

# Fast and Simple

## Observing Code and Infra Deployments at Honeycomb

**Shelby Spees**

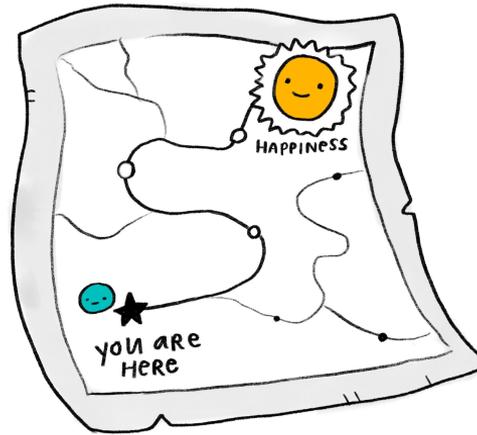
@shelbyspees

#SKILupDay

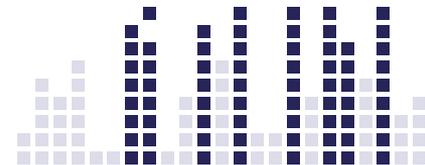
June 18, 2020



w/ illustrations by @emilywithcurls!



Observability is evolving quickly.





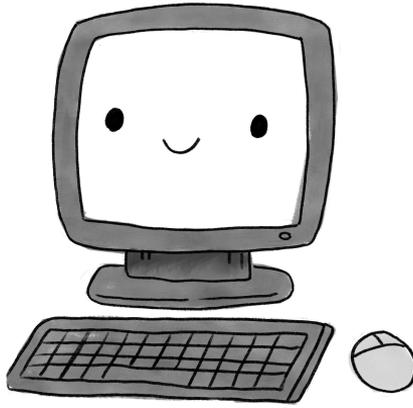
We need **velocity** and **reliability**.



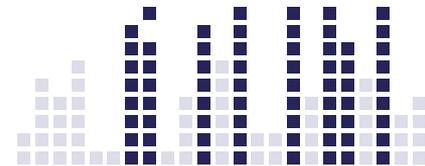


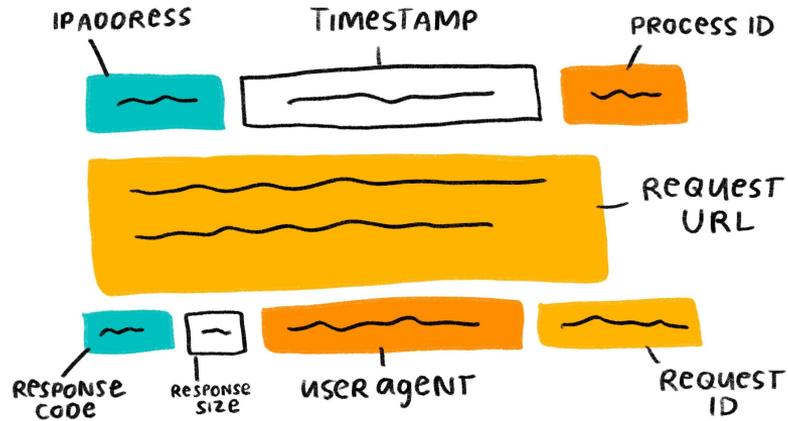
A dozen engineers **build** Honeycomb.



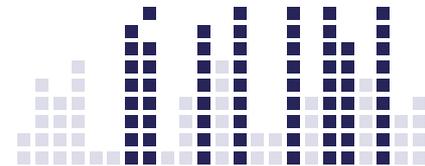


**We make systems humane to run,**



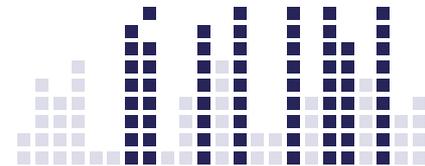


by **ingesting** telemetry,



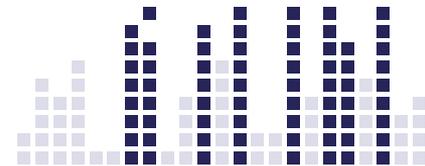


enabling data **exploration**,





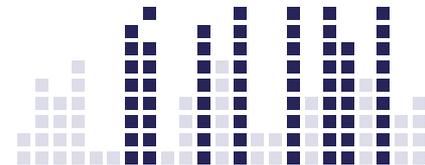
and **empowering** engineers.



A hand-drawn calendar for the month of July. The title 'JULY' is written in a yellow banner at the top. The calendar grid shows dates from 1 to 31. The 5th column, representing Fridays, is highlighted in yellow. The dates in this column are 5, 12, 19, and 26.

JULY						
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

Yes, we **deploy** on Fridays.



# distinct build\_id per day

Datasets / system\_stats

☆ Add to Board

➔ Share

11/1 12:00am to 6/11 8:03am

Add a description for this query

fx VISUALIZE

COUNT\_DISTINCT(build\_id)

WHERE

build\_id != 0  
host\_group != labrador

AND

GROUP BY

None; combine all rows

ORDER BY

None

LIMIT

None

Run Query

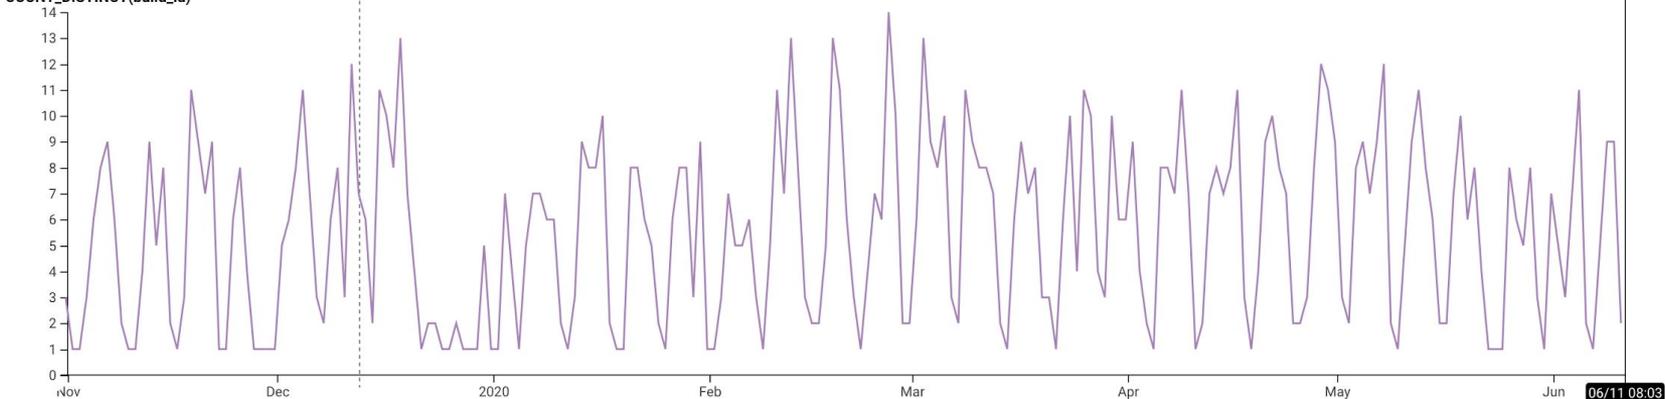
Run 8 minutes ago

Nov 1 2019, 12:00 AM – Jun 11 2020, 8:03 AM Granularity: 1 day

Results BubbleUp Traces Raw Data

Graph Settings

COUNT\_DISTINCT(build\_id)

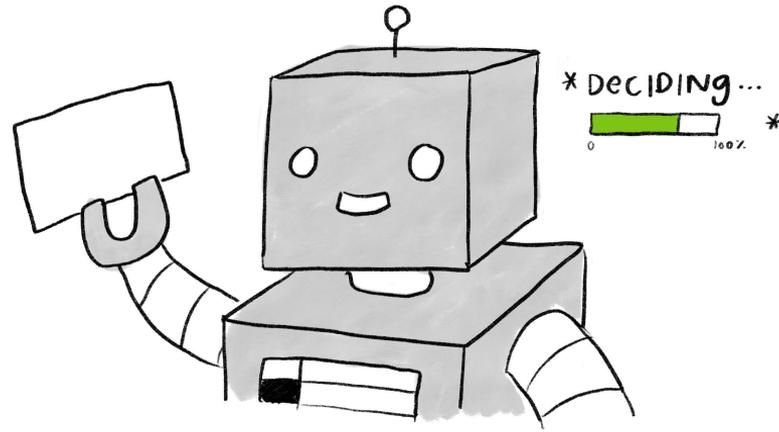


COUNT_DISTINCT(build_id)
879

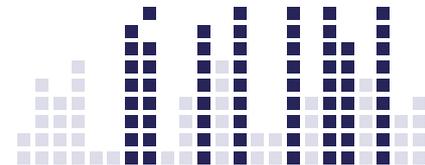
elapsed query time: 2.521905601s rows examined: 90,322,310 nodes reporting: 100%



@shelbyspees at #SKILupDay



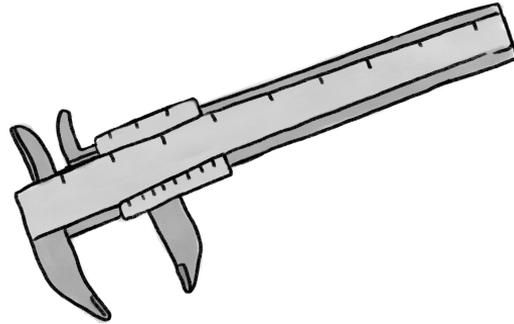
**Continuous delivery** is an investment.



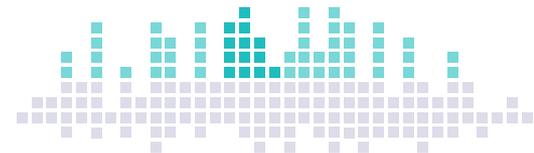


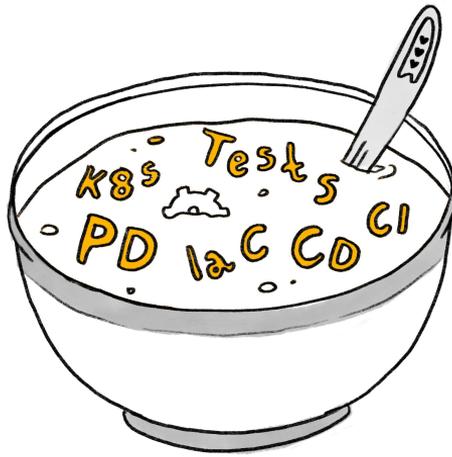
# How did we get **there**?



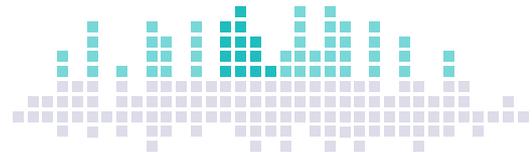


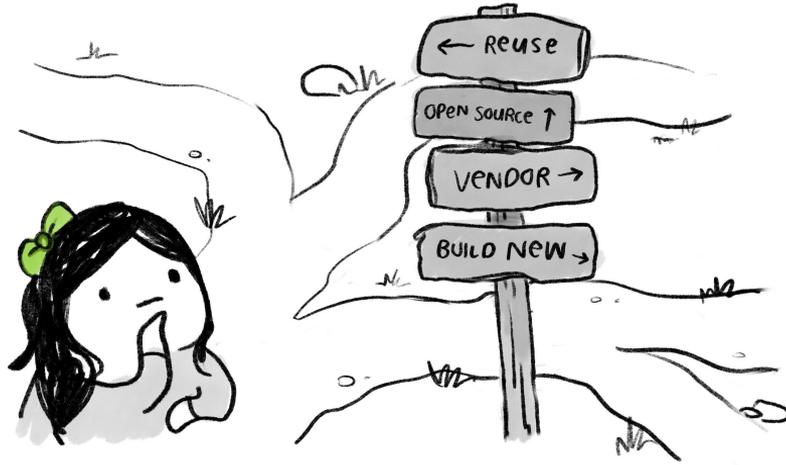
**Investment** in tooling paid off,



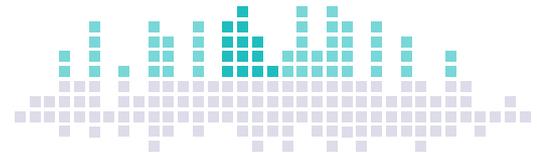


but we didn't need **all the soup.**



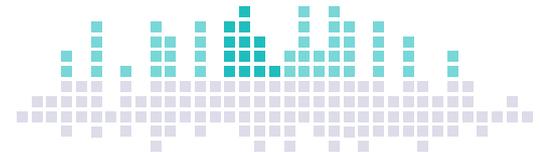


We needed to be thoughtful.



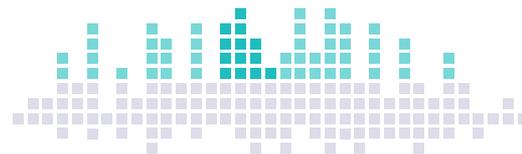


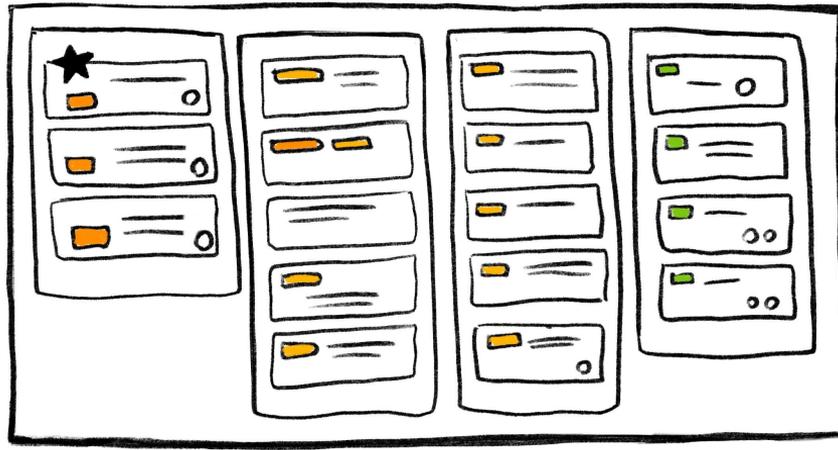
and we needed **cultural process** too!



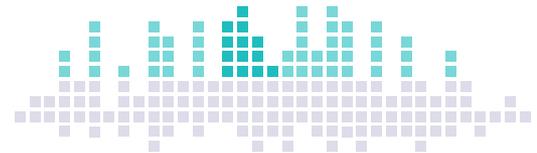


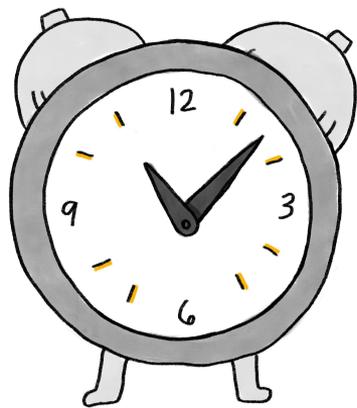
**Continuously evaluate** tech debt.



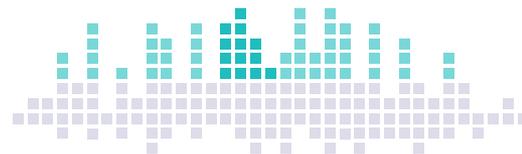


and improve **where it matters.**



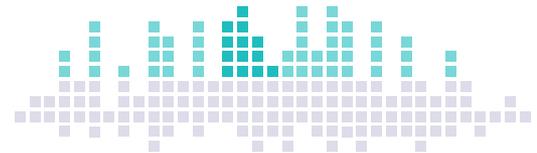


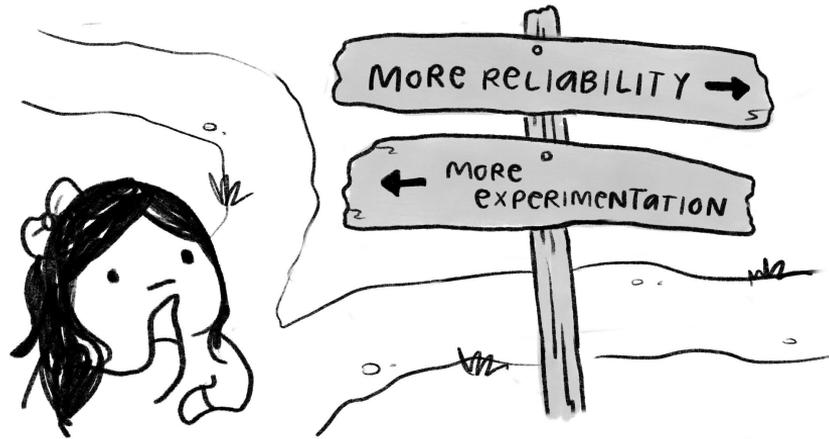
**Speed up product development.**



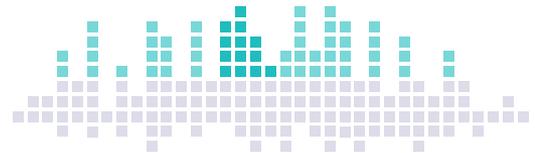


**and infrastructure safety.**



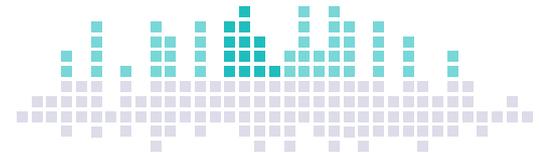


**Embrace risk, but mitigate it.**





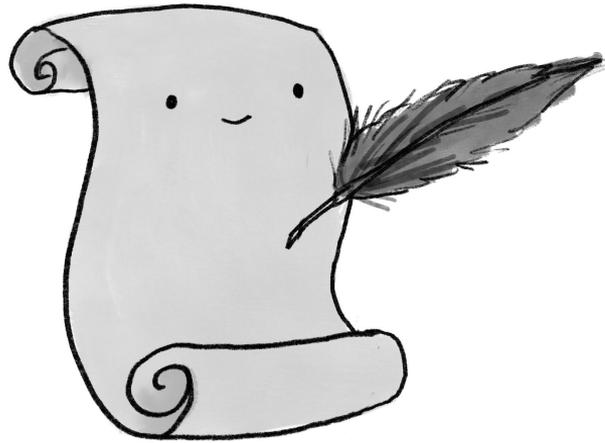
and never stop **improving.**



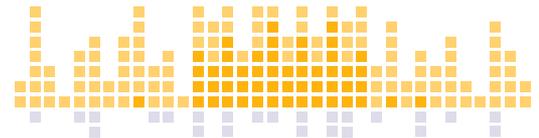
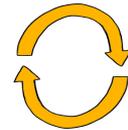
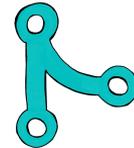
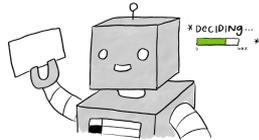
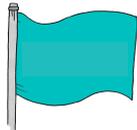


# Shipping prod features



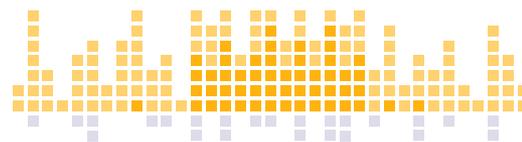
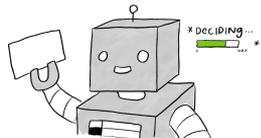
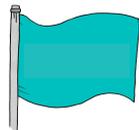


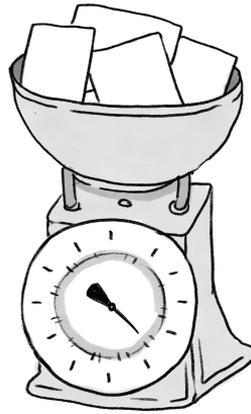
# What's our recipe?



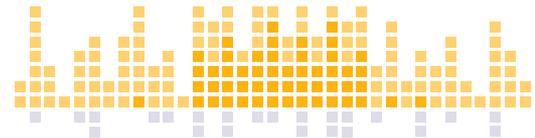
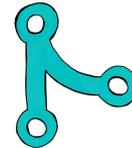
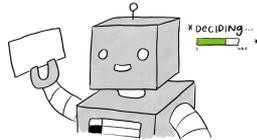
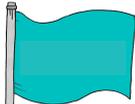
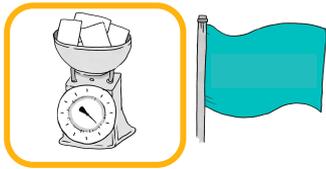


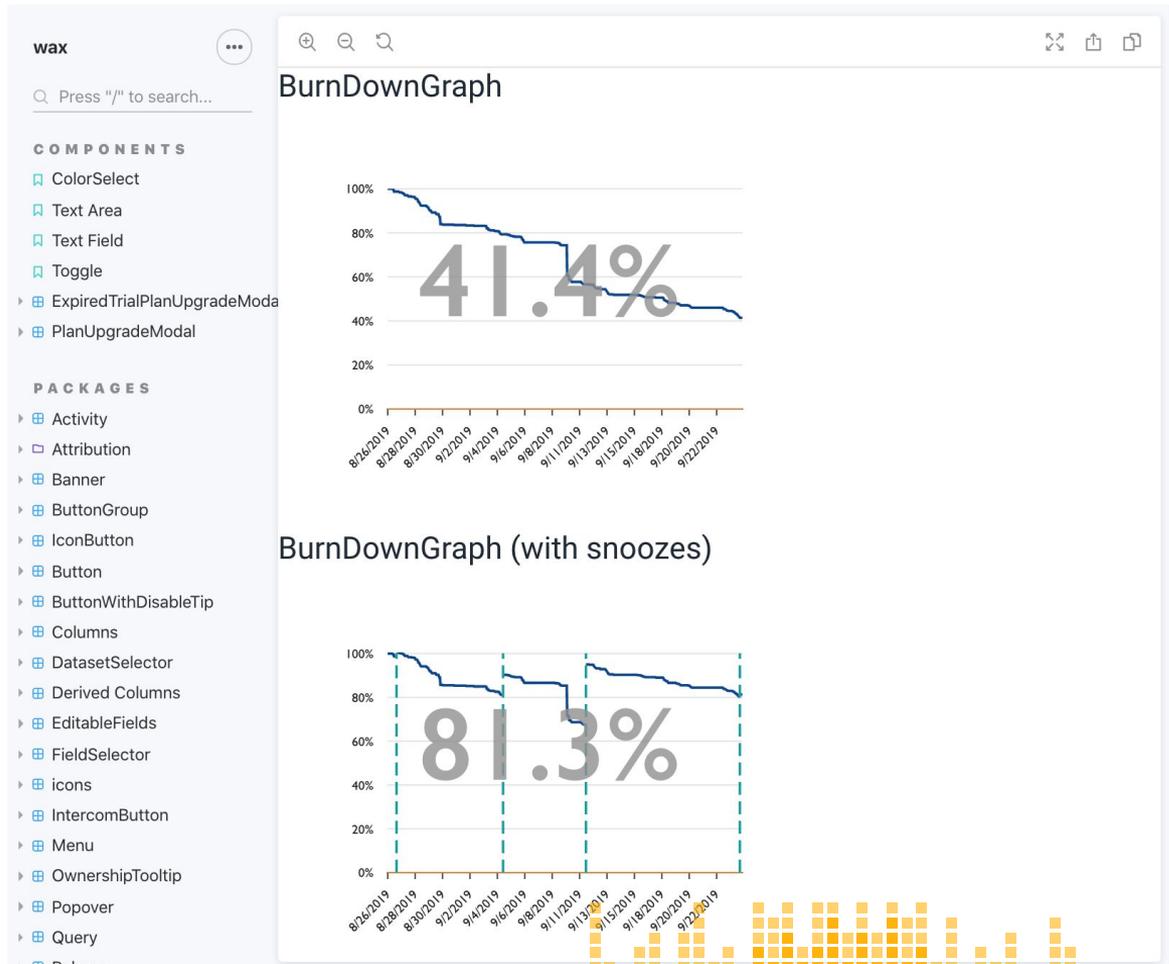
# Instrument as we code.

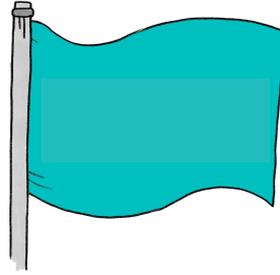




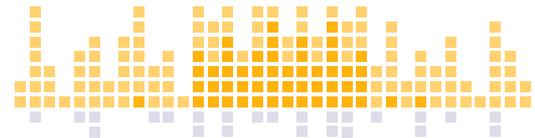
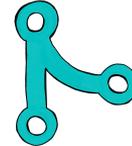
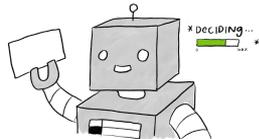
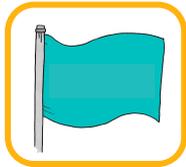
# Functional and visual testing.

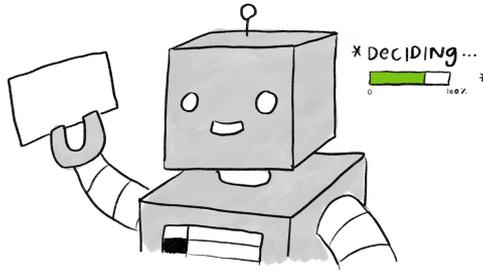




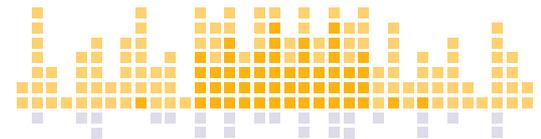
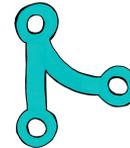
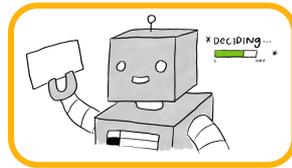
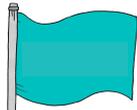


# Design for **feature flag** deployment.





# Automated integration.



Welcome to a preview of the new CircleCI Job experience

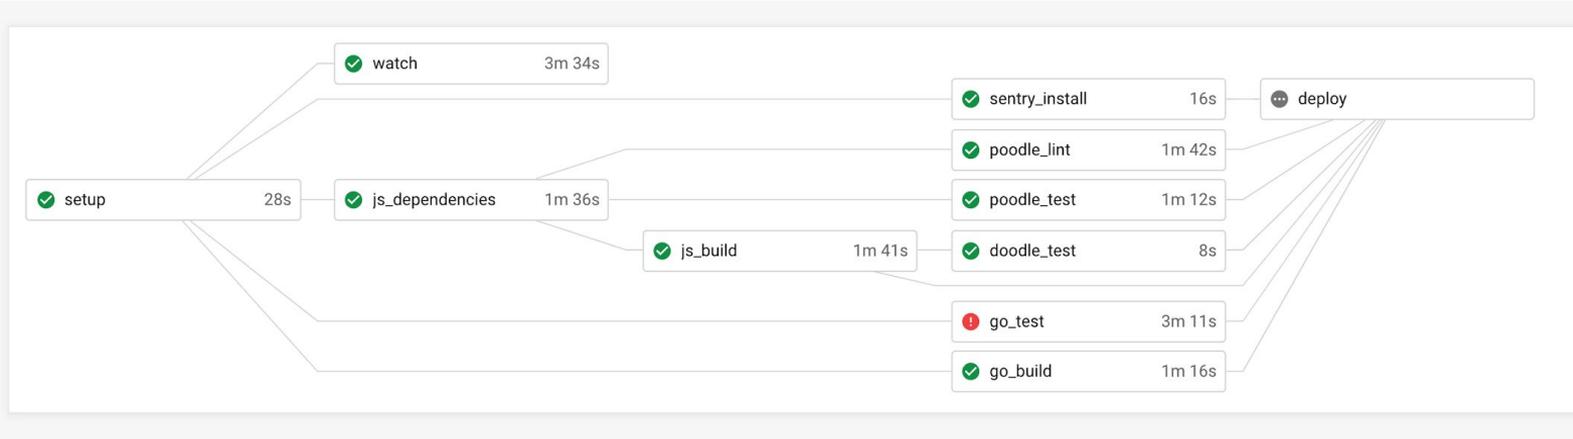
Please note that this is currently a work in progress and functionality may change over time.

Old Experience

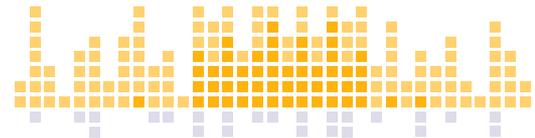
Duration 4m 9s  
Finished 3 days ago

2c6cf25

travis.limit-cleanup  
tredman



## CircleCI's view of Honeycomb build & deploy



Trace 0448bb12-2420-4e70-b677-addea31a5efe at 2020-01-15 11:41:58

Search spans

Fields

name	service_name	0s	200s	400s	600s	800s	940.4s	
10 watch 0448bb12-2...677-addea31a5efe	watch	940.420s						
start_trace	step	1.323s						
7 go_test	step	194.503s						
native-deps	cmd	2.571s						
wait-for-mysql	cmd	5.000ms						
database_conflicts	cmd	24.00ms						
honeycomb-migrations	cmd	1.003s						
retriever-migrations	cmd	16.00ms						
slos-migrations	cmd	12.00ms						
go_test	cmd	183.132s						
3 js_dependencies	step	487.830s						
native-deps	cmd	1.828s						
poodle-yarn	cmd	331.280s						
doodle-yarn	cmd	364.0ms						
1 go_build	step	68.675s						
go_install	cmd	22.175s						
sentry_install	step	2.415s						
2 poodle_test	step	117.752s						
install-yarn	cmd	1.219s						
poodle-test	cmd	40.233s						

watch >  
 watch 0448bb12-2420-4e70-b677-addea31a5efe

Distribution of span duration



Fields

Filter fields and values in span

Timestamp  
 2020-01-15T19:41:58Z

branch  
 toshok.measure-datatable-render

build\_num  
 62405

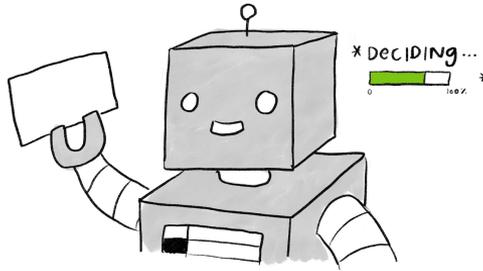
build\_url  
 https://circleci.com/gh/honeycombio/hound/62405

ci\_provider  
 CircleCI

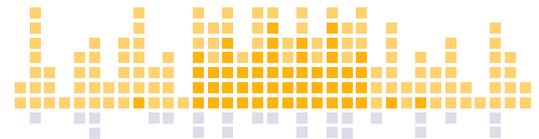
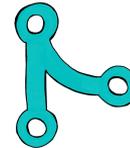
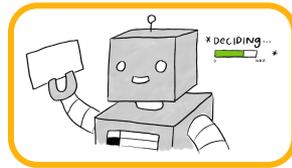
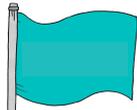
duration\_ms  
 940420

job\_name  
 watch

meta.version

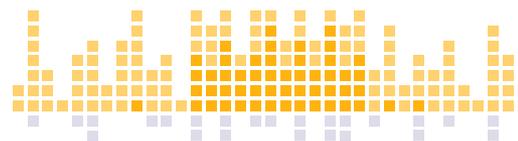
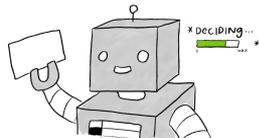
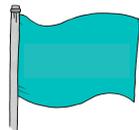


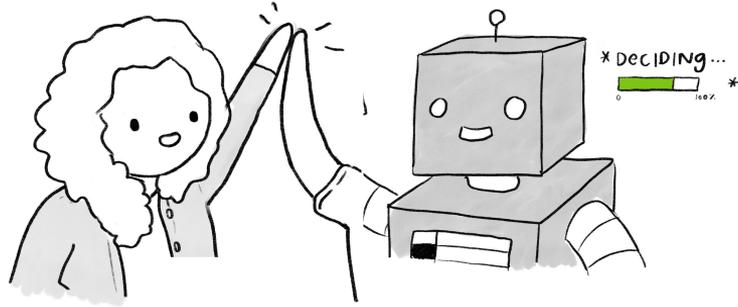
# Automated integration.



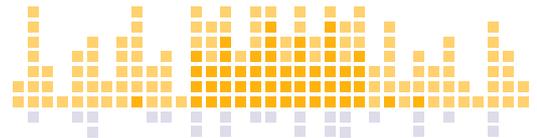
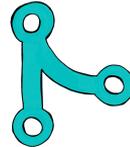
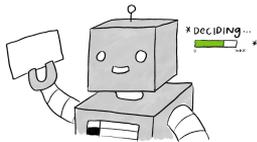
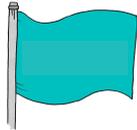


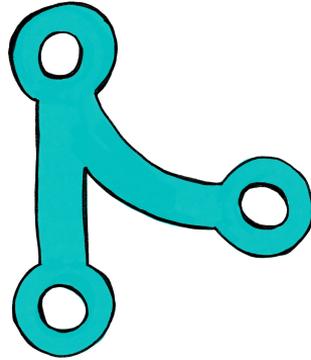
# Human PR review.



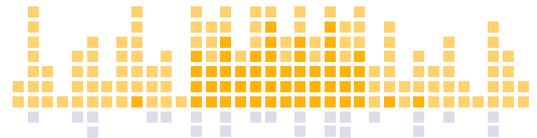
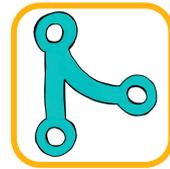
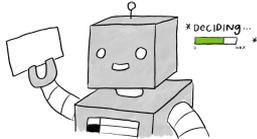
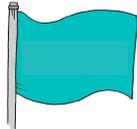


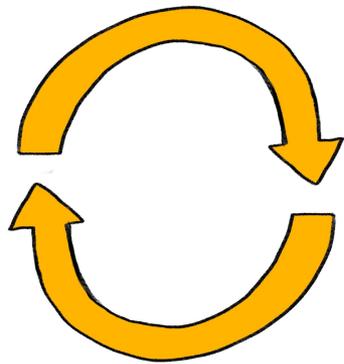
# Automated integration.



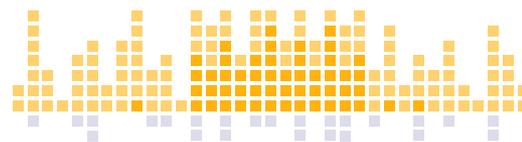
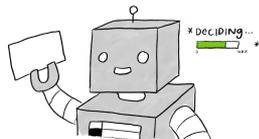
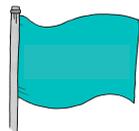


# Green button merge.



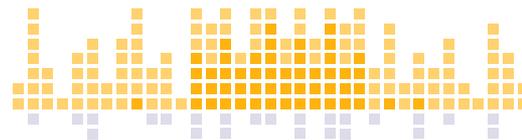
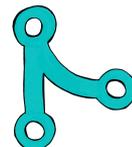
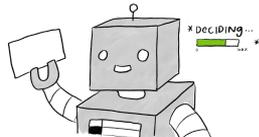
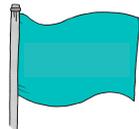


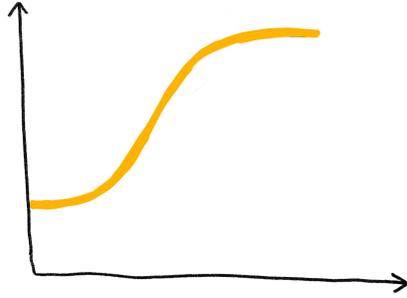
# Auto-updates, rollbacks, & pins.



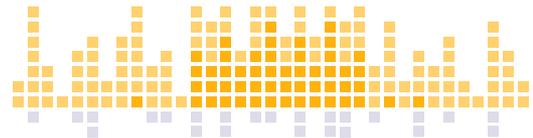
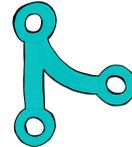
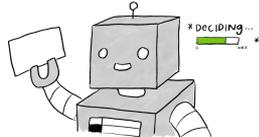
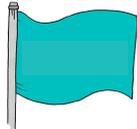


# Observe behavior in prod.



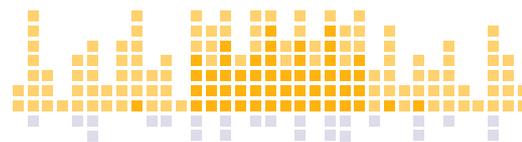
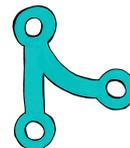
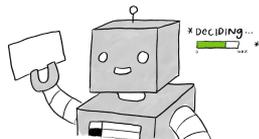
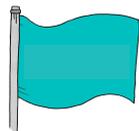


Prod: customers **observe** data.





# Dogfood **observes** prod.



# SLO Adoption: Page Views

Add to Board

Share

Last 1 month 4 days 7 hours 57 minutes 42 seconds

Datasets / poodle

Add a description for this query

VISUALIZE  
COUNT

WHERE  
request.url exists  
AND  
request.url contains /slo/

GROUP BY  
None; combine all rows

ORDER BY  
COUNT desc

LIMIT  
None

Run Query

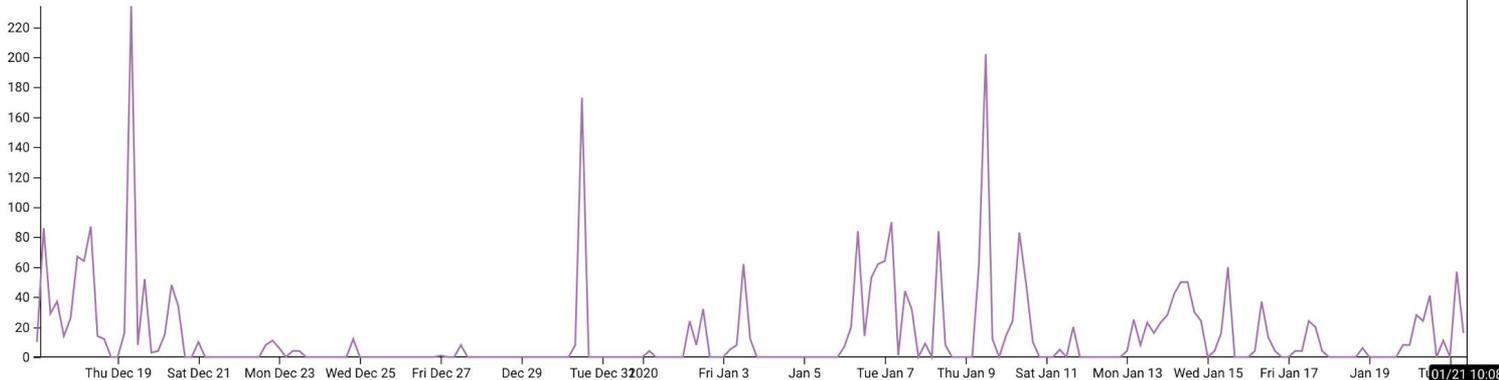
Run 2 minutes ago

Dec 17 2019, 2:10 AM - Jan 21 2020, 10:08 AM Granularity: 4 hr

Results BubbleUp Traces Raw Data

Graph Settings

COUNT

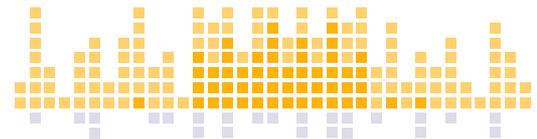
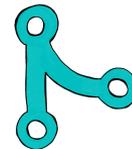
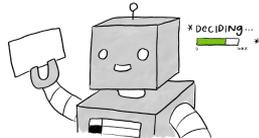
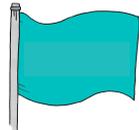


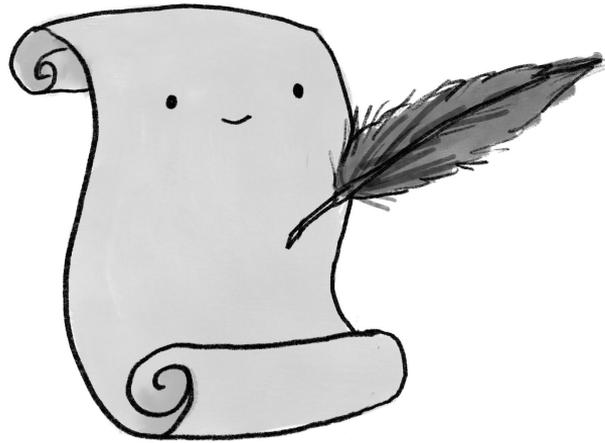
COUNT
3,029



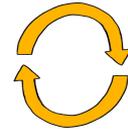
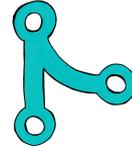
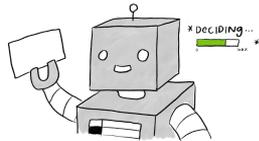
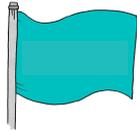


# Kibble **observes** dogfood.

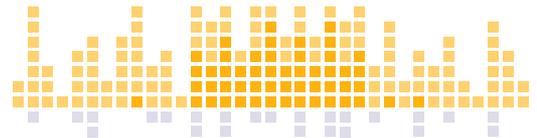




# That's how 12 eng **deploy** 12x/day!



@shelbyspees at #SKILupDay

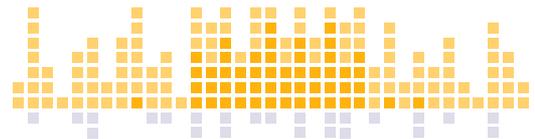




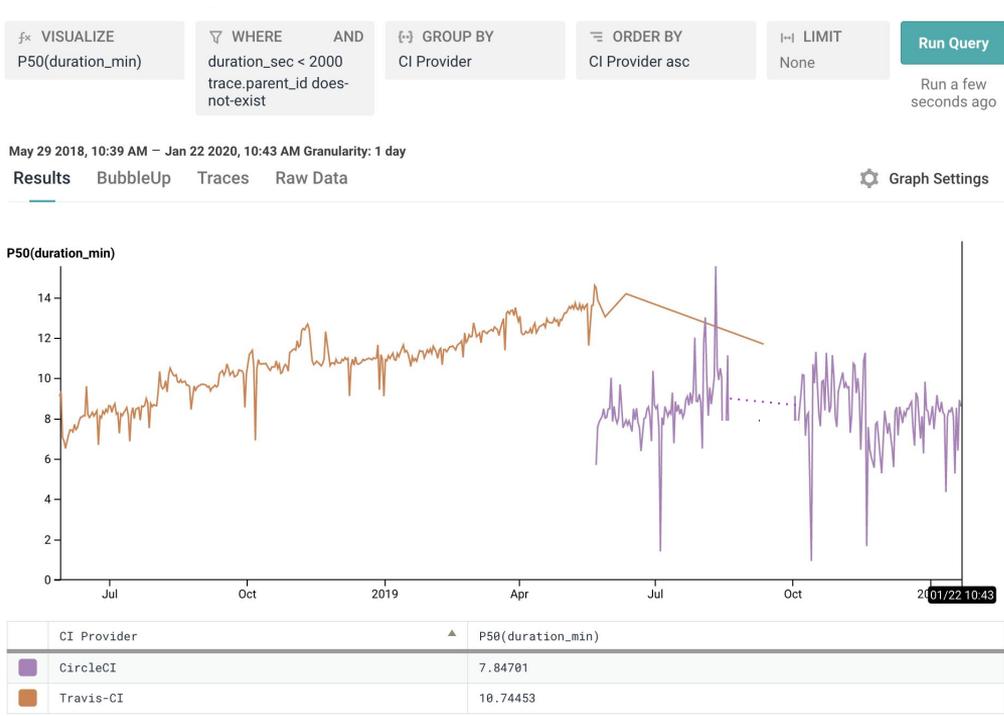
**DORA data describes feedback loops.**



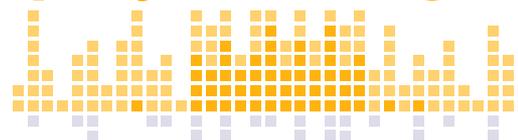
@shelbyspees at #SKILupDay



# start with lead time.



10 min builds (x3 at worst), 1h for peer review,  
hourly push train = **3 hours to deploy a change.**



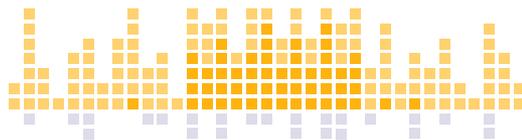
start with lead time.

deploy frequency goes up.

```
# update the deployable version at 19 minutes past the hour everywhere
# deploy the services every hour. stagger kibble, dogfood, and prod by
# at 15min each, 20, 35, and 50 minutes past the hour
# currently overridden for labrador (static site)
default["deploy"]["cron_minute"]["production"] = "50"
default["deploy"]["cron_minute"]["dogfood"] = "35"
default["deploy"]["cron_minute"]["kibble"] = "20"
```

```
# run the job that sets the deployable version at 19 past, one minute before the
# deploys will start running in kibble. NOTE this value is used in the
# `hound_util` cookbook, not here.
default["deploy"]["update_deploy"] = "19"
```

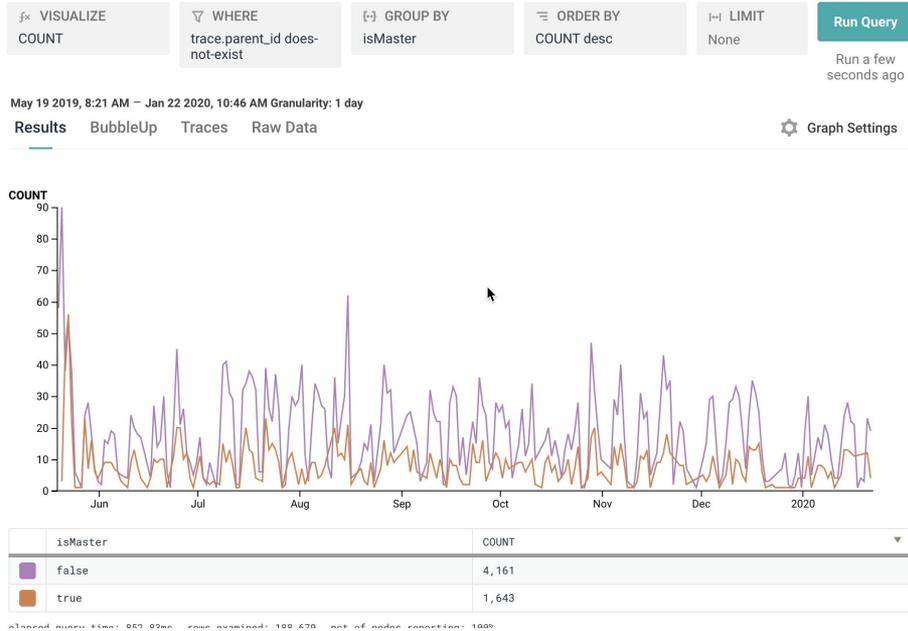
Builds go out **every hour** if there's a change.  
1-2 new commits per build artifact.



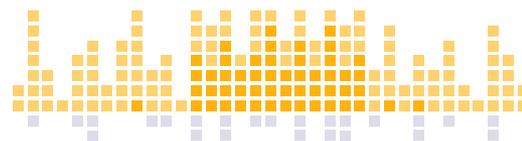
start with lead time.

deploy frequency goes up.

change fail rate goes down.



Increased confidence via testing. Flag-flip or fix-forward, not emergency rollback. **0.1% fail rate.**



start with lead time.

deploy frequency goes up.

change fail rate goes down.

time to restore goes down.

Flag flip takes **30 seconds**. Rollback to previous build takes **<10 min**. Fix-forward takes **20 min**.

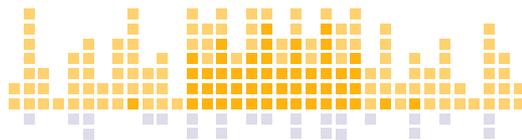
- [Trace Overview](#)

Added 10 months ago — Last requested 1 minute ago

If true, the new Trace Overview display in the Trace Inspector is displayed.

trace-overview

ON



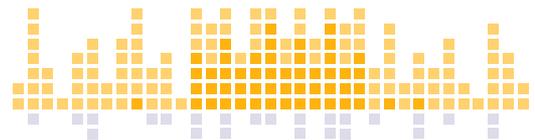
# High productivity product engineering:

start with lead time. (<3 hours)

deploy frequency goes up. (hourly, >12x/day)

change fail rate goes down. (<0.1%)

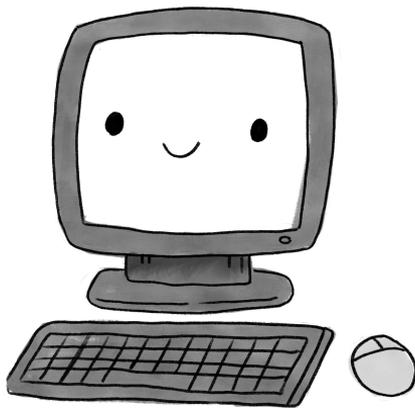
time to restore goes down. (seconds to minutes)





But what about **infra**?

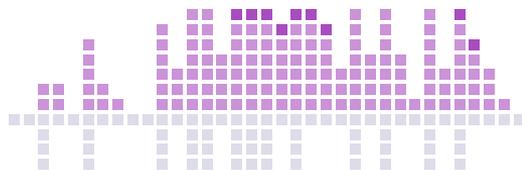




**Infrastructure empowers product.**



@shelbyspees at #SKILupDay

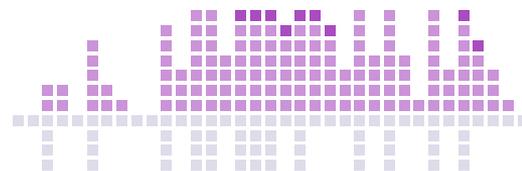


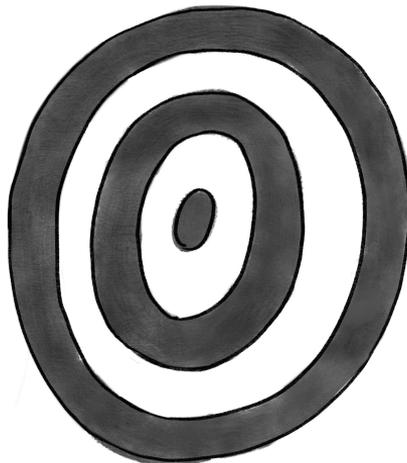


**Kubernetes isn't the goal.**



@shelbyspees at #SKILupDay

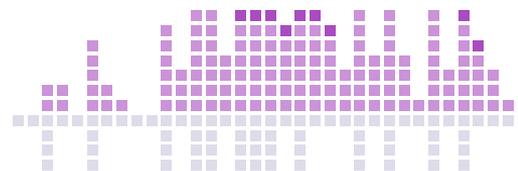




**Reliability and simplicity is.**



@shelbyspees at #SKILupDay

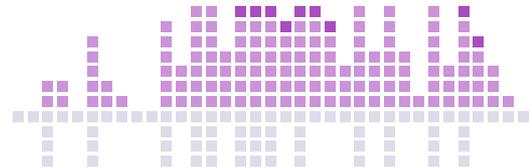


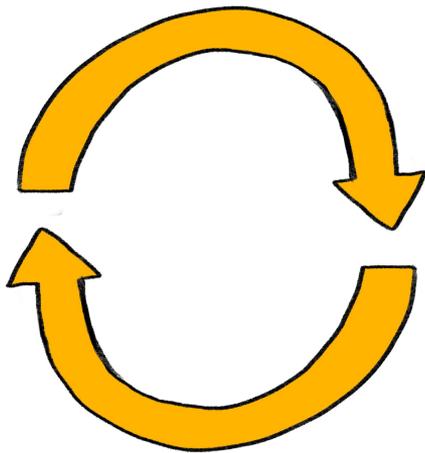


Everyone starts somewhere.



@shelbyspees at #SKILupDay

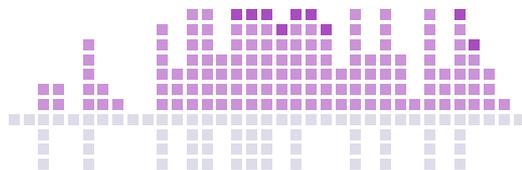


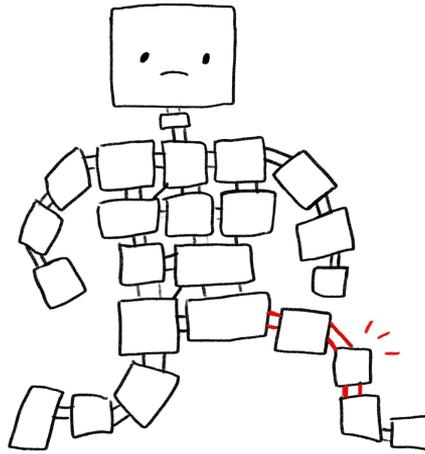


**Automate** the painful parts.

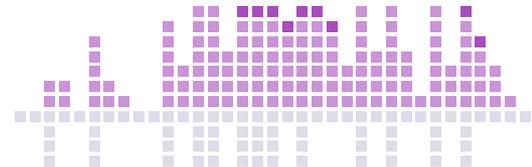


@shelbyspees at #SKILupDay





**Fix the duct tape!**

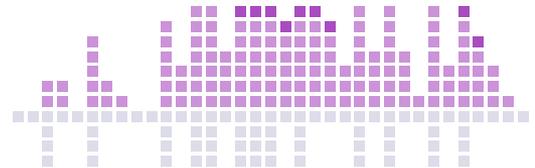


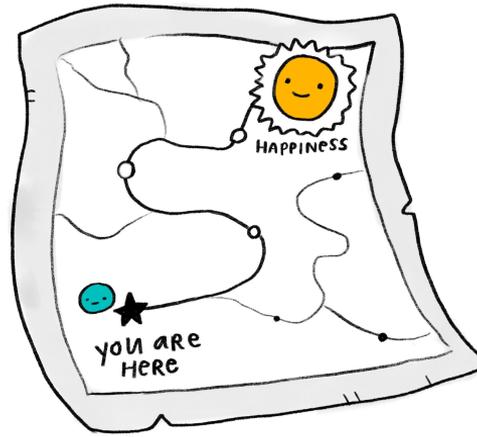


Keep the environment clean.

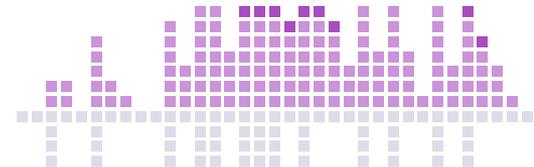


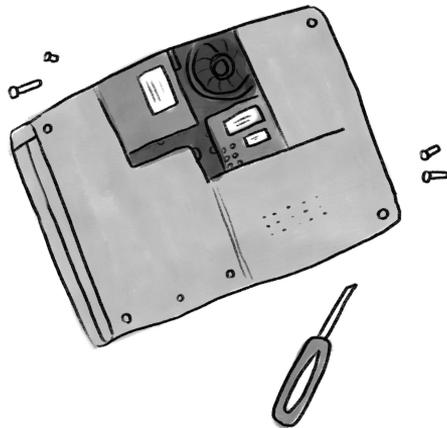
@shelbyspees at #SKILupDay



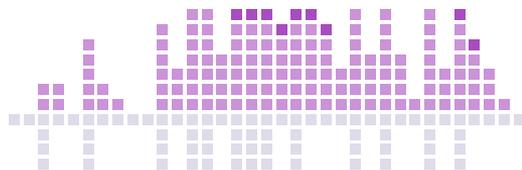


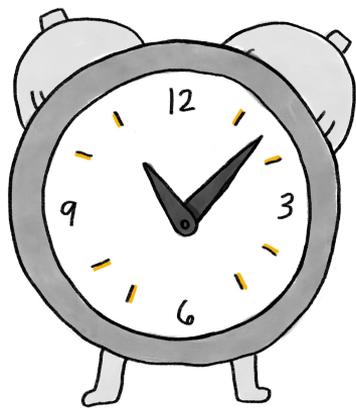
# How do we do it?



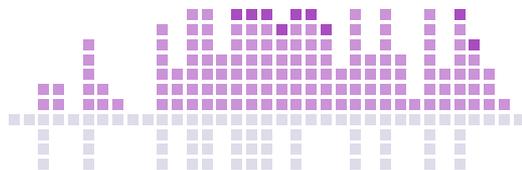


Raw VMs are **simpler** than containers.



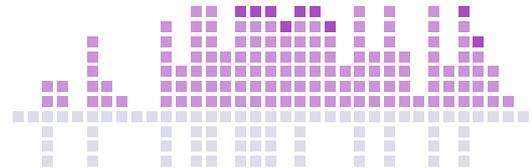


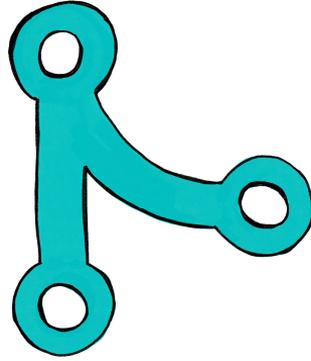
**Cold boot from Chef, use cron to sync.**



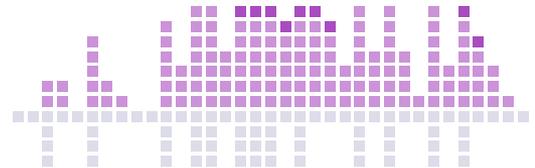


**Outsource utilities like blob storage.**



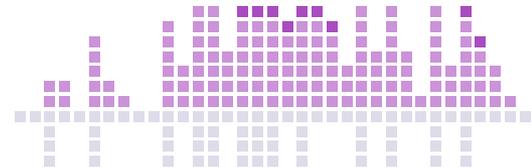


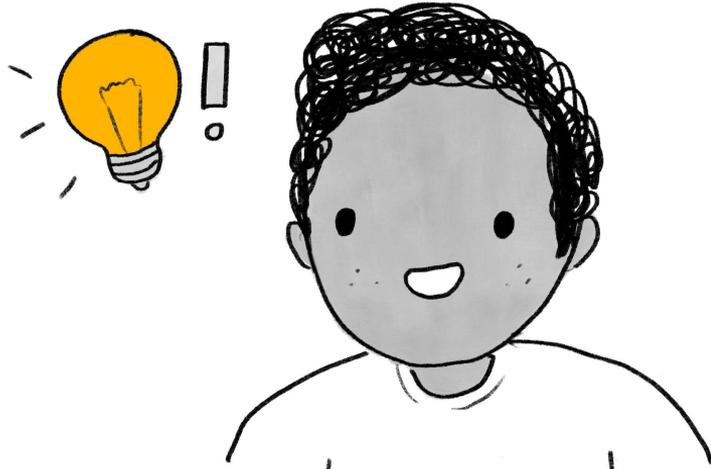
**Repeatable** infrastructure with code.





# Centralize state & locking.

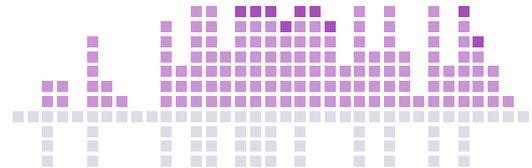




**Remove barriers to setup.**

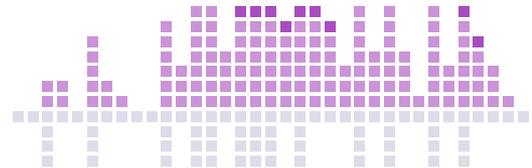


@shelbyspees at #SKILupDay





# Diff and release in browser.



# Allow 20% of prod instances to run on-demand #447

Edit

Merged lizthegrey merged 1 commit into master from lizf.fix-spot 6 hours ago

Conversation 0 Commits 1 Checks 1 Files changed 1 +1 -1

all commits File filter... Jump to...

No Whitespace 0 / 1 files viewed Review changes

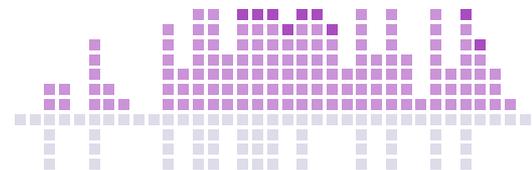
```
terraform/modules/autoscaling_groups/shepherd_asg.tf  
@@ -168,7 +168,7 @@ resource "aws_autoscaling_group" "shepherd_asg" {  
  mixed_instances_policy {  
    instances_distribution {  
      on_demand_percentage_above_base_capacity = 20  
    }  
  }  
}
```

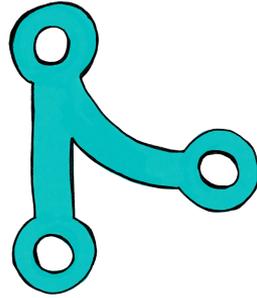
## All checks have passed

13 successful checks

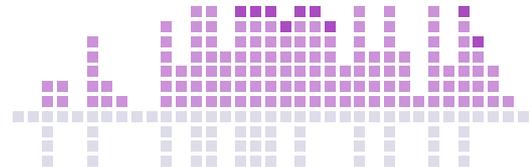
- ✓ atlas/honeycomb/kubernetes — Run not triggered: T... [Details](#)
- ✓ atlas/honeycomb/prod — Terraform plan: 0 to add, 1... [Details](#)
- ✓ check-infra — Workflow: check-infra [Details](#)
- ✓ sentinel/honeycomb/base — Policy checks passed [Details](#)

```
{  
  on_strategy = "prioritized"  
  capacity = 35  
  size_above_base_capacity = 10 -> 0  
  strategy = "lowest-price"  
  s = 2  
  ecification {  
    ...  
  }  
}
```



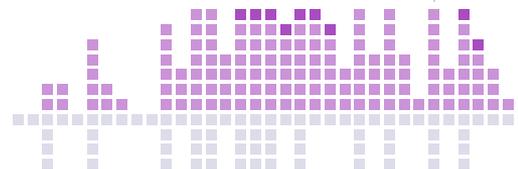


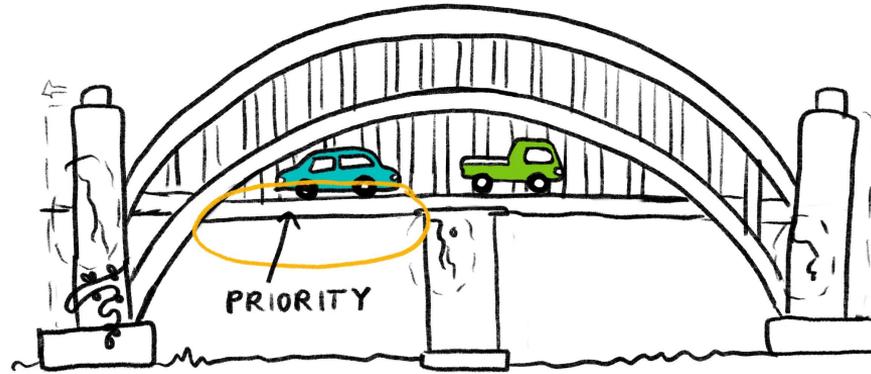
Remote **run** from git.



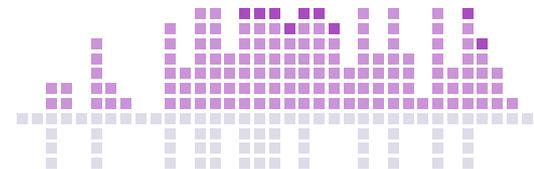
## Run List

	<b>Revert "Allow 20% of prod instances to run on-demand" (#448)</b> <span>CURRENT</span>	<span>✓ APPLIED</span>
#run-eMmN3VwovBP8NWX4   lizthegrey triggered from GitHub   Branch master   fd7fee0		
	<b>Allow 20% of prod instances to run on-demand (#447)</b>	<span>✓ APPLIED</span>
#run-Vd7miibVisRAwrB2   lizthegrey triggered from GitHub   Branch master   b80330f		
	<b>actually redirect HTTPS traffic. (#444)</b>	<span>✓ APPLIED</span>
#run-KAamLNTt51wji5Fe   lizthegrey triggered from GitHub   Branch master   944a14e		
	<b>fix remaining apply issues. (#443)</b>	<span>✓ APPLIED</span>
#run-xRdsnkWQDB4P6k4B   lizthegrey triggered from GitHub   Branch master   e7a4e14		
	<b>fix tags &amp; syntax (#442)</b>	<span>✗ ERRORED</span>
#run-Z9K3KSb71ZS9jnx5   lizthegrey triggered from GitHub   Branch master   cf1c12f		
	<b>add quarantine service. (#441)</b>	<span>✗ ERRORED</span>
#run-tSjr7wofSVGvSje7   lizthegrey triggered from GitHub   Branch master   1dca77c		
	<b>[terraform] add buckets for storing trace decisions between restarts in ...</b>	<span>✓ PLANNED</span>
#run-hj34neLhBPPCpQ3j   tredman triggered from GitHub   Branch master   9c367d0		



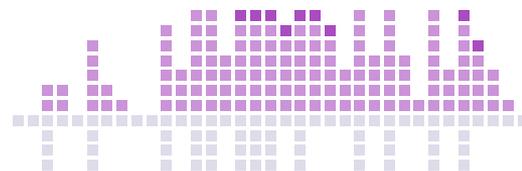


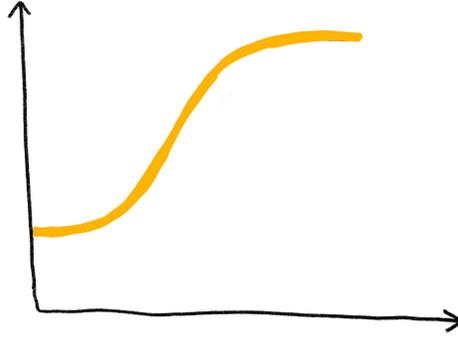
# Sentinel guardrails.





**Notify only on risky changes.**

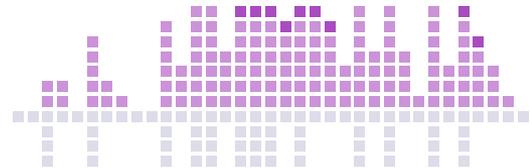


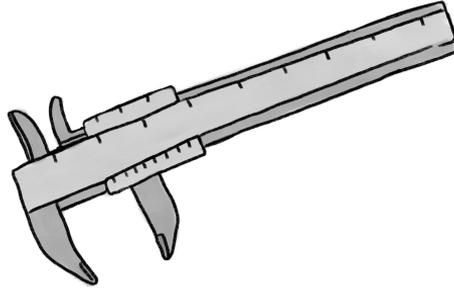


# Deploy changes **incrementally!**

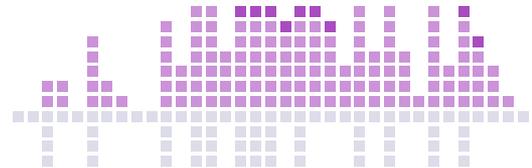


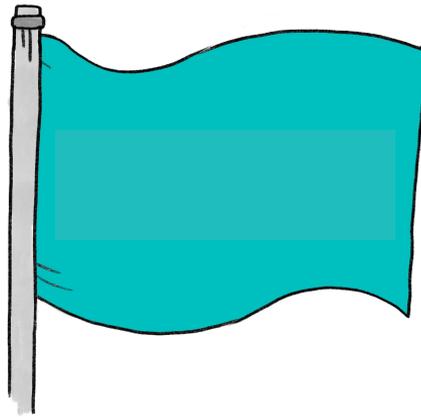
@shelbyspees at #SKILupDay



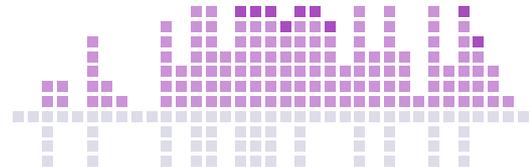


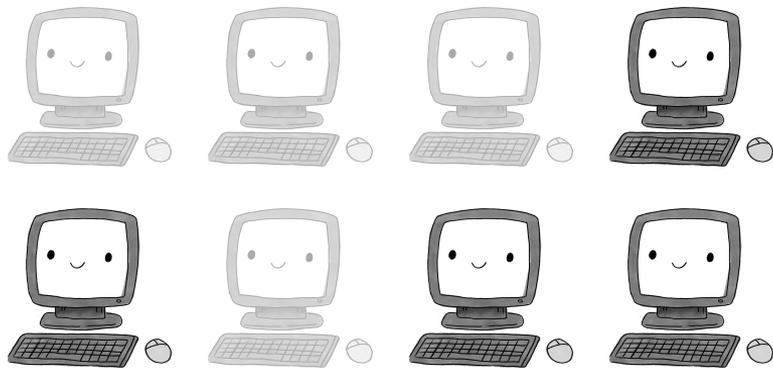
Use **modern** components.



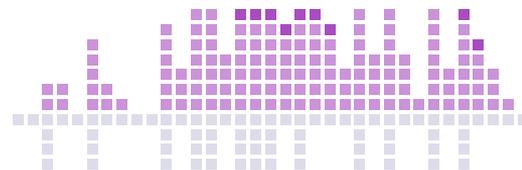


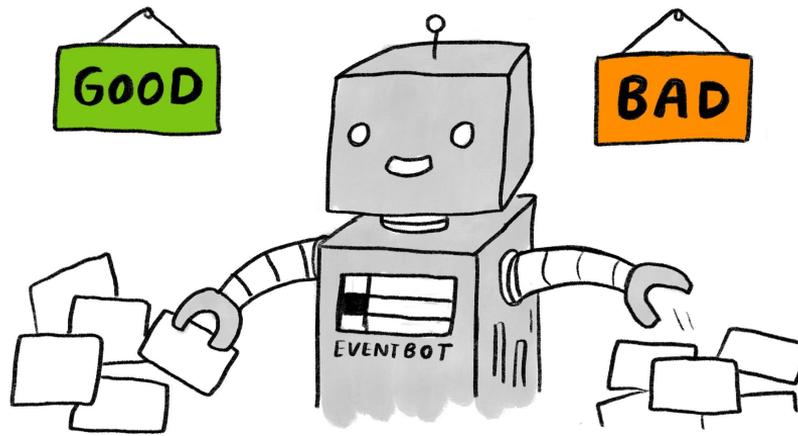
# Feature **flags**... for infra!





# Ephemeral fleets & autoscaling.

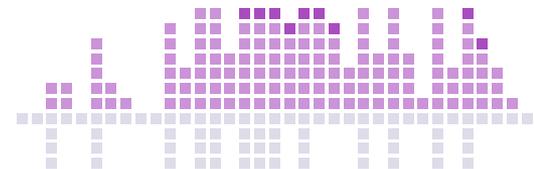




**Quarantine** bad traffic.

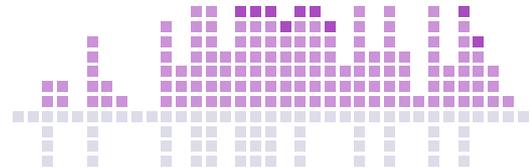


@shelbyspees at #SKILupDay



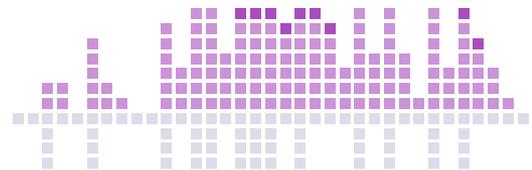


**Delete** unused code & components.





# Refactor continuously!



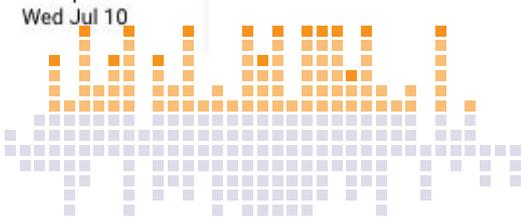
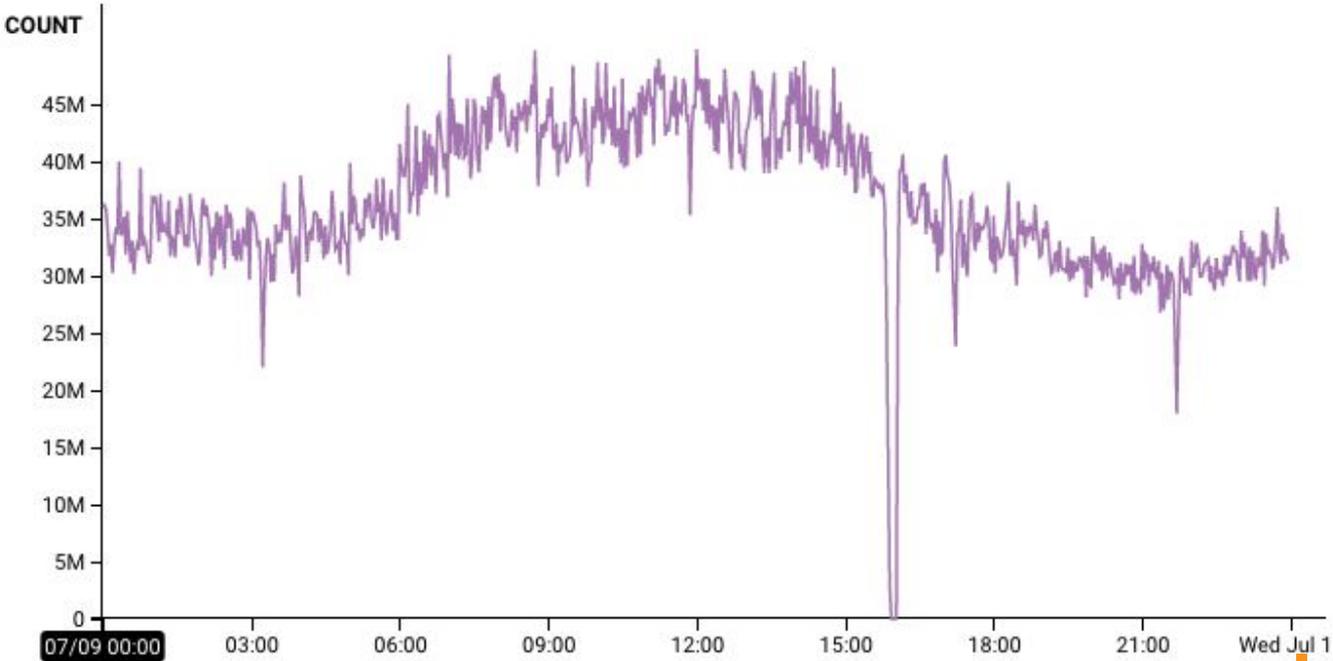


# What if it goes wrong?

Jul 9 2019, 12:00 AM – Jul 10 2019, 12:00 AM

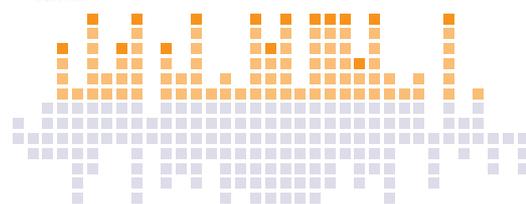
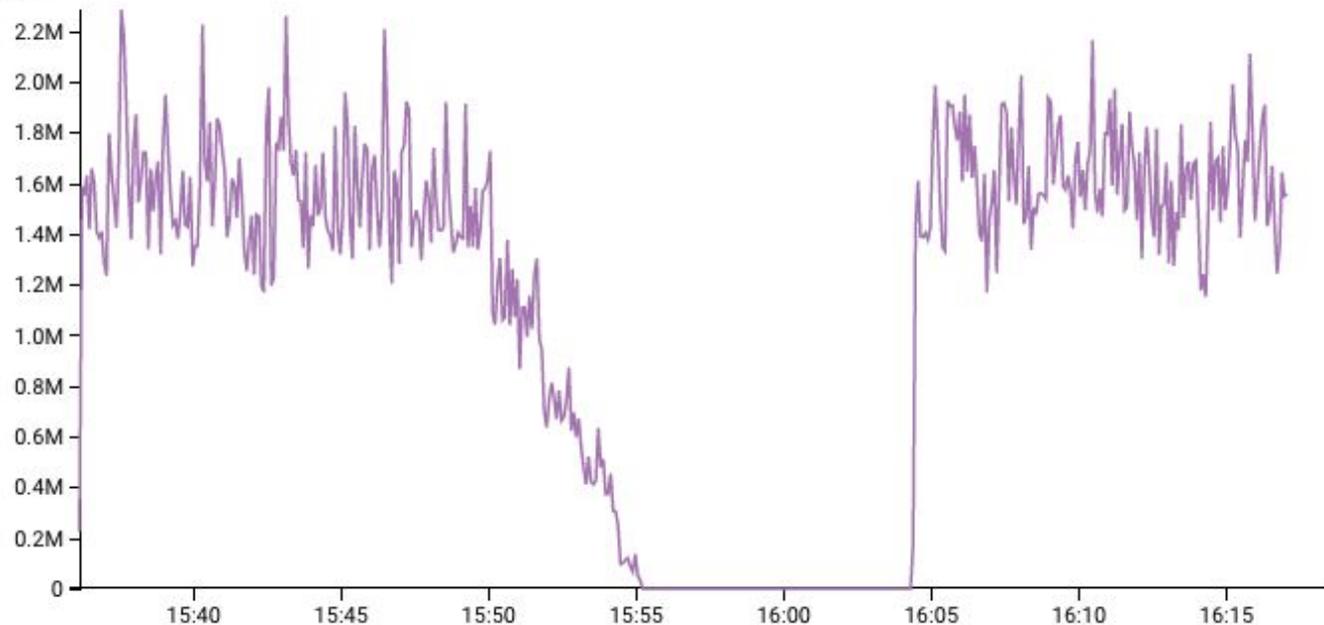
Results BubbleUp Traces Raw Data

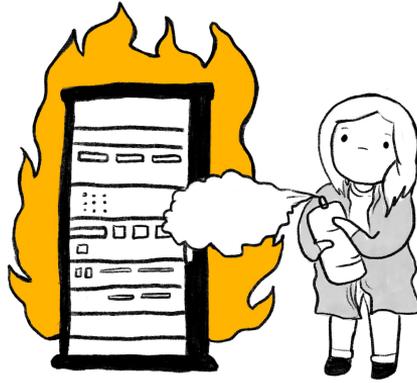
⚙ Graph Settings



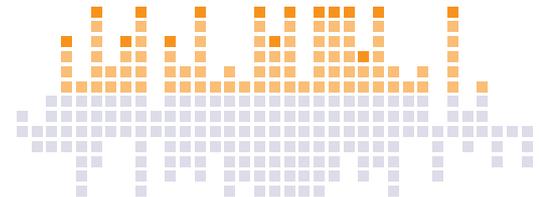
@shelbyspees at #SKILupDay

COUNT



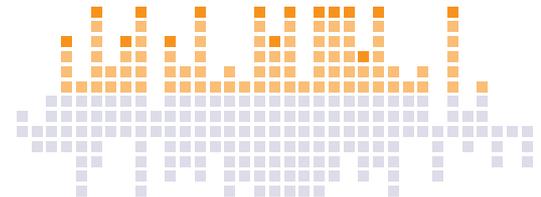


# How broken is “too broken”?



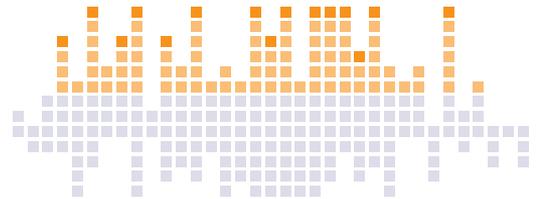


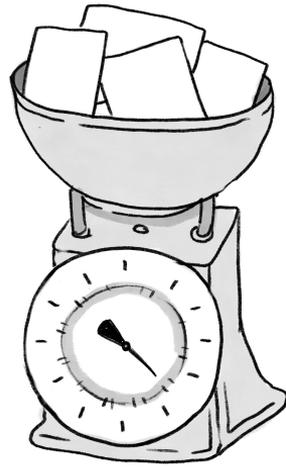
**Service Level Objectives** define success.





**SLOs are common language.**

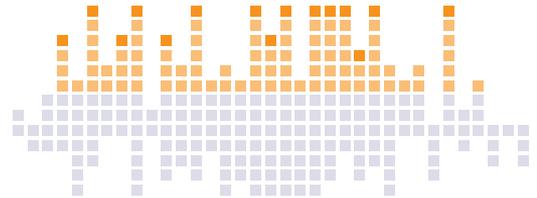


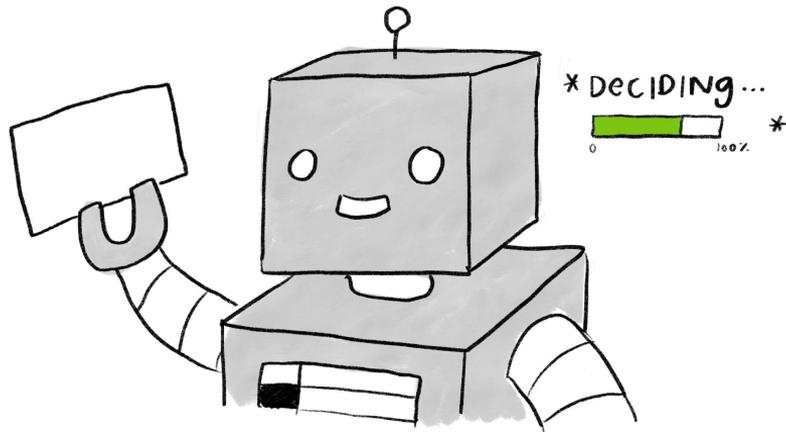


How many **eligible events** did we see?

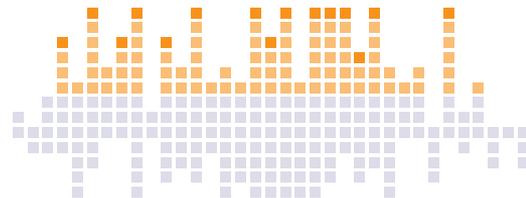


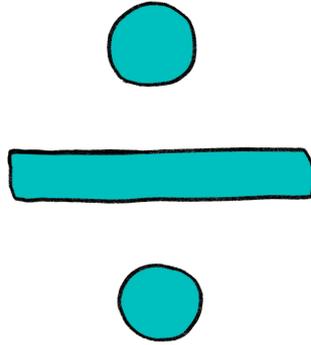
@shelbyspees at #SKILupDay



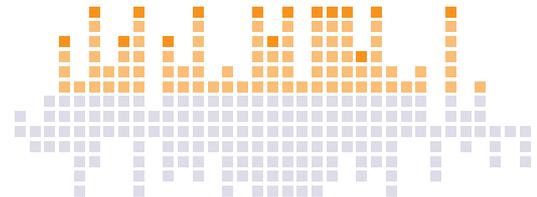


**HTTP Code 200? Latency < 100ms?**

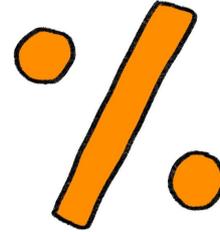




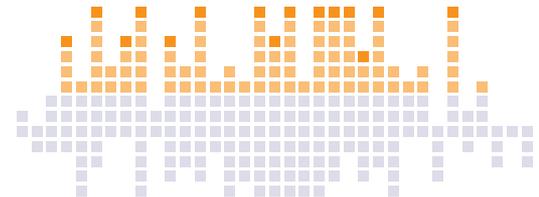
**Availability:** Good / Eligible Events



JULY						
	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31			

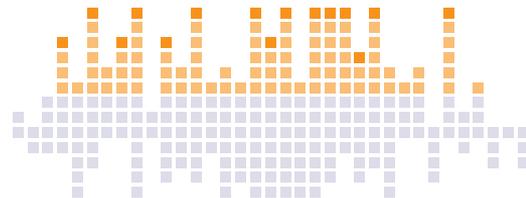


Use a **window** and **target percentage**.



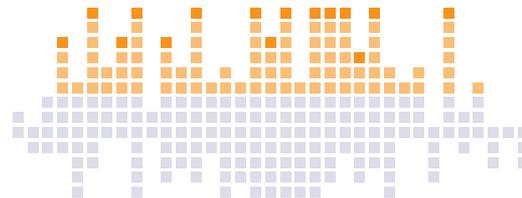
$$1 - x$$

Error budget: **allowed** unavailability



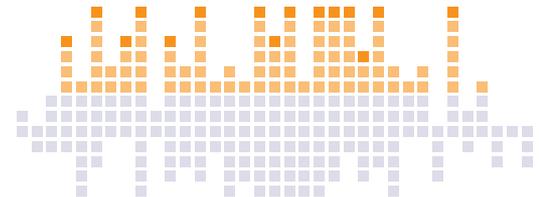


Drive **alerting** with SLOs.



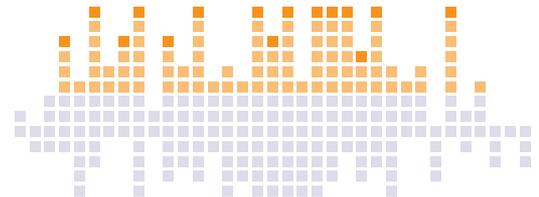


We keep **SLOs** at Honeycomb.



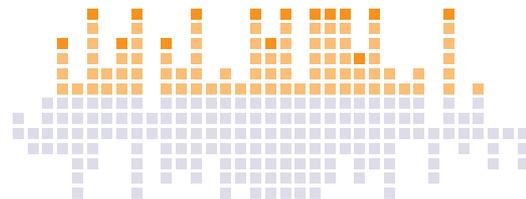
99.99%

We **store** incoming telemetry.



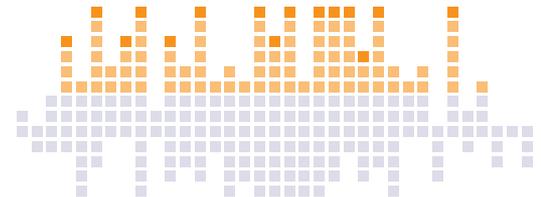
99.9%

Default dashboards usually **load** in 1s.



99%

Often, **queries** come back under 10s.

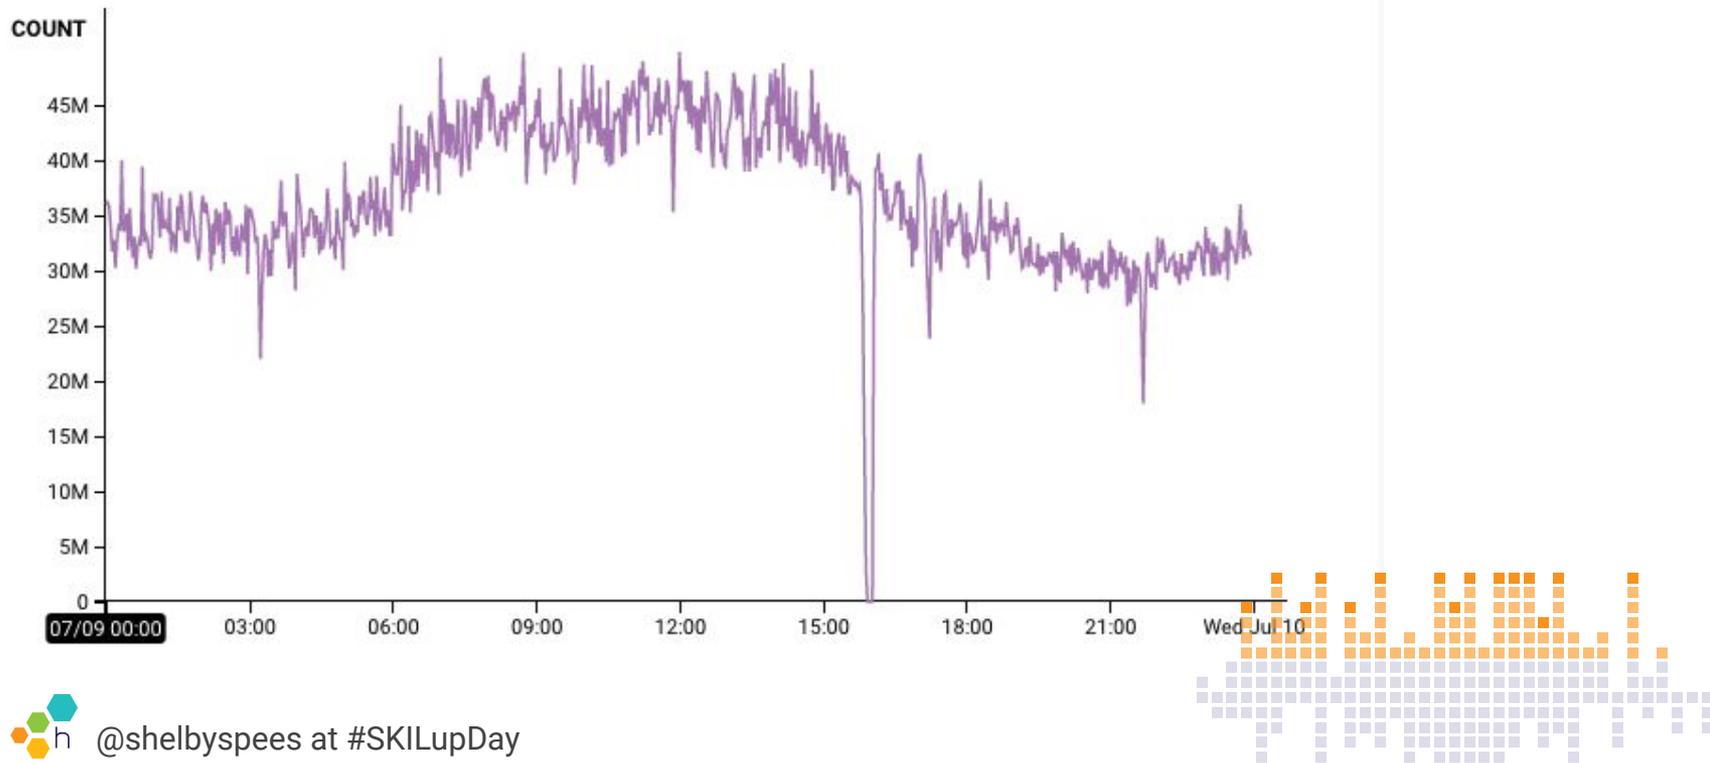


# User Data Throughput

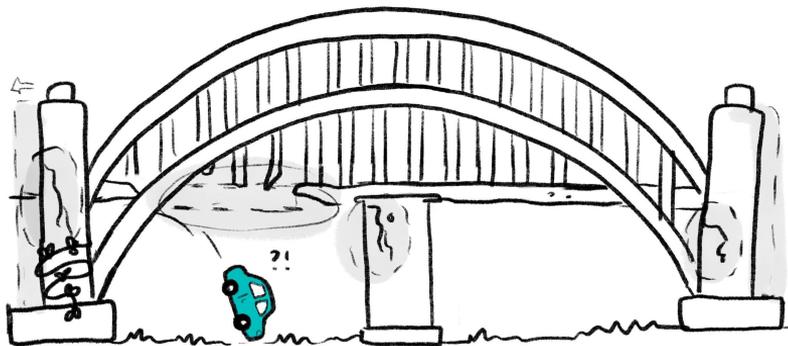
Jul 9 2019, 12:00 AM – Jul 10 2019, 12:00 AM

Results BubbleUp Traces Raw Data

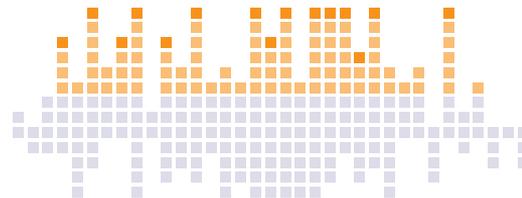
⚙ Graph Settings

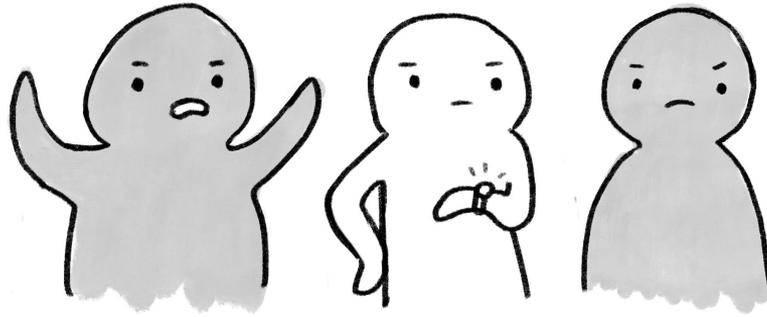


@shelbyspees at #SKILupDay

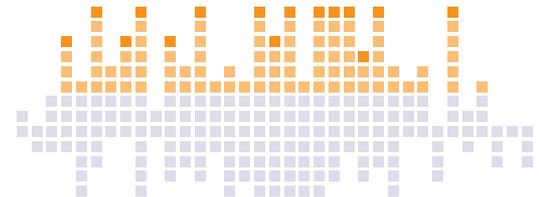


**We **dropped** customer data.**



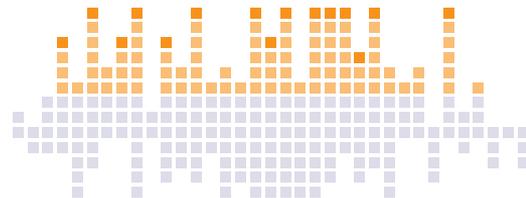


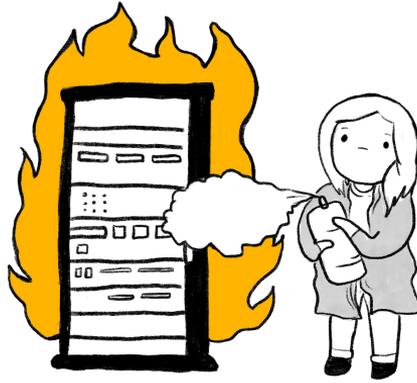
but **rolled back** (at human speed)



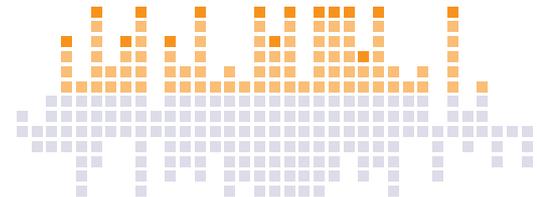


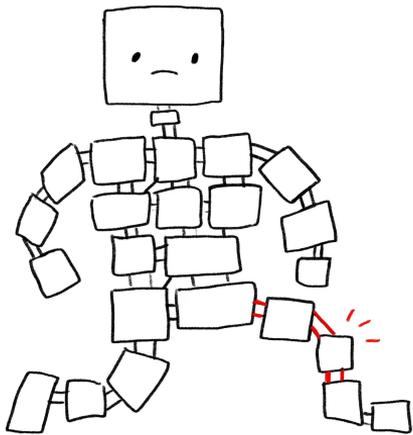
**We communicated to customers.**



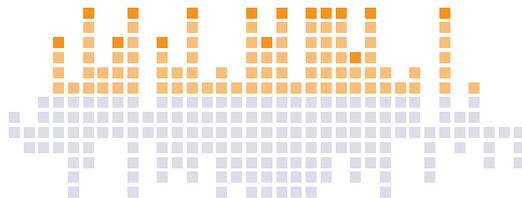


**We'd burned triple our error budget.**





**We halted deploys.**

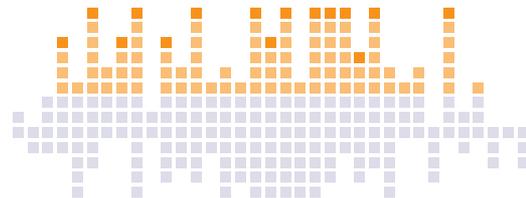




**How** did this happen?



@shelbyspees at #SKILupDay

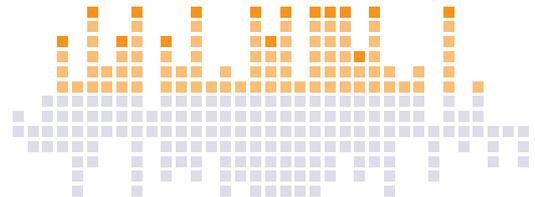


**We checked in code that **didn't build**.**

**We had **experimental CI** build wiring.**

**Our scripts deployed **empty binaries**.**

**There was no **health check & rollback**.**

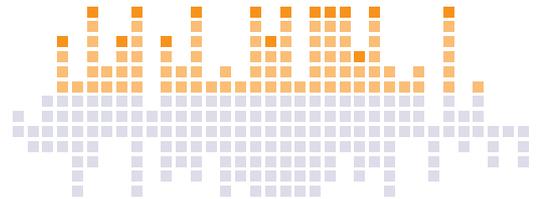


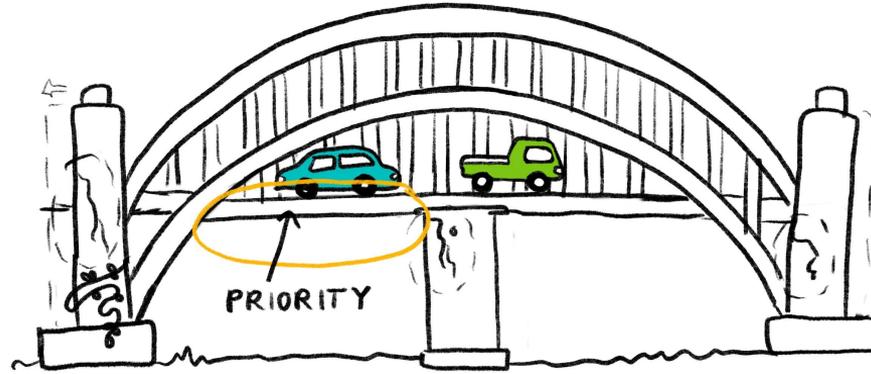


**We re-prioritized stability.**

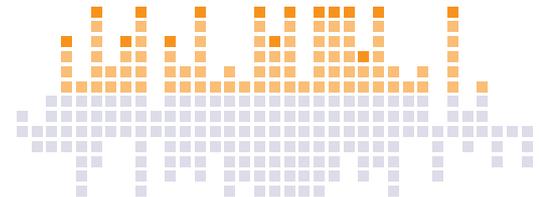


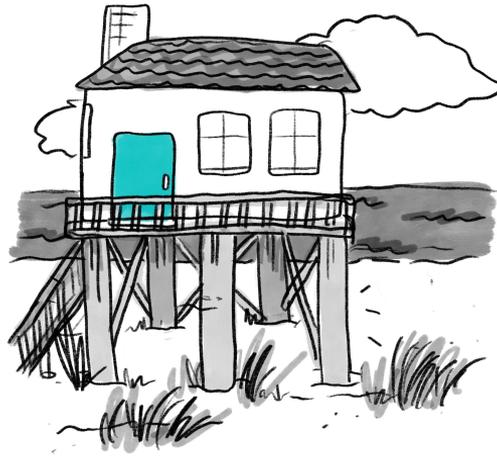
@shelbyspees at #SKILupDay





We **mitigated** the key risks,

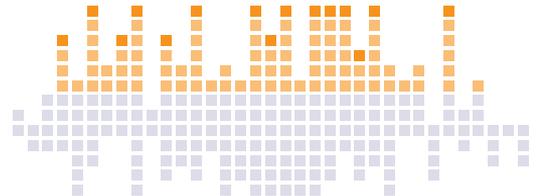




then resumed **building.**



@shelbyspees at #SKILupDay





# What's ahead for us?

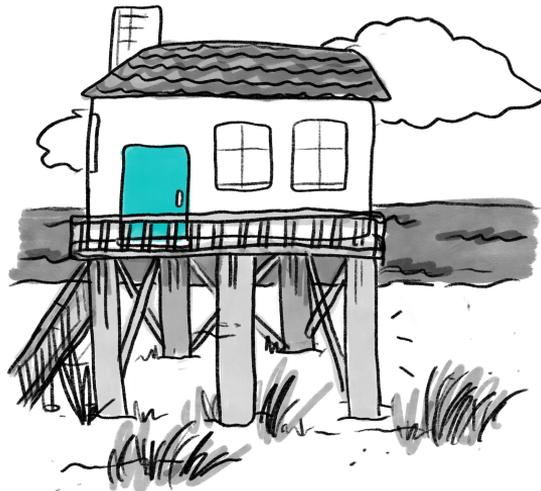


Be more **reliable & scalable.**



@shelbyspees at #SKILupDay

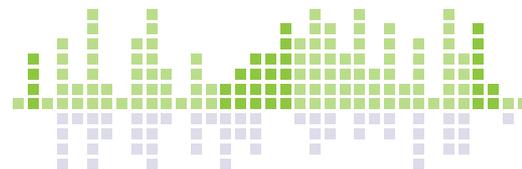


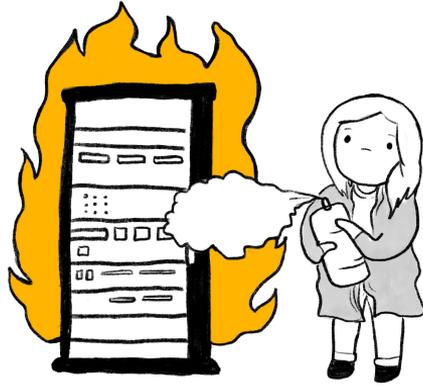


Launch **new services** easily.

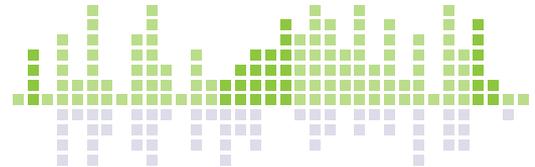


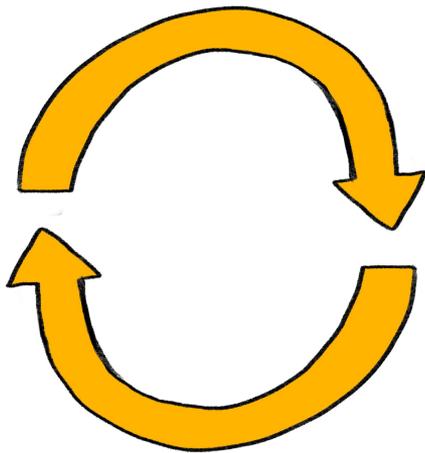
@shelbyspees at #SKILupDay



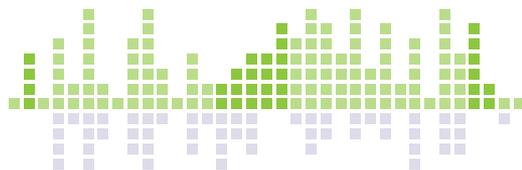


**Burn less money.**



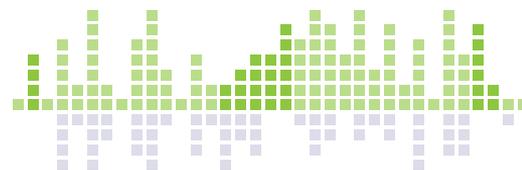


**Continue modernizing & refactoring.**



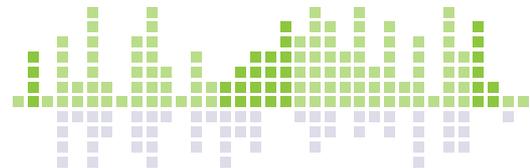


**Sleep easily at night.**





**You can do this too, step by step.**





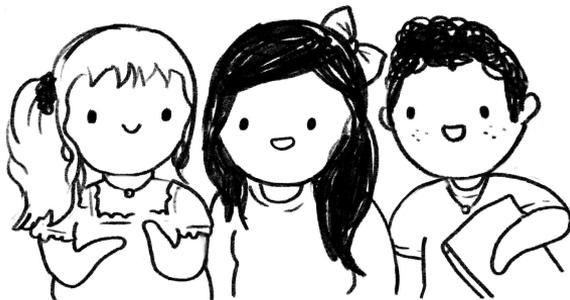
Read our blog! [hny.co/blog](https://hny.co/blog)



@shelbyspees at #SKILupDay



# Understand & control production.



**Go faster on stable infra.**  
**Manage risk and iterate.**

[shelbyspees.com](http://shelbyspees.com)  
[@shelbyspees](https://twitter.com/shelbyspees)

