


# Lessons from Automatic Incident Resolution for a Million Databases

---

SRECon EU, July 2016


Greg Burek



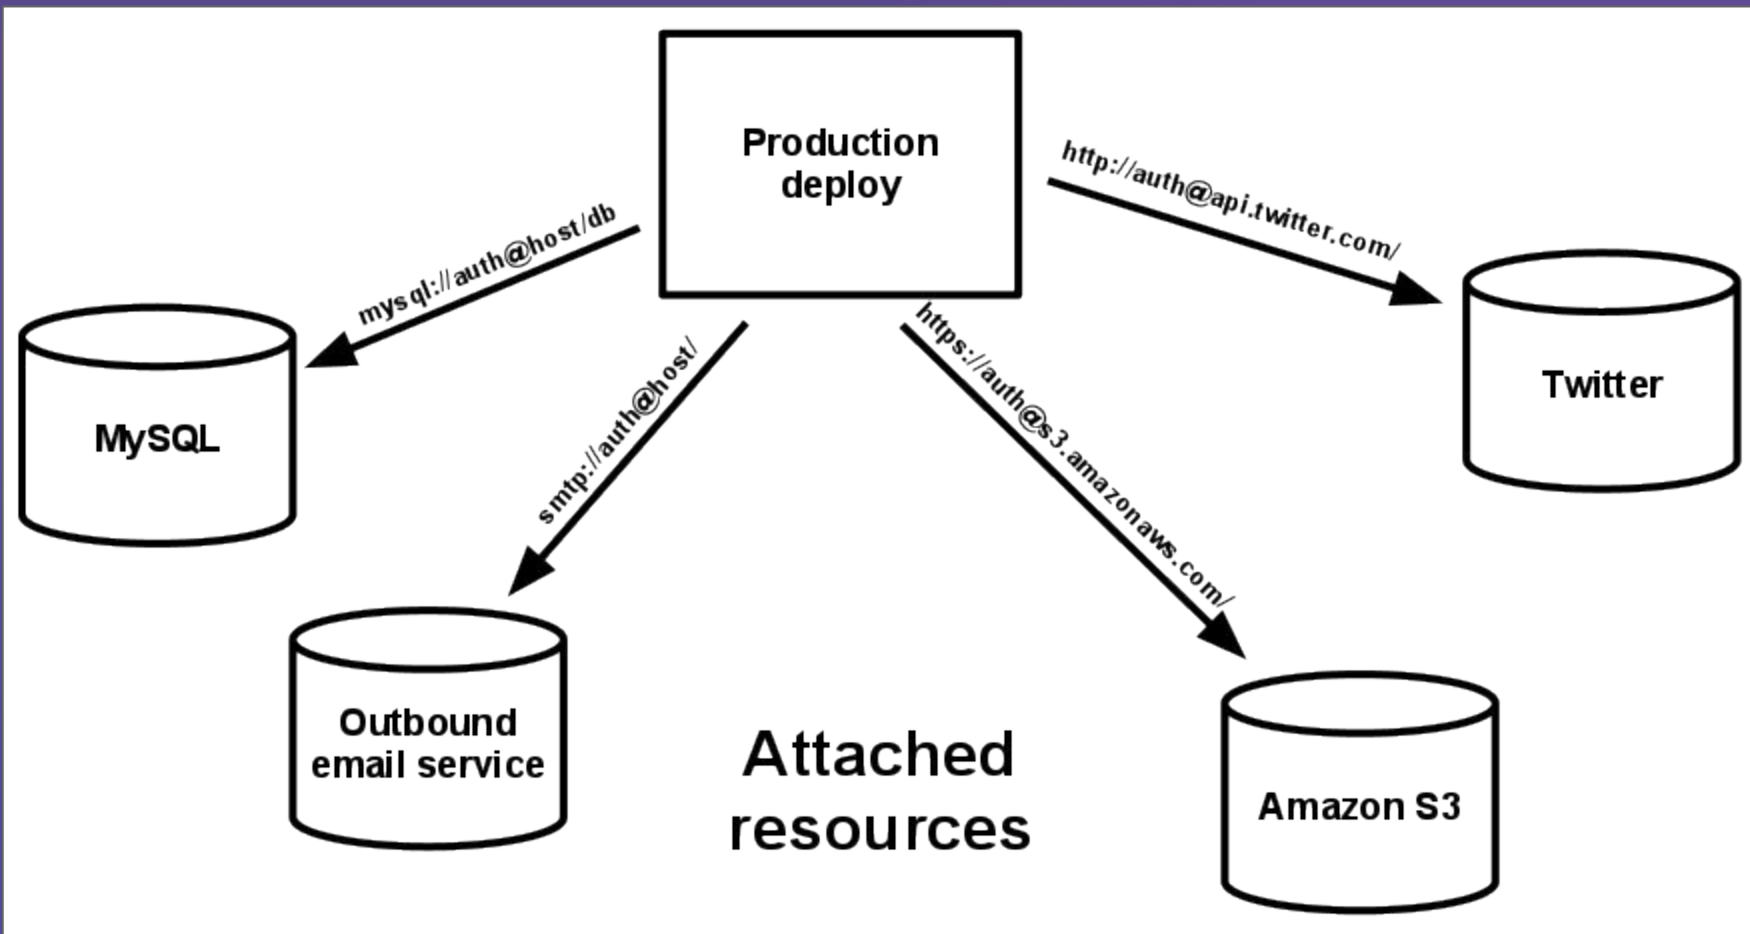



**HEROKU**

# The Twelve-Factor App




**HEROKU**





**HEROKU**

Department of Data



**HEROKU**

Postgresql

Redis

Kafka

~ Million Databases

Tens of thousands of AWS Instances



Some Databases

Hundreds of AWS Instances



“The goal is to build systems that can scale linearly with machines & sub-linearly with people” - Caitie McCaffrey

Tackling Alert Fatigue

Monitor and alert on your  
business



**Alice Goldfuss**

@alicegoldfuss

Me: oh sorry, I got paged

Date: are you a doctor?

Me: uh

Nagios: holy SHIT this cert expires in SIXTY DAYS

Me: ...yes

Dec 28, 2015, 6:23 PM

---

**842** RETWEETS

**1,478** LIKES

---

Monitor and alert on your  
business

Usually, don't alert on machine  
specific metrics

Write runbooks and playbooks

Turn playbooks into code





**Sam Kottler**

@samkottler



Stop writing playbooks and start writing scripts. If the steps to respond to an alert can be written down they can also be automated.

Jul 5, 2016, 1:58 PM

---

**26** RETWEETS

**36** LIKES

---

“The goal is not to never get paged, the goal is to never get paged for the same thing twice”  
- Astrid Atkinson

Engineering for the long game

Verify monitoring before  
restarting the world



# Circuit breakers

Automation can't handle the  
unknown

Wake someone up on  
exceptions and timeouts



Have a REPL/console



Aggregate and review trends

## Trusty Kernel Panics in Last Week

No Rows Returned

## Stale Incidents (over 2 days old)

type	state	count
old_wale_base_backup	await_resolution	1
resource_down	ignored	1

## Total Incidents

type	count
follower_behind	1
old_wale_base_backup	1
blocking_queries	1
resource_down	1
changeover	1
server_down	1

Humans can break

Automation can be simplistic

Humans + Automation for a  
resilient and operable system

1. Monitor and alert on your business
2. Write playbooks
3. Make playbooks into automation
4. Checks and balances of automation
5. Circuit breakers
6. Alert on exceptions and timeouts
7. Admin console
8. Aggregate and review trends



gregburek@heroku.com  
@gregburek

# State Machines