1.069ms	1	lib/github/sql.rb:264 ✖ github_ro_2@github-... /github_production (github-...)
91.207ms	0.844ms	1	lib/github/sql.rb:264 ✖ github_ro_2@github-... /github_production (github-...)



Chat



~137 remote employees



100+ chat rooms



Hubot



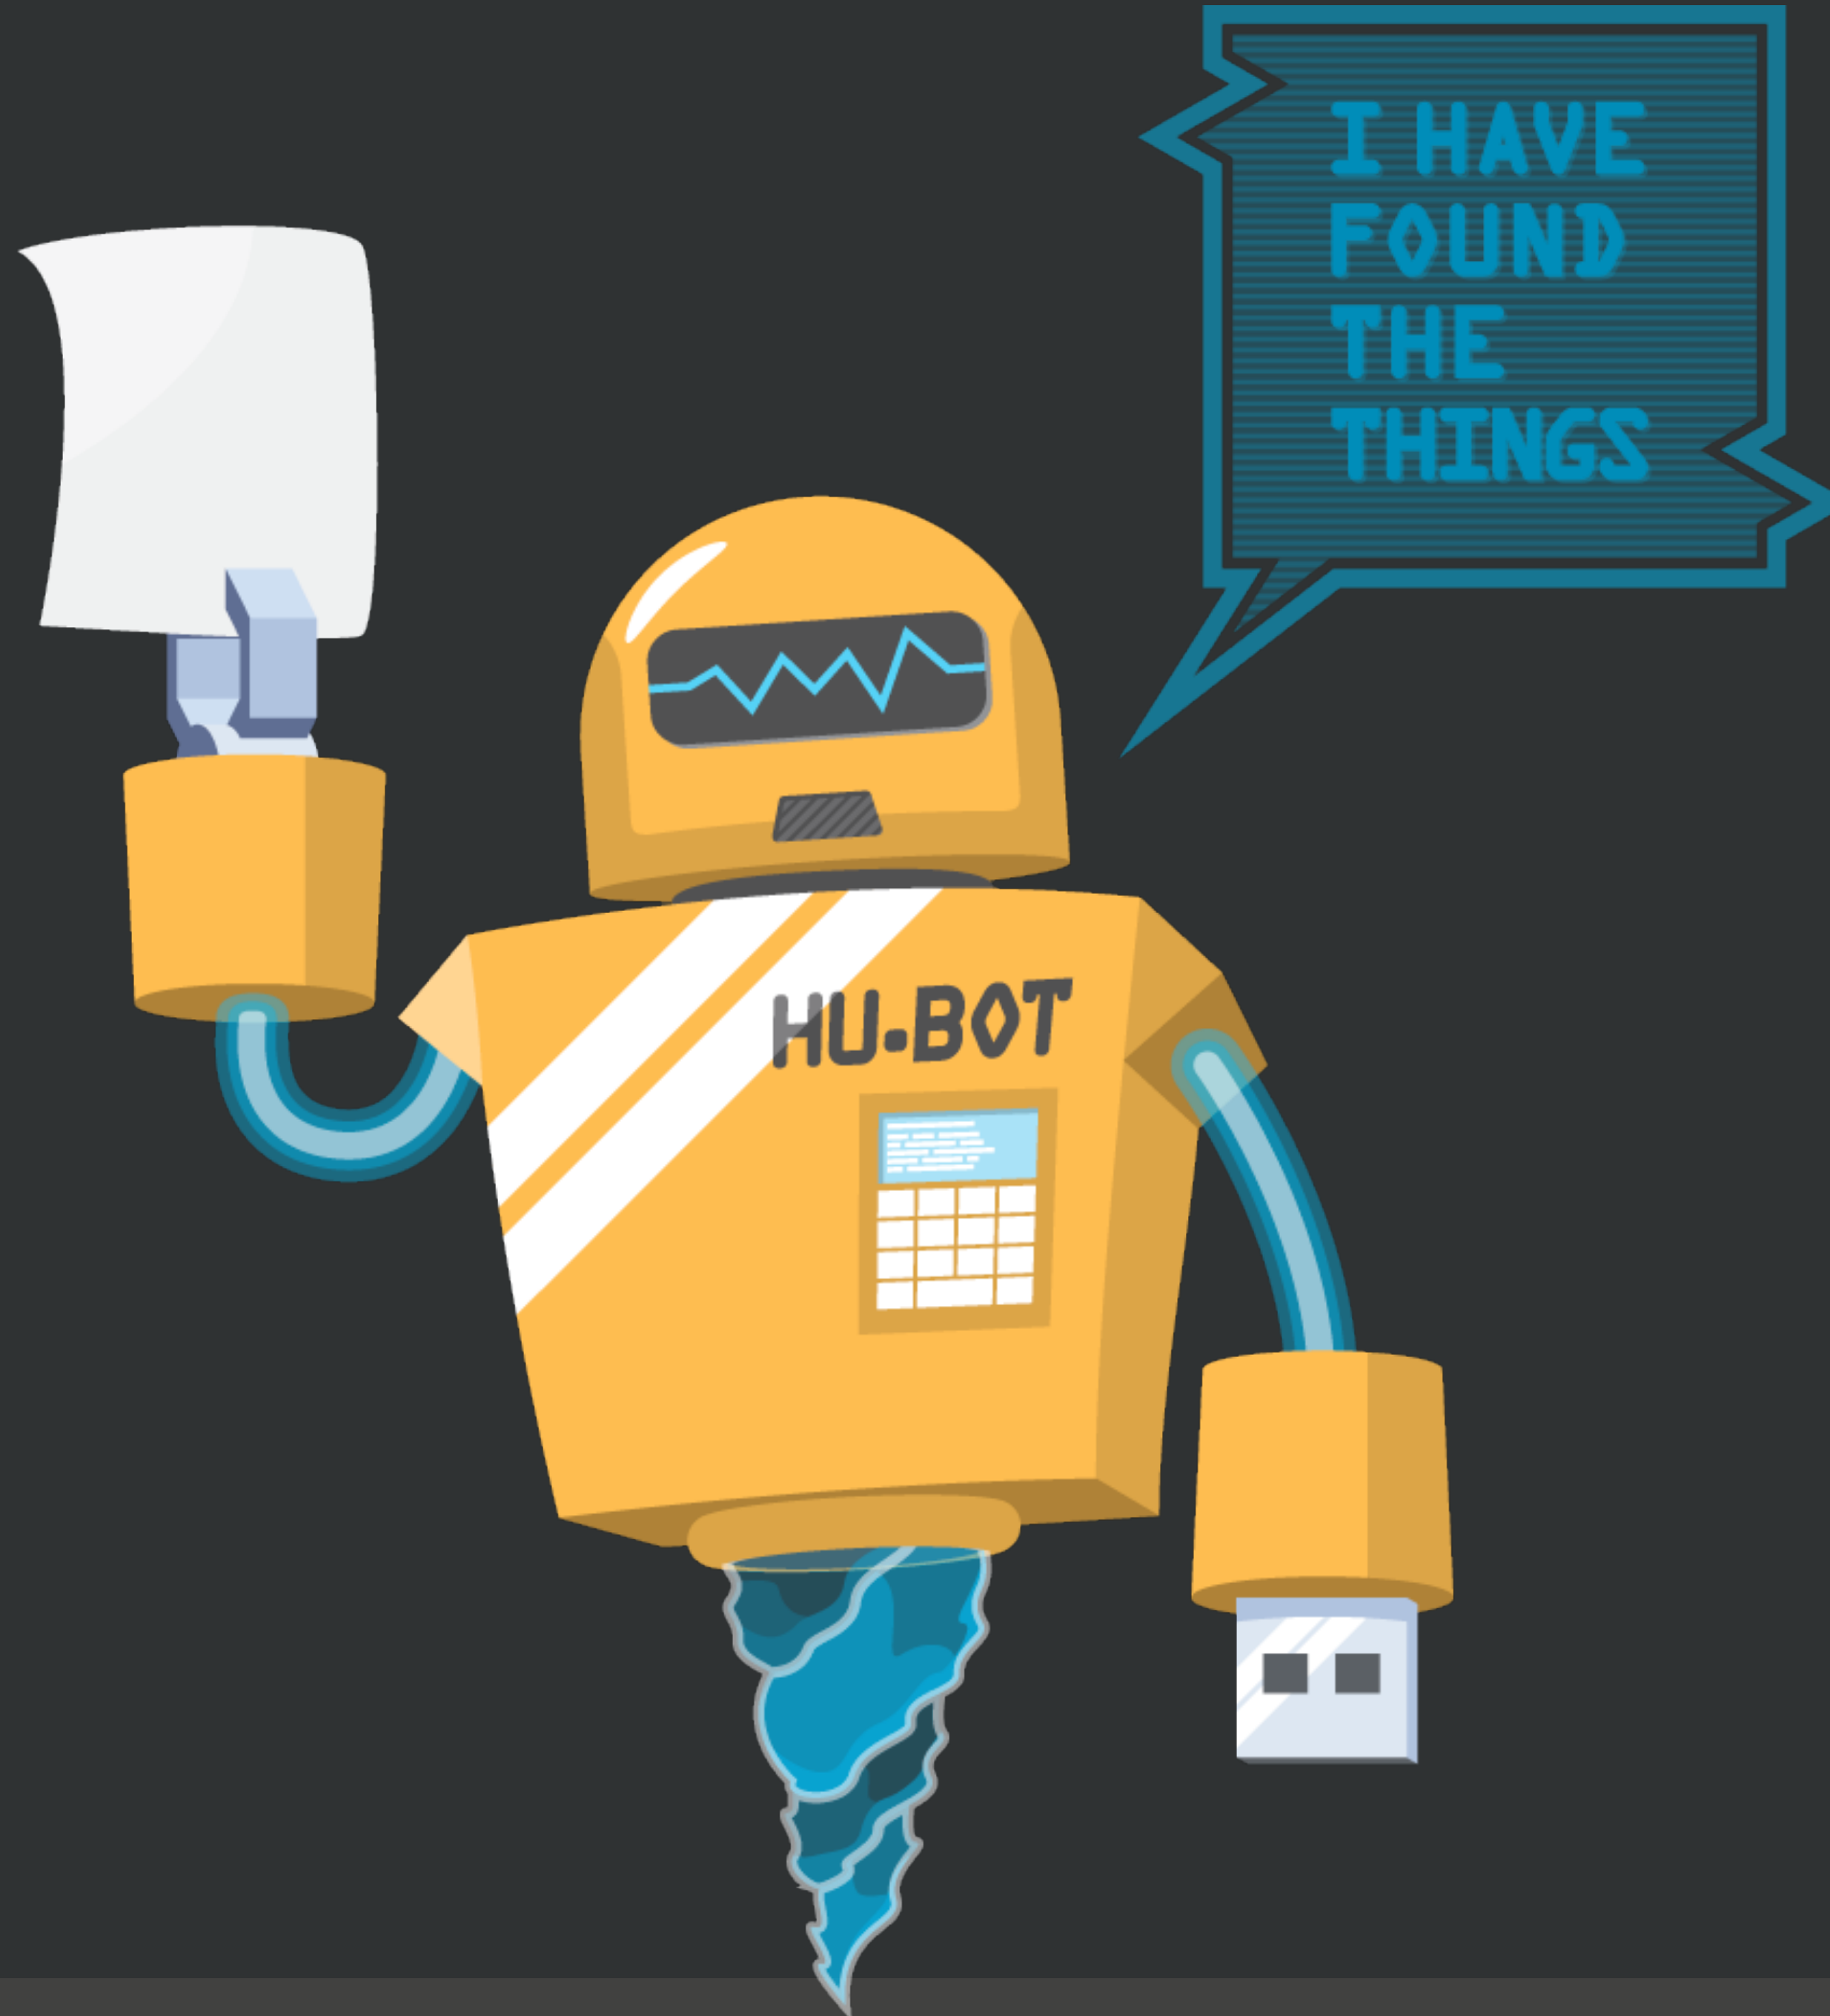
Who is Hubot?

Who is Hubot?



GITHUB, INC., WROTE THE FIRST VERSION OF HUBOT TO AUTOMATE OUR COMPANY CHAT ROOM. HUBOT KNEW HOW TO DEPLOY THE SITE, AUTOMATE A LOT OF TASKS, AND BE A SOURCE OF FUN IN THE COMPANY. EVENTUALLY HE GREW TO BECOME A FORMIDABLE FORCE IN GITHUB. BUT HE LED A PRIVATE, MESSY LIFE. SO WE REWROTE HIM.







Open Source



CoffeeScript



hubot.github.com



github.com/hubot-scripts




Super easy to set up



📖 README.md

Hubot, on Heroku, with the Slack adaptor

This is a version of GitHub's Campfire bot, [Hubot](#), which is designed to be deployed on [Heroku](#) with the [Slack adaptor](#). To deploy your own Hubot, just click the button below:

 Deploy to Heroku

License

[MIT](#), just like the upstream [Hubot license](#), so go 🍻s.



github.com/cobyism/hubot-heroku-slack



What can hubot do?



850+ commands



Nearly everything



Part of shipping is adding
ChatOps



samlambert
/hq me




Hubot

4:

3:

2:

1:





ChatOps is core to our
culture

You are on your own



```
→ ~ mysql
```

```
Welcome to the MySQL monitor.  Commands end with ; or \g.
```

```
Your MySQL connection id is 1
```

```
Server version: 5.6.23 Homebrew
```

```
Copyright (c) 2000, 2015, Oracle and/or its affiliates. All rights reserved.
```

```
Oracle is a registered trademark of Oracle Corporation and/or its  
affiliates. Other names may be trademarks of their respective  
owners.
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```


```
mysql>
```


Invite your friends




TwP
oh we're definitely exercising those disks
`/graph me -2h @collectd.disk_utilization(codesearch2-storage)`

grantr
well we can't hold everything in memory even with 215gb per node. that's still less than 10% of the index



Hubot

mute



grantr
we'll have the same amount of file cache as in the current i2.4xl cluster

grantr
the point of trying 8xls is to see if there's some virtualization issue causing 2 i2.4xls to perform worse than 1 i2.8xl

samlambert
Aiming for in memory is expensive and probably not scalable for anyone large



It's all about context



Context is implicit



Incident response




DDoS



/mitigation enable



/mysql



```
51 /mysql analyze          analyze a table
52 /mysql archive        start or stop archiving on a MySQL cluster
53 /mysql backup         backup a table
54 /mysql backup-list    list backups for a table
55 /mysql clone          clone a table to the MySQL staging environment
56 /mysql cluster        show cluster topology
57 /mysql drop           safely drops a mysql table
58 /mysql drop-all      drops all eligible tables
59 /mysql drop-check     checks if a table is safe to drop
60 /mysql drop-list      lists the tables that are ready to be dropped
61 /mysql drop-pending   lists the tables that are pending dropping
62 /mysql dummy-drop     drops a table from production by renaming it
63 /mysql engine         shows output of show engine innodb status on a given host
64 /mysql explain        explain a query
65 /mysql gather         gathers MySQL diagnostic information
66 /mysql log-migration-version insert the version of a migration into production
67 /mysql inno           list queries
68 /mysql index-stats    Prints out index stats for the specified table
69 /mysql kill           kills queries matching supplied criteria
70 /mysql kills          shows victims of pt-kill on given date
71 /mysql maintenance   begin and end the maintenance of a mysql node
72 /mysql master-swap    returns the binlog position to start replication from on a new master
73 /mysql mutexes       prints a list of mutex events for a MySQL host
74 /mysql nibble         deletes all rows in a table
75 /mysql nibble-all    nibble all non nibbled migration tables
76 /mysql panic         stops the delayed replica to prevent propagation of data loss
77 /mysql pk-pct         shows the % of the primary key range that is used
78 /mysql pool           saves or restores the innodb buffer pool of a host
```




samlambert

`/explain SELECT * FROM users LIMIT 1`



Hubot

samlambert:

Select Type	Table	Type	Possibl...	Key	Key Len	Ref	Rows	Ext...
SIMPLE	users	ALL	null	null	null	null		



Learn together



Collaborative workspace

/mysql backup



samlambert

/mysql backup yubicats



Hubot

A backup job has been created for the yubicats table. You will be notified in [+The Database Infrastructure Room](#) when the backup is complete.



Hubot

[2015-04-13T14:18:43-0700] `github:mysql@10.1.1.1:3306 - @samlambert: yubicats has been backed up for you and sent away to the cloud fairies with the file name github:mysql@10.1.1.1:3306-github_production-yubicats-1428959921.sql.gz`

/mysql clone



samlambert

/mysql clone yubicats



Hubot

[2015-04-13T14:31:54-0700] `github.com:samlambert/yubicats.git`.github.net - @samlambert: yubicats has been transferred into the samlambert database on `mysql://localhost:3306/yubicats`



Killed queries

Killed queries



Hubot

The following queries have been killed in the last minute:

```
+-----+
| Info
+-----+
| SELECT [REDACTED] FROM [REDACTED] WHERE [REDACTED]
+-----+
```





User access

User access



Hubot

probablycorey connected to g[REDACTED] ("analytics") as console_rw via gh-dbconsole on c[REDACTED]



Graphite Editor

Use shift-tab for autocomplete | [General API](#) | [Target API](#)

```
from: "-570d"
until: "-575d"
max: 200
areaMode: "stacked"
areaAlpha: "0.5"
hideGrid: true
bgcolor: "2F3233"
thickness: 1
fontSize: 15
hideLegend: "true"
hideAxes: true

template: "plain"
targets:
  "github.unicorn.{browser,anon,api,raw,googlek
    summarize: ["lmin","avg"]
    substr: [2,3]
  "githauth.unicorn.githauth.mysql.time.mean":
    summarize: ["lmin","avg"]
    substr: [2,3]
```

[Delete](#) [Save](#)

Hubot

```
/graph me -570d..-575d substr(summarize(github.ui
```




samlambert

/graph me -1h @mysql.commands.all



Hubot

Commands across the mysql1 cluster





Amen



- Monitoring app
- Set up backends
- Easily gather metrics
- Pushed to Graphme




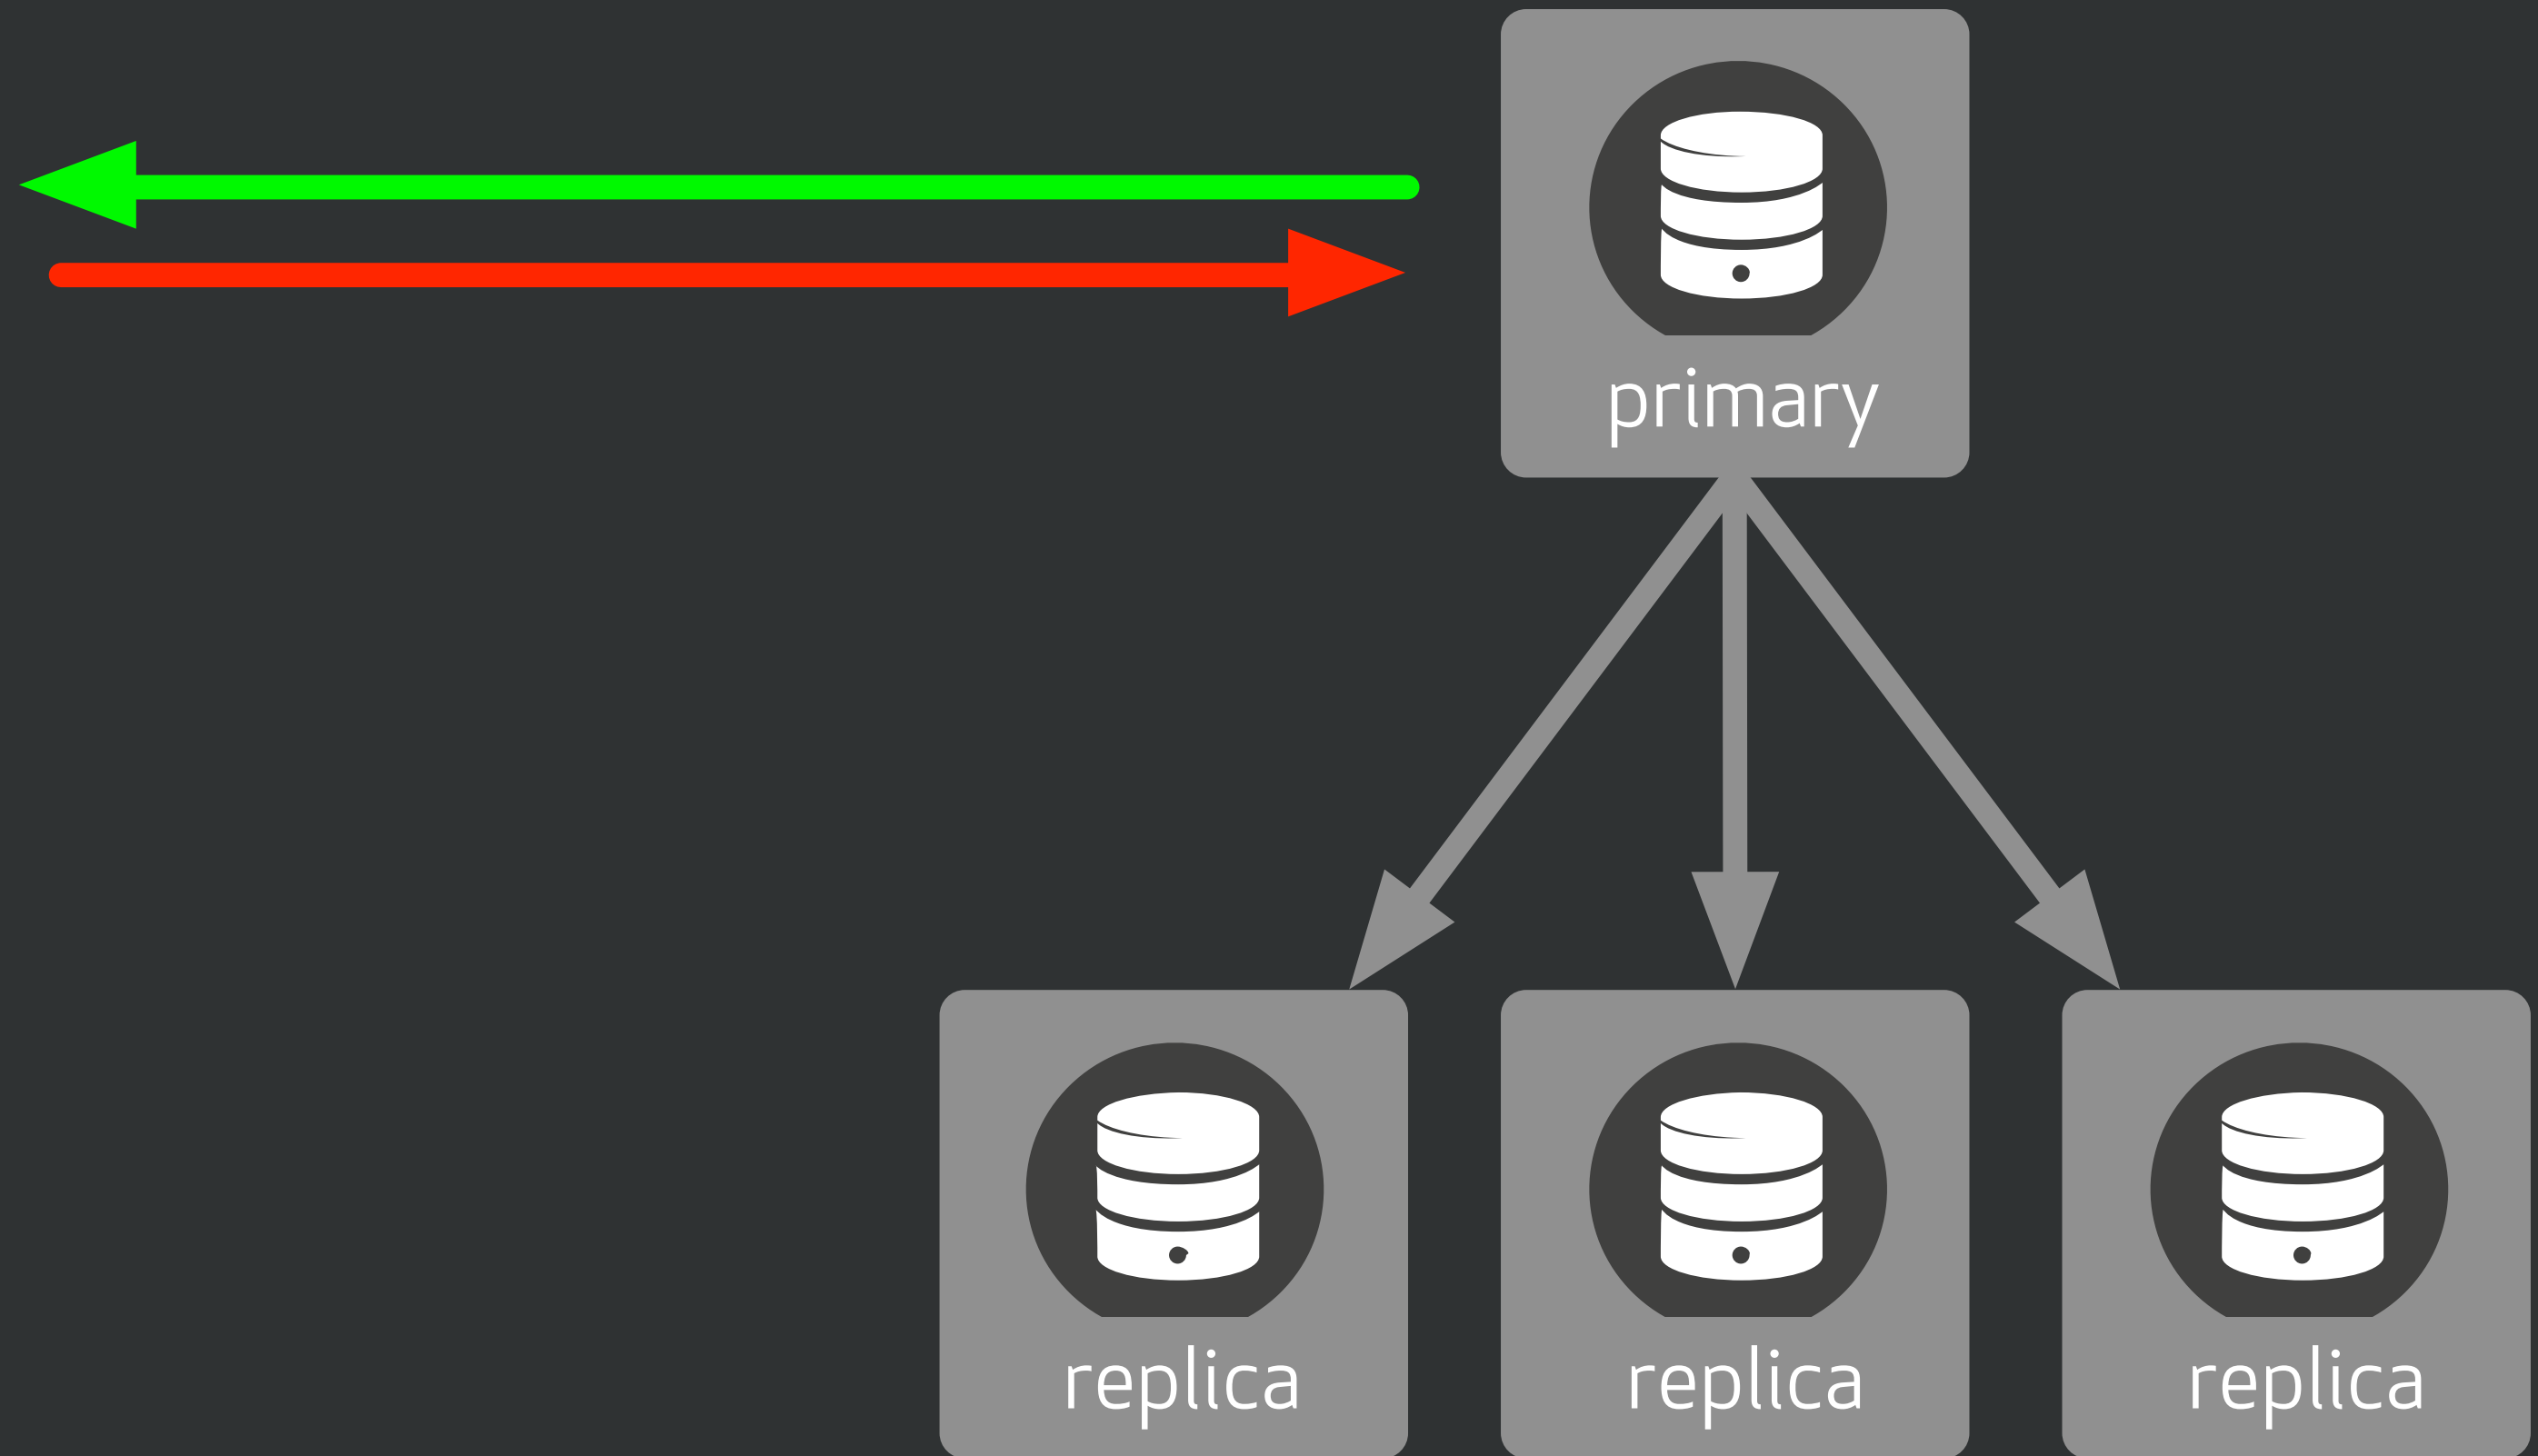
We have these tools. Let's use them



All the hardware

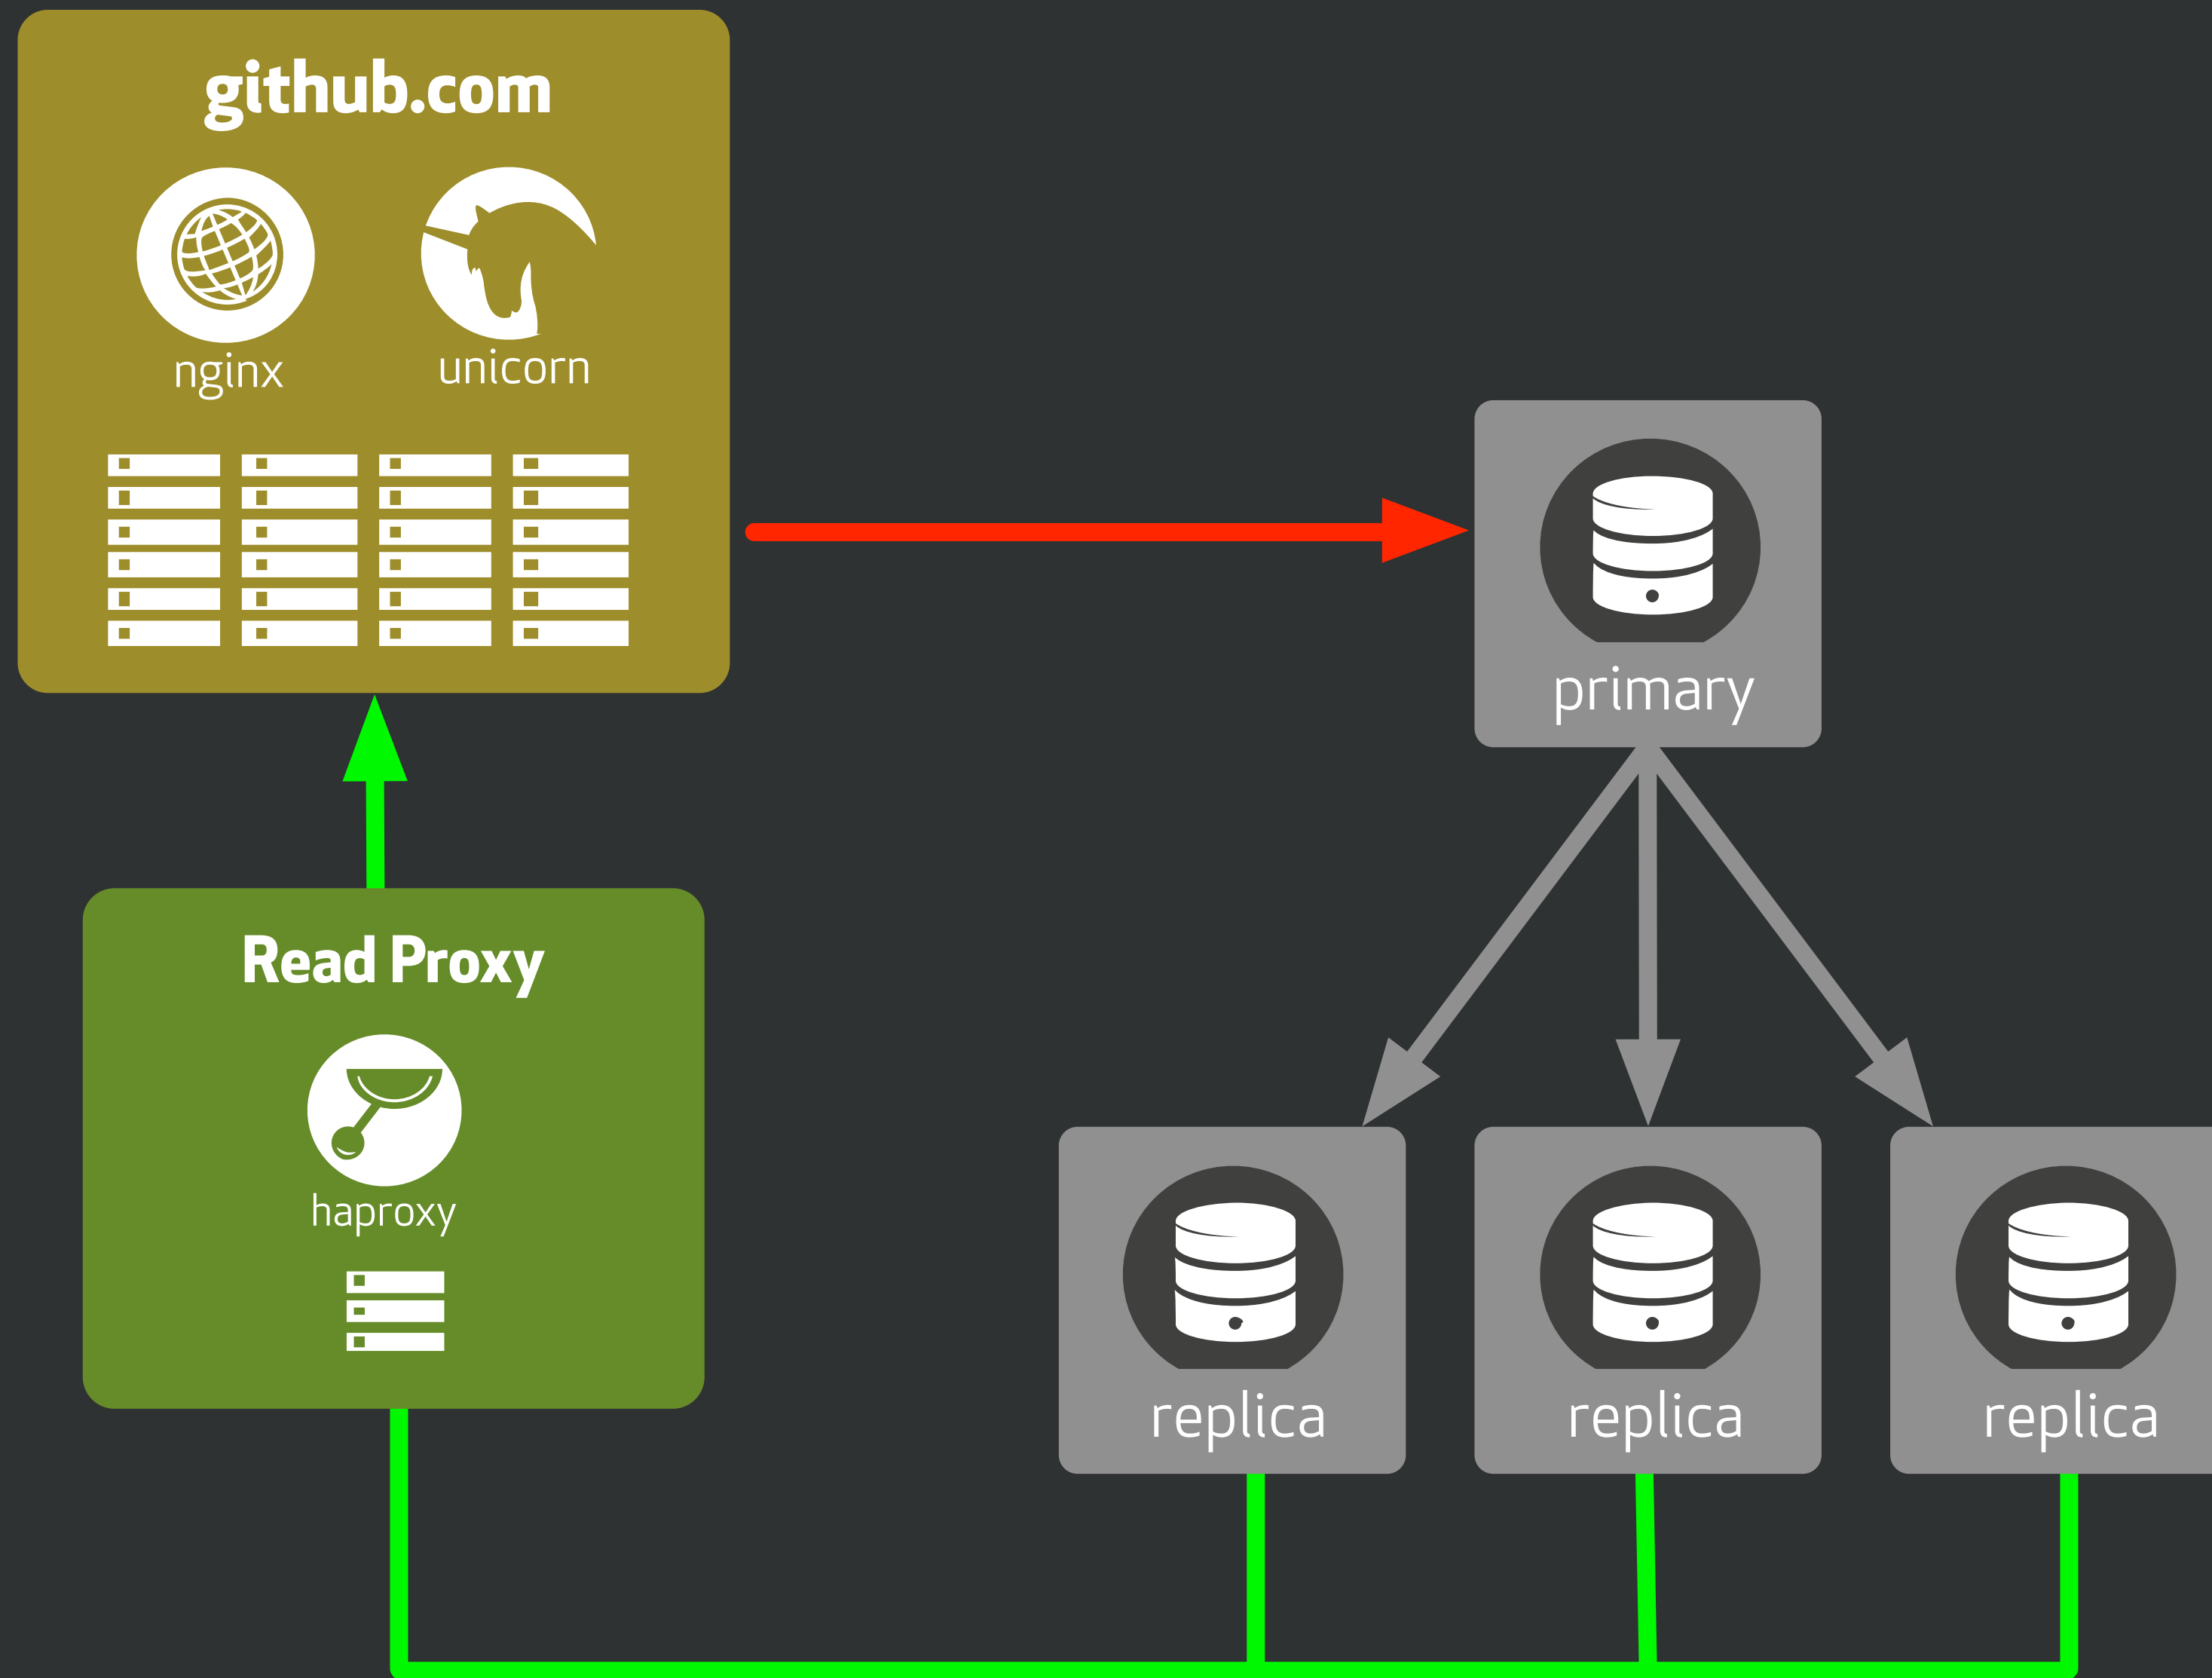
Single host







Doesn't scale





Use existing patterns



/mysqlproxy

/mysqlproxy status



samlambert

/mysqlproxy status



Hubot

haproxy status for `mysql`:

	sess	peak	status	weight	chkfails
*** mysql					
github-	54	75	UP	10	18
github-	55	76	UP	10	51
github-	58	75	UP	10	17
github-	55	72	UP	10	26
github-	0	0	DRAIN	0	4
github-	0	0	DRAIN	0	336
github-	54	75	UP	10	18
github-	0	0	DRAIN	0	0
github-	0	0	DRAIN	0	27
github-	56	75	UP	10	108
github-	0	0	DRAIN	0	5



/mysqlproxy weight



`/mysqlproxy disable`



`/mysqlproxy enable`



App refactor



Stick to master for a second
after POST




Write alerts go to Haystack

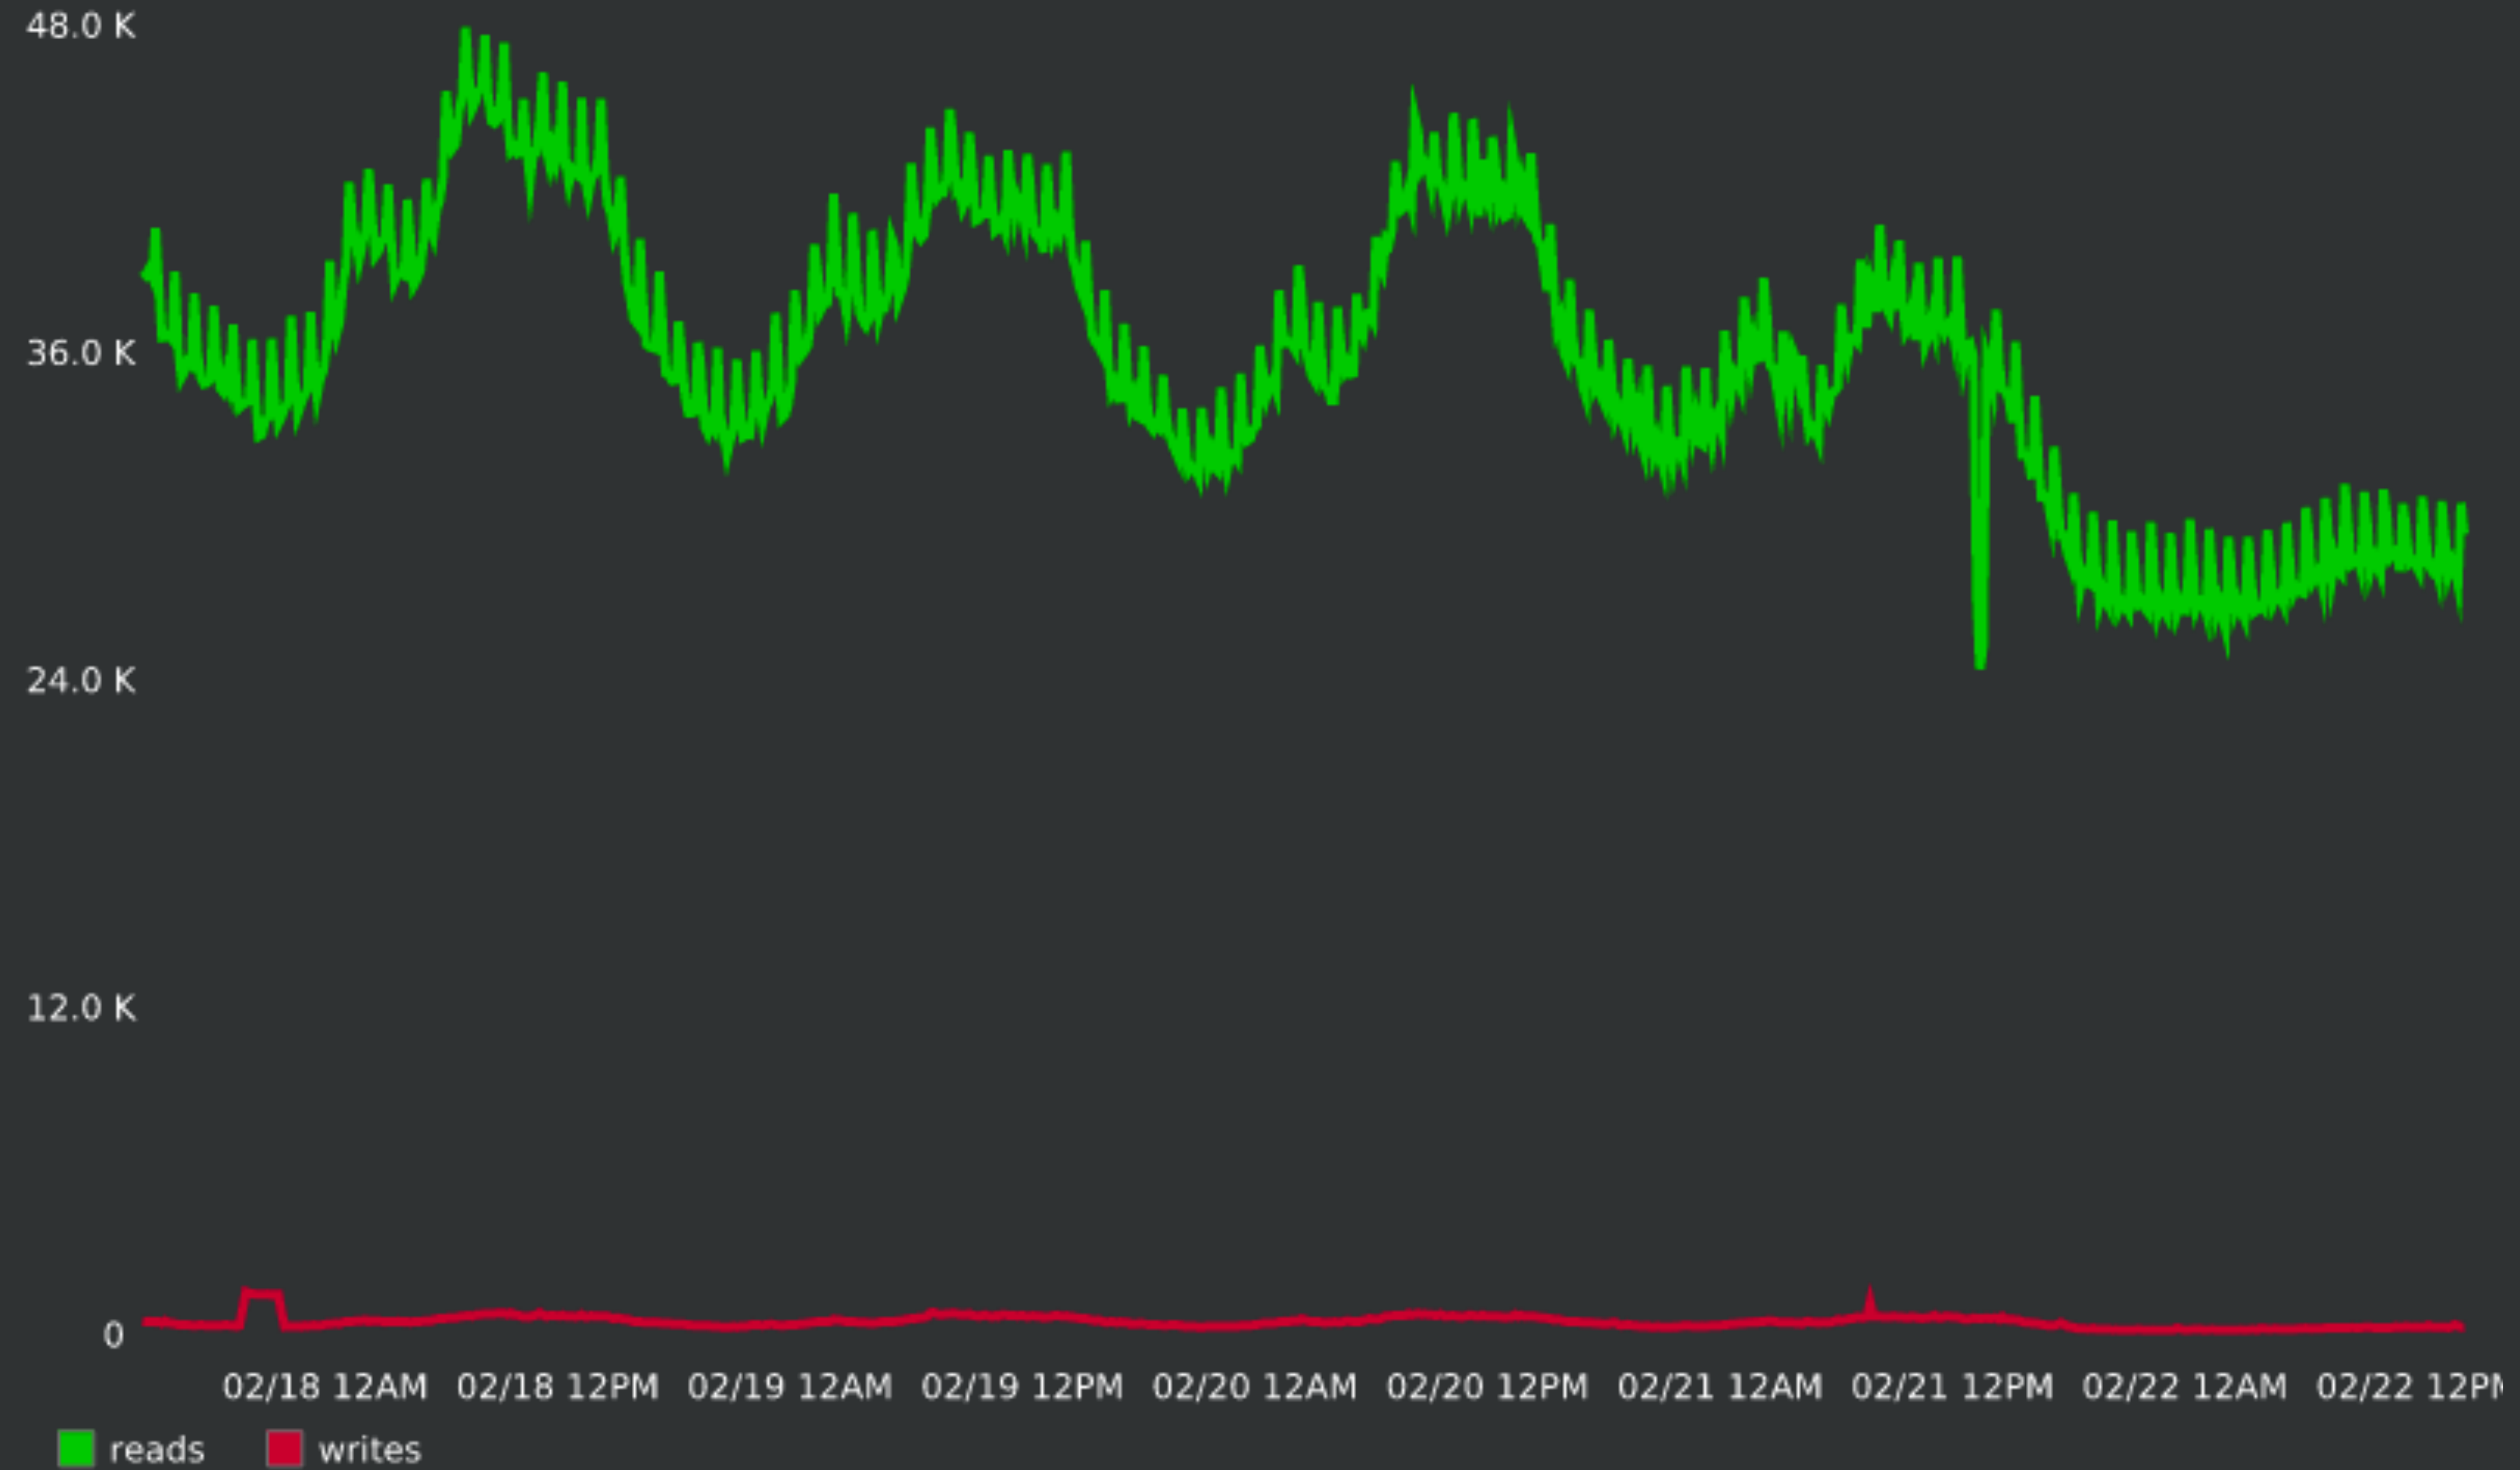


Gradual iteration over the
codebase


Change



Single host




Master





Now we are sensitive to delay


Delay!!





What is happening?

Delay!!




Amen - IO waits by index



```
51 @client.query "SELECT object_name AS table_name,\n52 LOWER(index_name) as index_name,\n53 count_fetch AS select_count,\n54 count_insert AS insert_count,\n55 count_update AS update_count,\n56 count_delete AS delete_count\n57 FROM performance_schema.table_io_waits_summary_by_index_usage\n58 WHERE object_schema = 'github_production'\n59 AND object_name NOT REGEXP '^_.'", (err, results) =>\n60   if err\n61     log 'error', err\n62   else\n63     prefix = "#{metric}.table_stats.waits"\n64     for row in results\n65       @gauge "#{prefix}.#{row.table_name}.#{row.index_name}.select.count", parseInt(row.select_count)\n66       @gauge "#{prefix}.#{row.table_name}.#{row.index_name}.insert.count", parseInt(row.insert_count)\n67       @gauge "#{prefix}.#{row.table_name}.#{row.index_name}.update.count", parseInt(row.update_count)\n68       @gauge "#{prefix}.#{row.table_name}.#{row.index_name}.delete.count", parseInt(row.delete_count)
```



Delete wait volume per table (Top 10) by most deviant





Haystack



Affected rows bucket

October 16, 2014 08:35:22



samlambert mentioned in [Repository deletes create replication lag](#)

October 16, 2014 08:47:39

github/github#34418

Changed 9665 rows

```
DELETE FROM `commit_contributions` WHERE ((repository_id = [REDACTED]))
```

class_name	QueryAffectedRowsLogger::QueryAffectedRows	🔍
current_ref	master	🔍
network_id	[REDACTED]	🔍
query	DELETE FROM `commit_contributions` WHERE ((repository_id = [REDACTED]))	🔍
queue_time	0.09425826991015625	🔍
queued_from	Api::Repos#DELETE	🔍
rails	3.0.20.github12	🔍
request_id	[REDACTED]	🔍
resque_args	[25307443, nil, [REDACTED]]	🔍
resque_job	GitHub::Jobs::RepositoryDelete	🔍
ruby	ruby 2.1.2p95-github-tcmalloc (71c08a81c0) [x86_64-linux]	🔍
server	[REDACTED]	🔍



samlambert mentioned in **Repository deletes create replication lag**
October 16, 2014 08:47:39



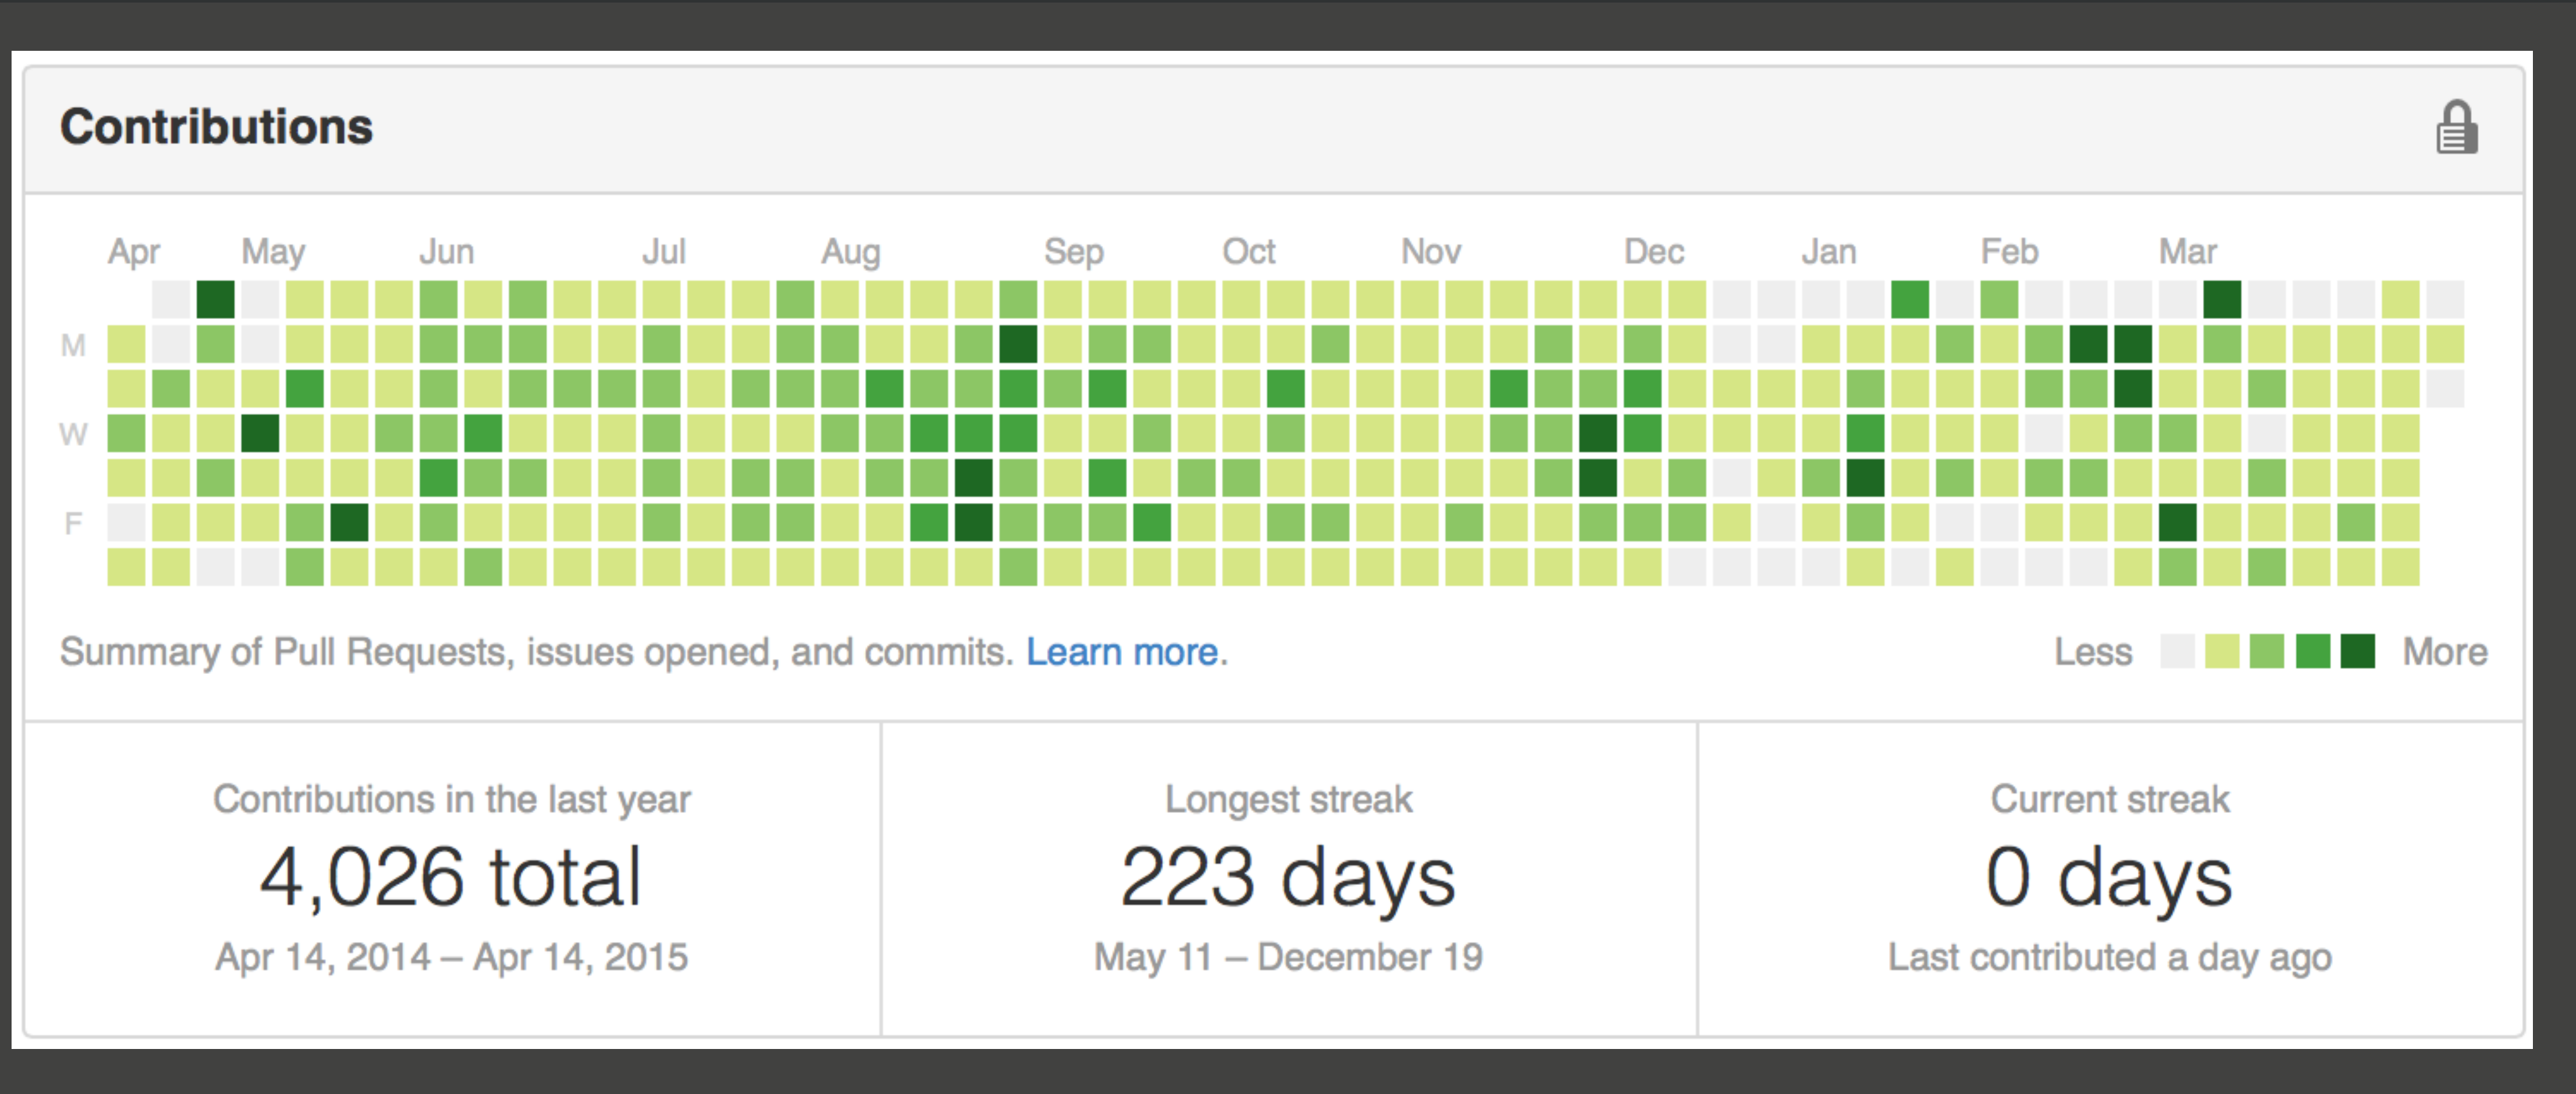
Delete wait volume per table (Top 10) by most deviant






Too many affected rows

Contributions





Problem found




Throttler



```
51 # Public: throttle based on replication delay
52 #
53 # Waits until replication delay is at or below ALLOWABLE_DELAY, then yields.
54 def throttle
55   GitHub.stats.increment "replication-delay.called"
56   slept = 0
57   total_slept = 0
58
59   loop do
60     begin
61       hostname, delay = maximum_replication_delay
62
63       if delay <= ALLOWABLE_DELAY
64         return yield
65       end
66
67       log "#{hostname} delay #{delay} > #{ALLOWABLE_DELAY}, sleeping"
68
69       sleep ALLOWABLE_DELAY
70       slept += ALLOWABLE_DELAY
71       total_slept += ALLOWABLE_DELAY
72       GitHub.stats.increment "replication-delay.slept"
73
74       if slept >= REPLICAS_RESET_TIME
75         GitHub.stats.increment "replication-delay.reload-replicas"
76         reset_replicas
77         slept = 0
78       end
79     end
80   end
81 end
```




```
51
52
53
54 def self.clear_contributions(repo)
55   repo.commit_contribution_ids.each_slice(DELETE_BATCH_SIZE) do |slice|
56     throttle do
57       GitHub::SQL.run <<-SQL, :ids => slice
58         DELETE FROM commit_contributions WHERE id IN :ids
59       SQL
60     end
61   end
62 end
63
64 end
65
66
67
```



Problem solved



We are hiring!!



- GTIDs
- MySQL 5.7
- Cluster aware monitoring
- Failover improvements
- So much more



Questions?



Sam Lambert

Director of Technology

 github.com/samlambert

 twitter.com/isamlambert

 samlambert.com