



# Arrested Development

The awkward adolescence of a microservices-based application

Europython 2015  
Scott Triglia



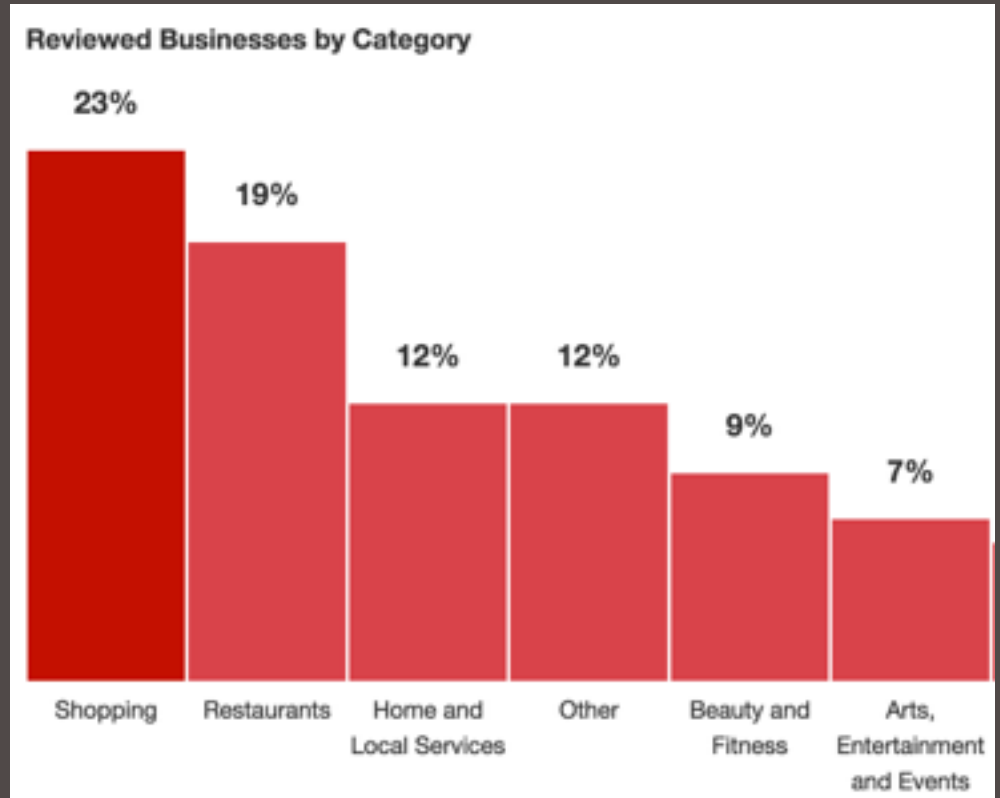
The Company





77M reviews

142M monthly unique users



Your Speaker

4 years with Yelp  
Search, ML, Services

Scott Triglia  
@scott\_triglia



Yelp  
Transaction  
Platform



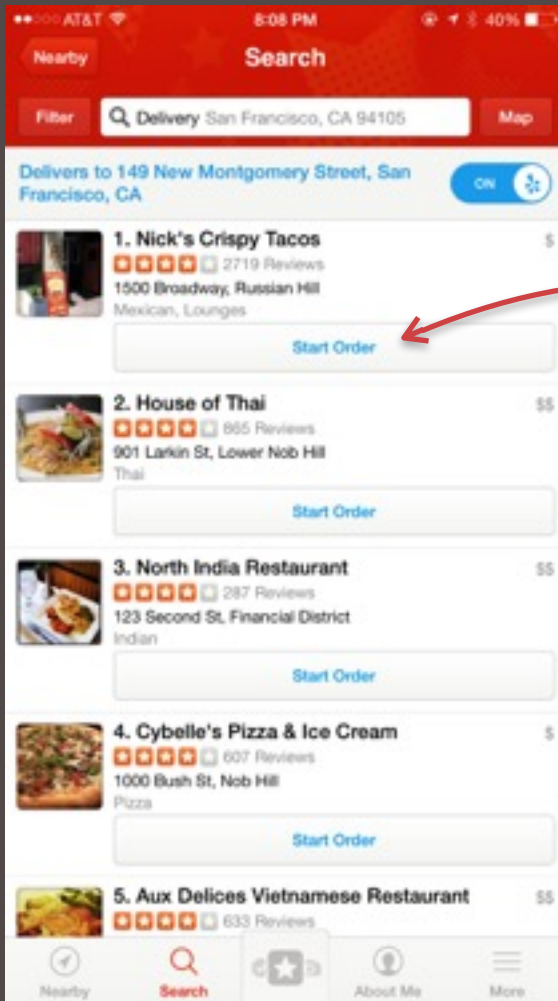
The Product

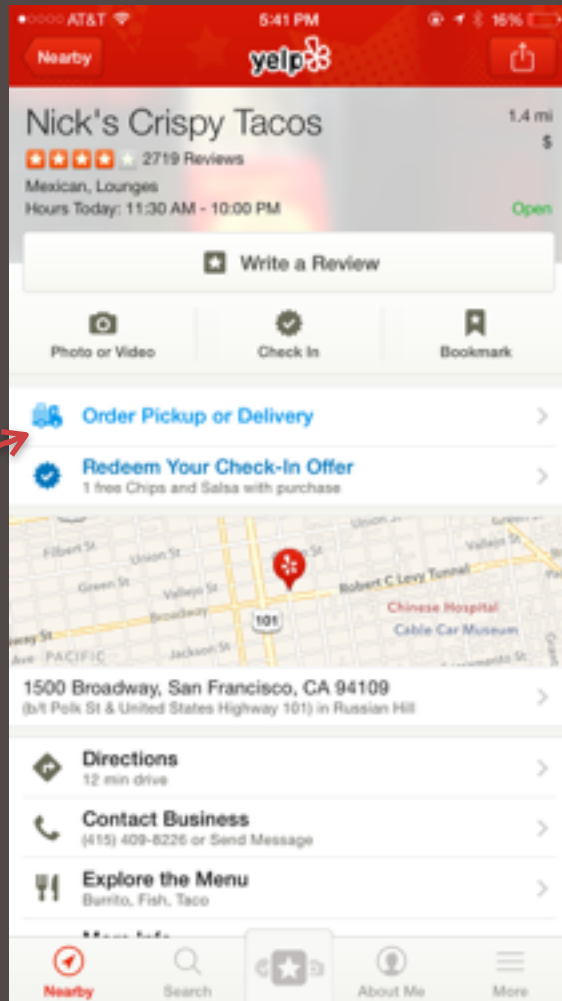
Yelp  
Transaction  
Platform



(or just “Platform”)

The Product







AT&T8:08 PM40%

CancelMake Your Selection

Nick's Crispy Tacos

Delivery \$5.00 • Min \$9.00


Estimated Delivery Time 45-75 min

ORDERING BY EAT24

Full Menu

Menu hours: 10:30 AM to 9:00 PM


What's Good



Pescado Taco

Baja style fried fish in a corn tortilla with salsa, lime mayonnaise, cilantro, cabbage and onions.


\$4.95



Carne Asada Taco


Grilled steak in corn tortilla with pico de gallo and pinto beans.

\$3.95



Chips and Salsa

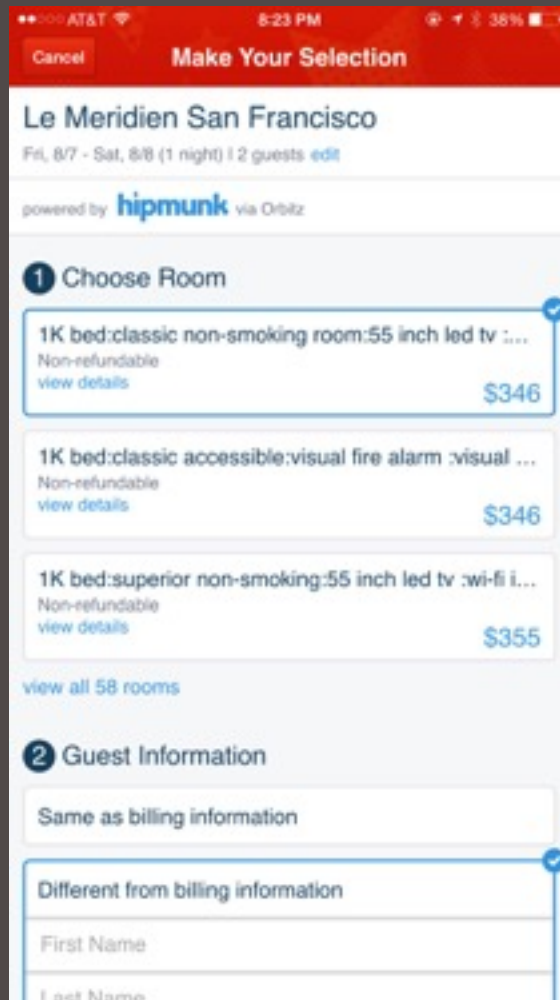
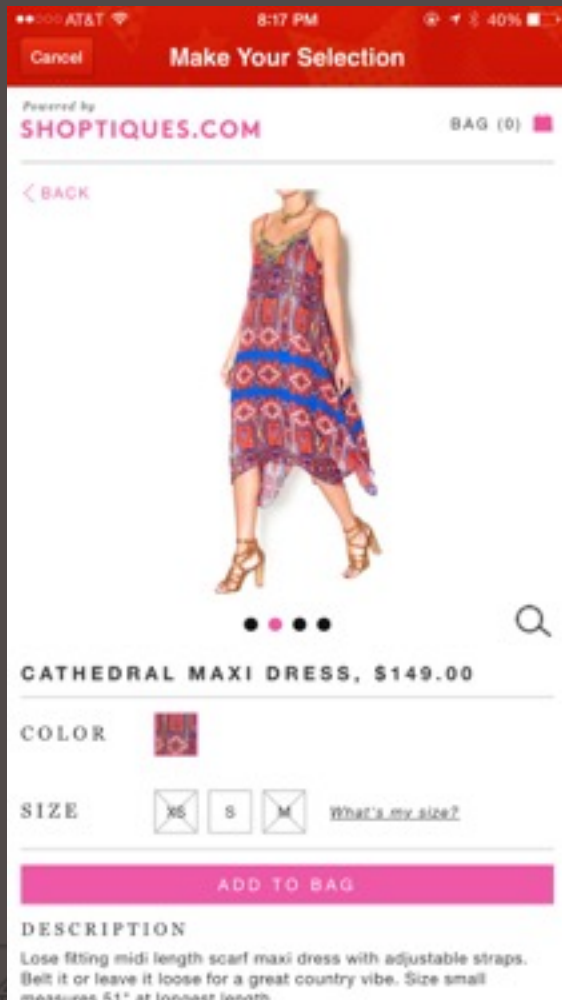
\$1.25



Nachos

Our house made chips topped with a traditional cheese sauce, black beans,

\$7.95



AT&T

\$48 PM

13%

CancelSummary

Pickup

Nick's Crispy Tacos

Use coupon

Carne Asada Taco

\$3.95 x 2 = \$7.90

Subtotal

\$7.90

Tax

\$0.69

SF Healthy Mandate

\$0.32

Total

\$8.91

Edit

Your Contact Information

Phone Number

555-4379125

Payment Information

Use VISA \*\*\*\*-6444

>

By proceeding, you agree to our [Terms of Service](#) and [Privacy Policy](#). We'll send your transaction to a partner company to fulfill.

Complete Purchase

# Microservices

That Hot Trend



“...an approach to developing a single application as **a suite of small services**, each running in its **own process** and communicating with **lightweight mechanisms**...”

<http://martinfowler.com/articles/microservices.html>





(clarkmaxwell via Flickr; CC BY-NC-ND 2.0)

Monolithic python code  
resisted decoupling



Monolithic python code catered to  
the lowest common denominator

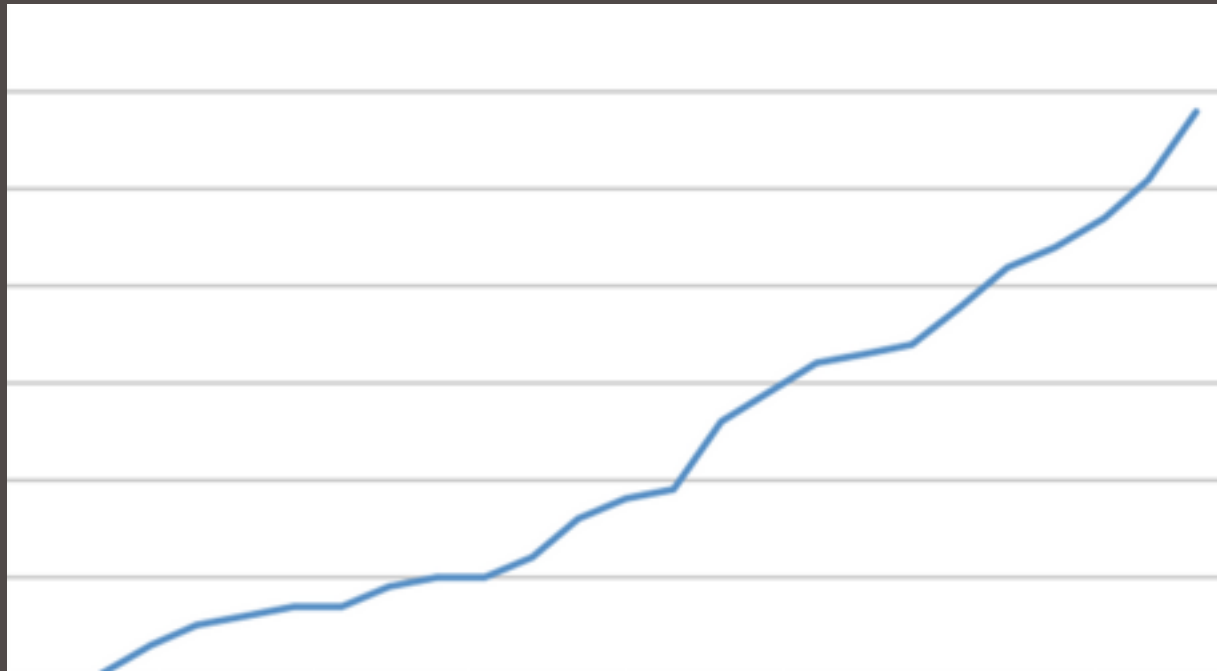




Monolithic python code **was**  
**anti-agile**



Services



Time





Pinterest Gingerbread House



Pinterest Gingerbread House

# API complexity increases



coupling rises



interactions get murky



process does not scale





# So what's an engineer to do?



- Decoupling
- Defining
- Understanding Production
- Staying Agile



- Decoupling
- Defining
- Understanding Production
- Staying Agile



# Old boring problem

## Monolithic spaghetti code



# Solution: microservices!



# New exciting problem

how to share concepts  
across services



New exciting problem

distributed tech debt



# service\_type





# service\_type

What product does your business provide and how do they provide it?



# service\_type

pickup

delivery



booking\_at\_customer  
service\_type

pickup

booking\_at\_business

delivery



booking\_at\_customer

service\_type

hotel\_reservation

goods\_at\_business

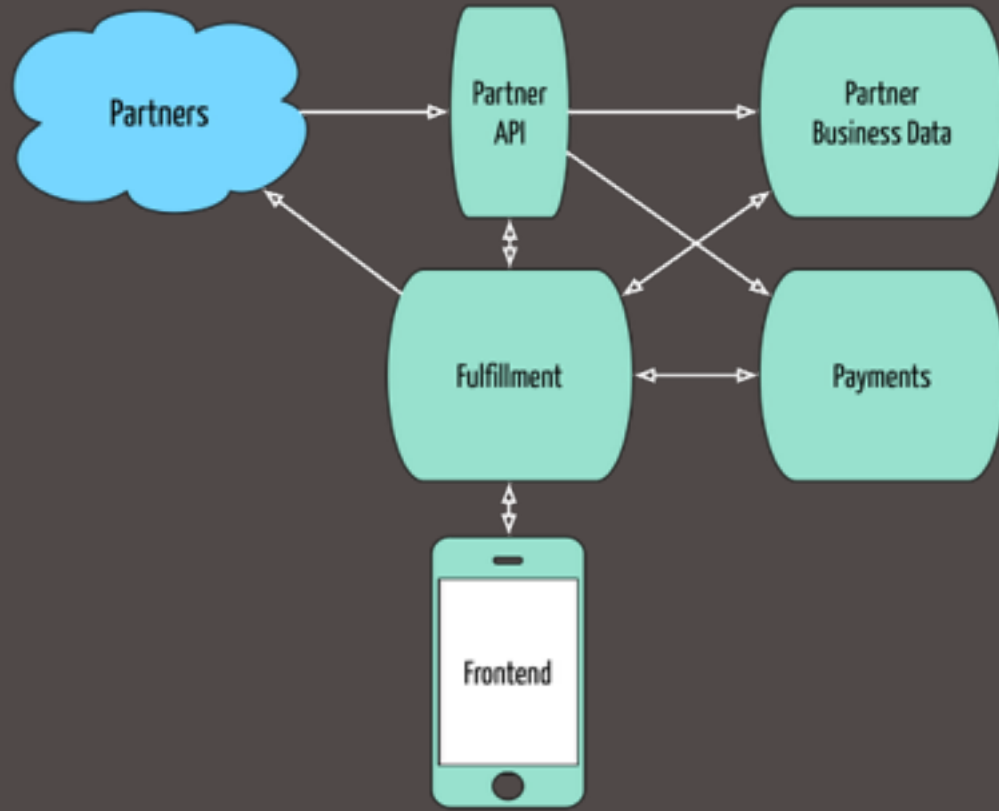
pickup

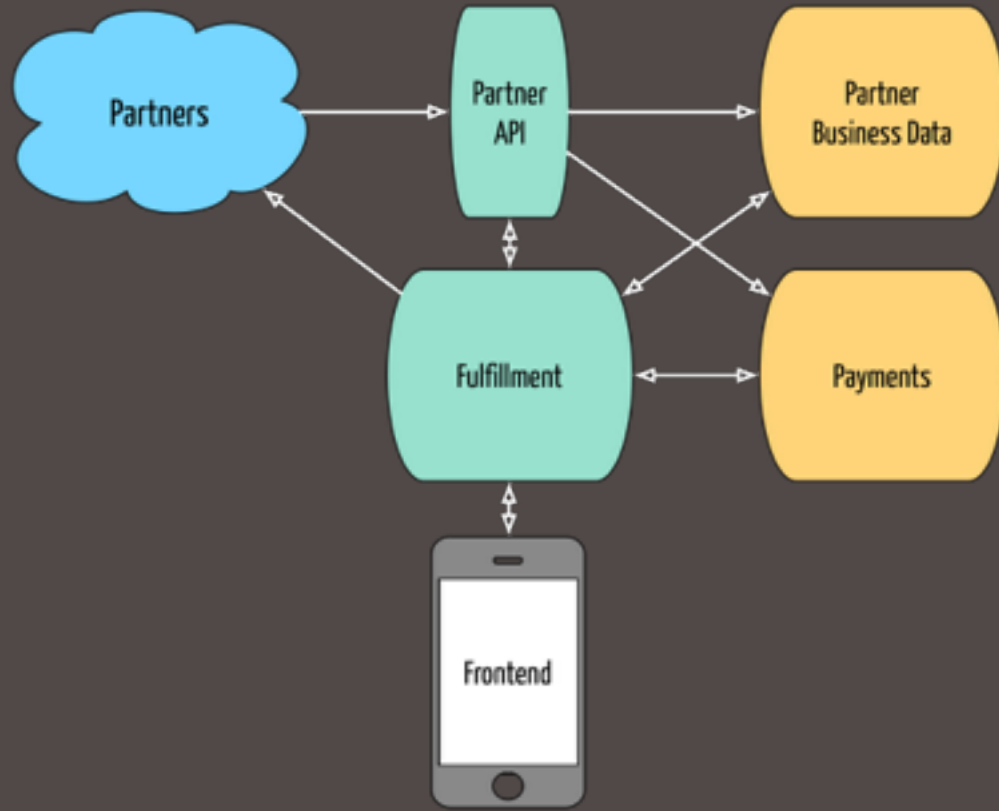
booking\_at\_business

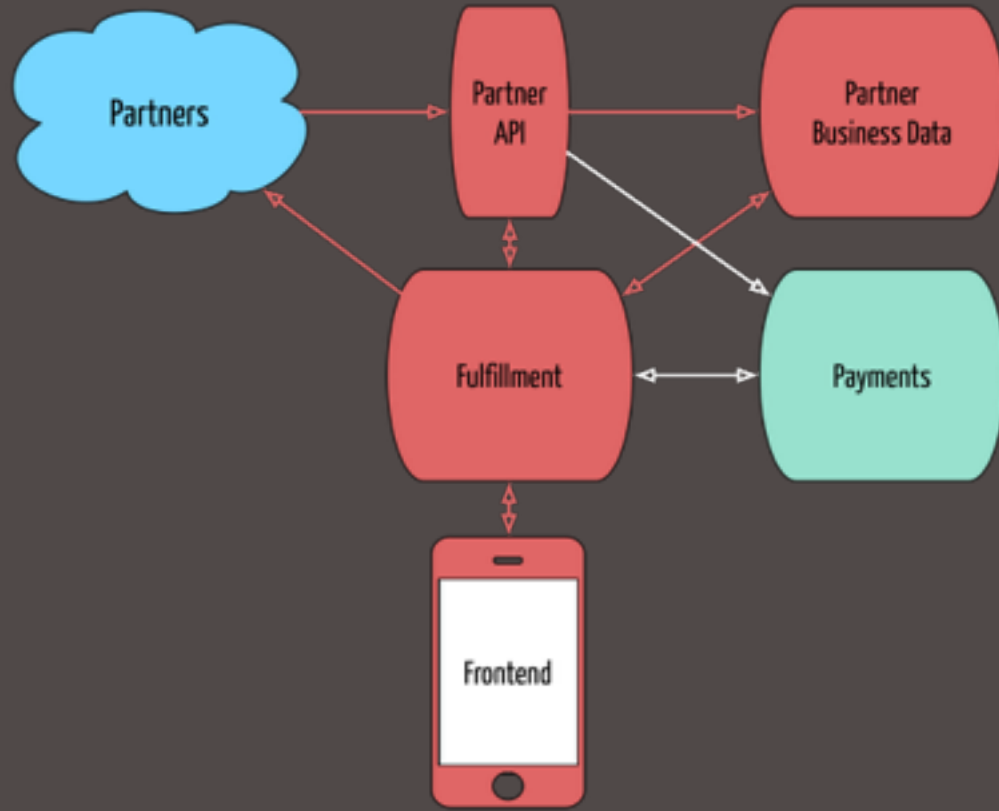
goods\_at\_customer

delivery







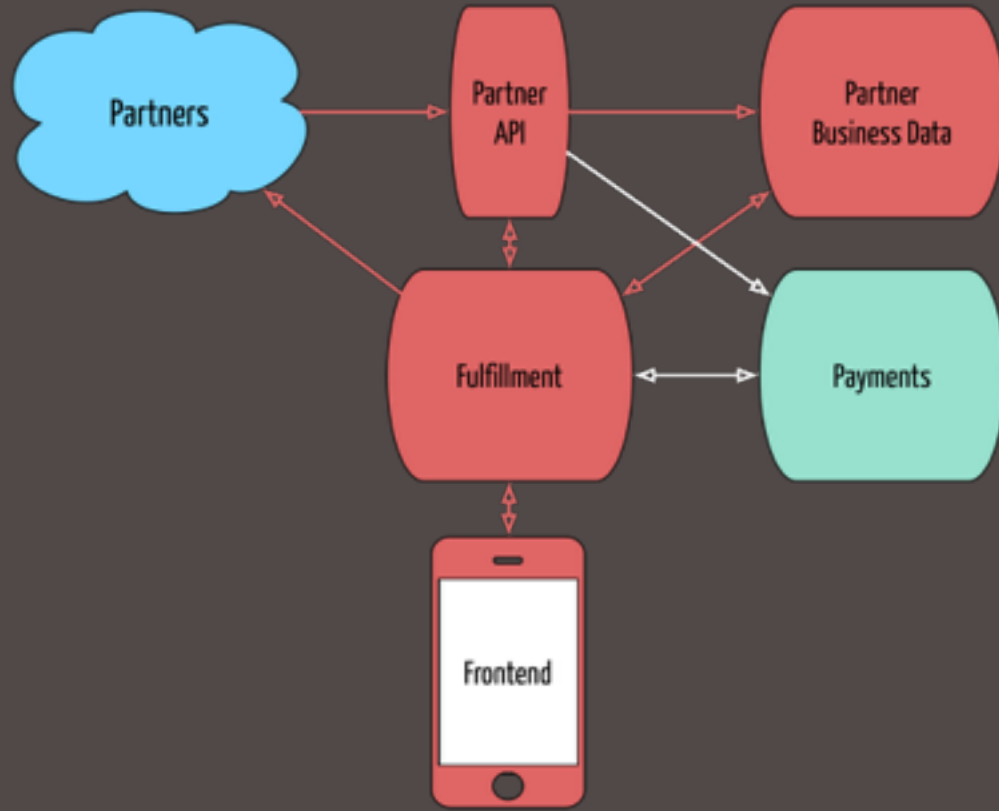


Confusing  
Pervasive

Convenient, but not designed

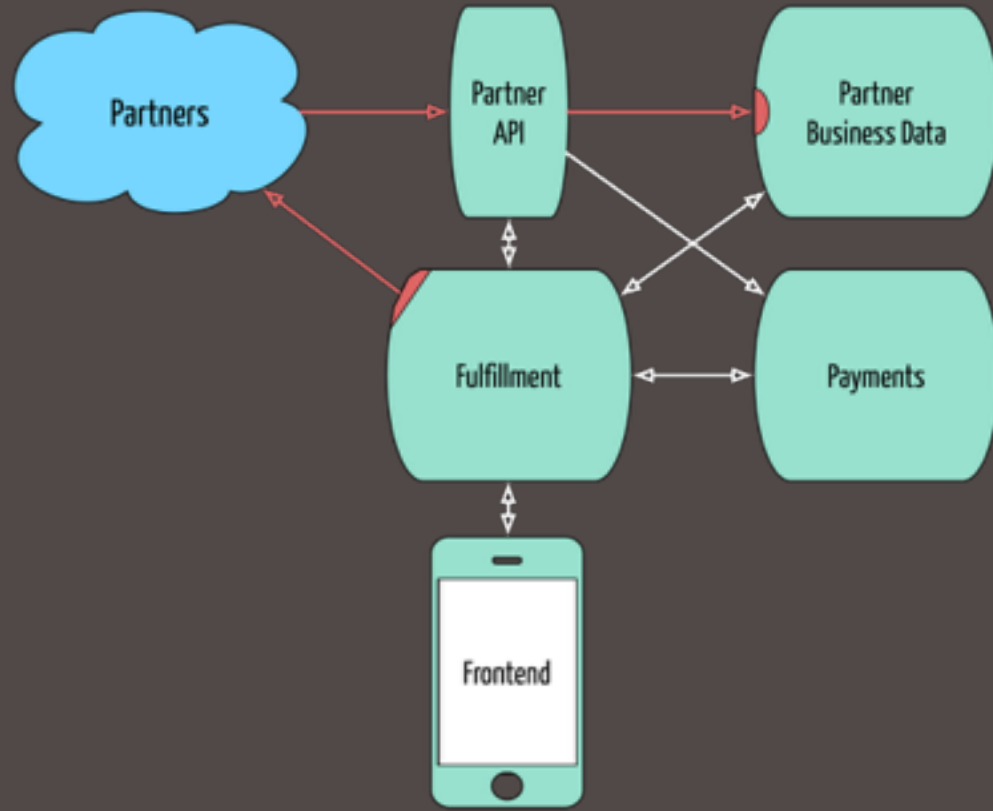






Draw boundaries, introduce  
**domain-specific** concepts  
tied to functionality





# Lessons



Interfaces are the sum of  
APIs, shared libraries, and the  
data that flows through them



Sacrificing DRYness can be the  
best choice for overall design



Service interfaces are a great  
opportunity to intentionally  
decouple systems



- Decoupling
- **Defining**
- Understanding Production
- Staying Agile





Have you ever needed to  
understand a system and  
been told go read the source?



What about a system which  
only validates half its  
interface?

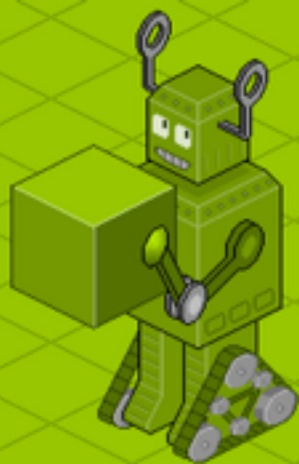


Coming from a python  
monolith, strong interfaces  
were quite rare



```
def checkout(order, price, **kwargs):  
    """Process an order."""  
  
    validate_order(order)  
    charge_credit_card(order.user, price)  
    notify_user(order, **kwargs)
```





# ***SWAGGER***

The World's Most Popular Framework for APIs.

```
13 # will be prefixed to all paths
14 basePath: /v1
15 produces:
16   - application/json
17 paths:
18   /products:
19     get:
20       summary: Product Types
21       description: |
22         The Products endpoint returns information about the *Uber* products
23         offered at a given location. The response includes the display name
24         and other details about each product, and lists the products in the
25         proper display order.
26     parameters:
27       - name: latitude
28         in: query
29         description: Latitude component of location.
30         required: true
31         type: number
32         format: double
33       - name: longitude
34         in: query
35         description: Longitude component of location.
36         required: true
37         type: number
38         format: double
39     tags:
40       - Products
```

## Swagger Petstore

This is a sample server Petstore server. You can find out more about Swagger at <http://swagger.io> or on [#swagger](irc.freenode.net). For this sample, you can use the api key `special-key` to test the authorization filters.

Find out more about Swagger

<http://swagger.io>

[Contact the developer](#)

[Apache 2.0](#)

### pet : Everything about your Pets

[Show/Hide](#)[List Operations](#)[Expand Operations](#)

POST

/pet

Add a new pet to the store

PUT

/pet

Update an existing pet

GET

/pet/findByStatus

Finds Pets by status

GET

/pet/findByTags

Finds Pets by tags

DELETE

/pet/{petId}

Deletes a pet

GET

/pet/{petId}

Find pet by ID

# Client side - Yelp/bravado

```
from bravado.client import SwaggerClient

client = SwaggerClient.from_url(
    "www.myservice.com/swagger.json"
)

pet = client.pet.getPetById(petId=42).result()
```



# Server side - striglia/pyramid\_swagger

```
# In your Pyramid webapp.py  
config.include('pyramid_swagger')
```

# Lessons



Interfaces should be intentional



Interfaces should be **explicit**



Find the **mechanical** things which don't  
scale and **automate** them **mercilessly**



- Decoupling
- Defining
- Understanding Production
- Staying Agile



Real customer bug report:

“We’re seeing 504s talking  
to the /user\_info API”



Ancient times:

Use logic and whatever logs  
happen to exist







(drbethsnow via Flickr; CC BY-NC-ND 2.0)

Better:

Log all incoming API  
requests to any service



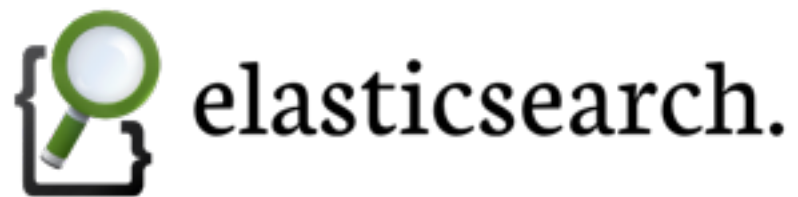


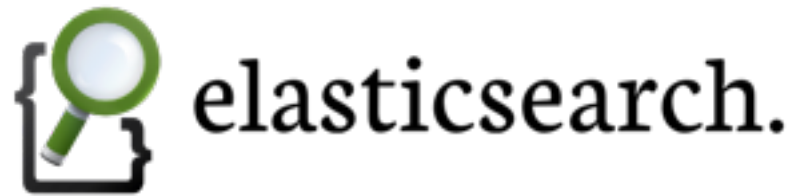
(spam via Flickr; CC by 2.0)

Best:

Every service has a detailed access/  
error log and tooling to examine them



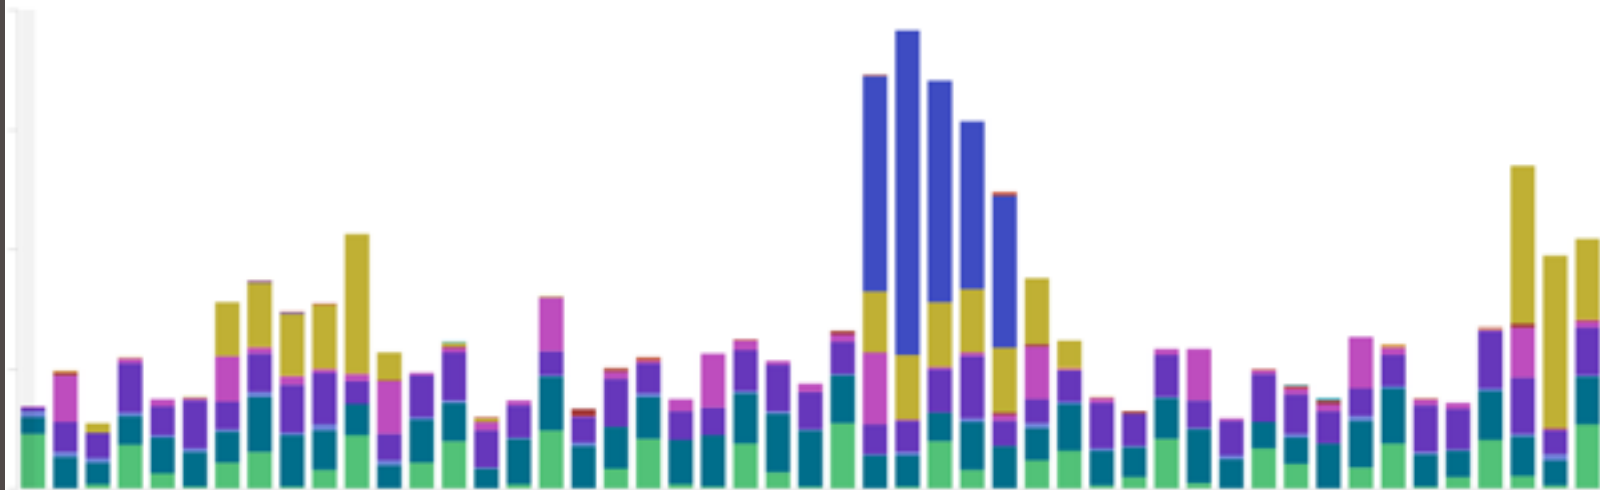




## Checkout Fulfillment Dashboard

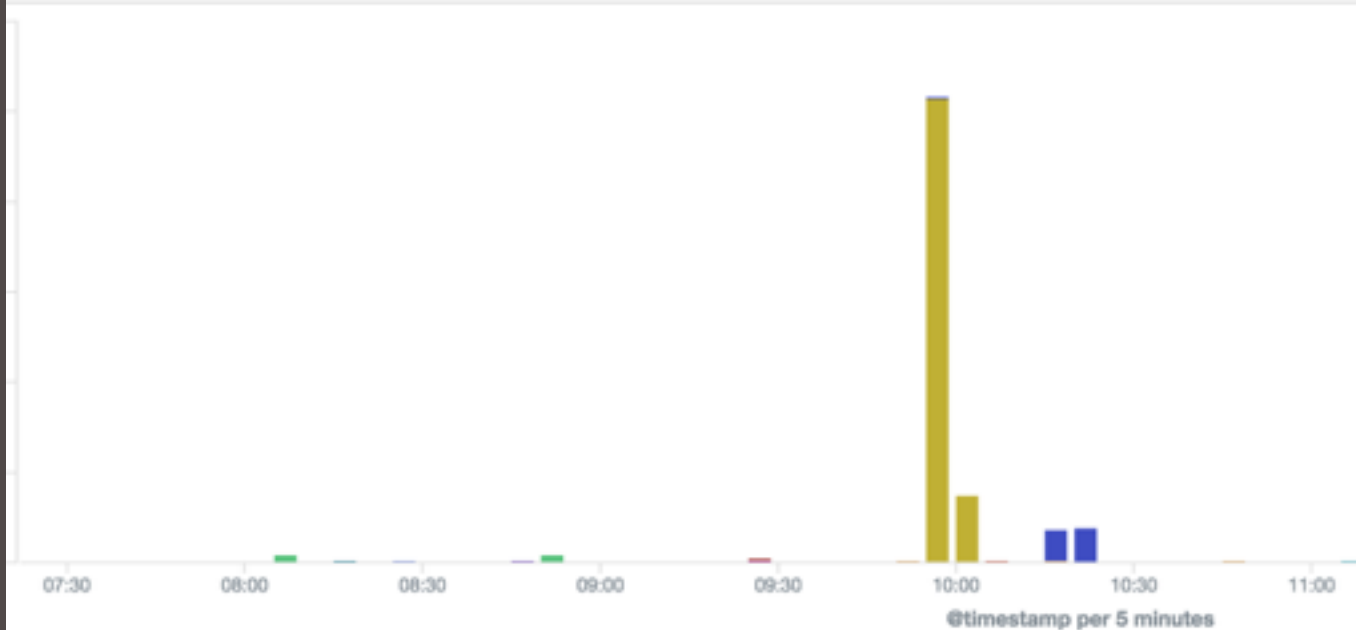
request.method\_name:checkout.v2.\*

API calls by partner





errors over time split by exception class

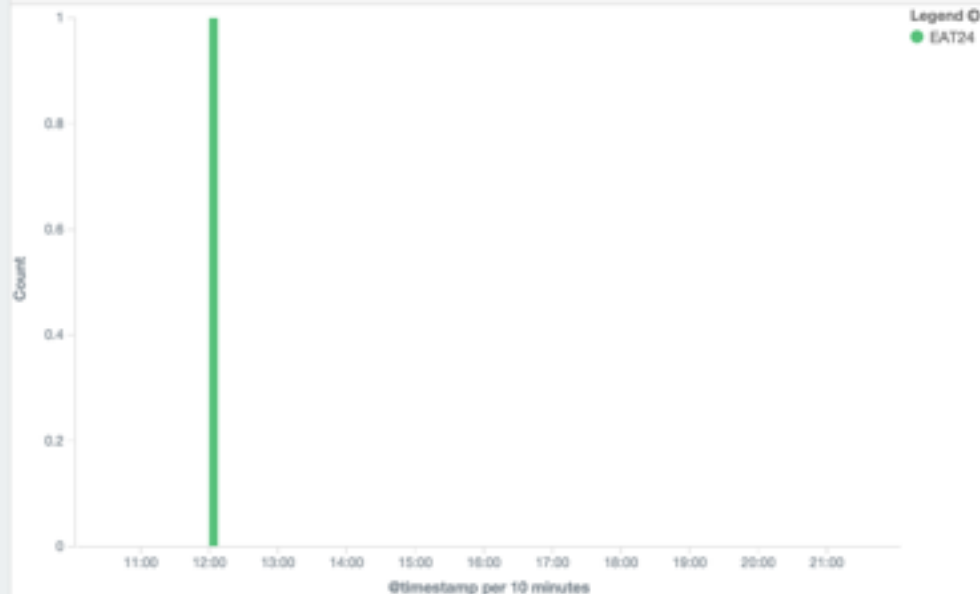




So what about that customer  
with the mystery 504?



YTP: API calls by partner

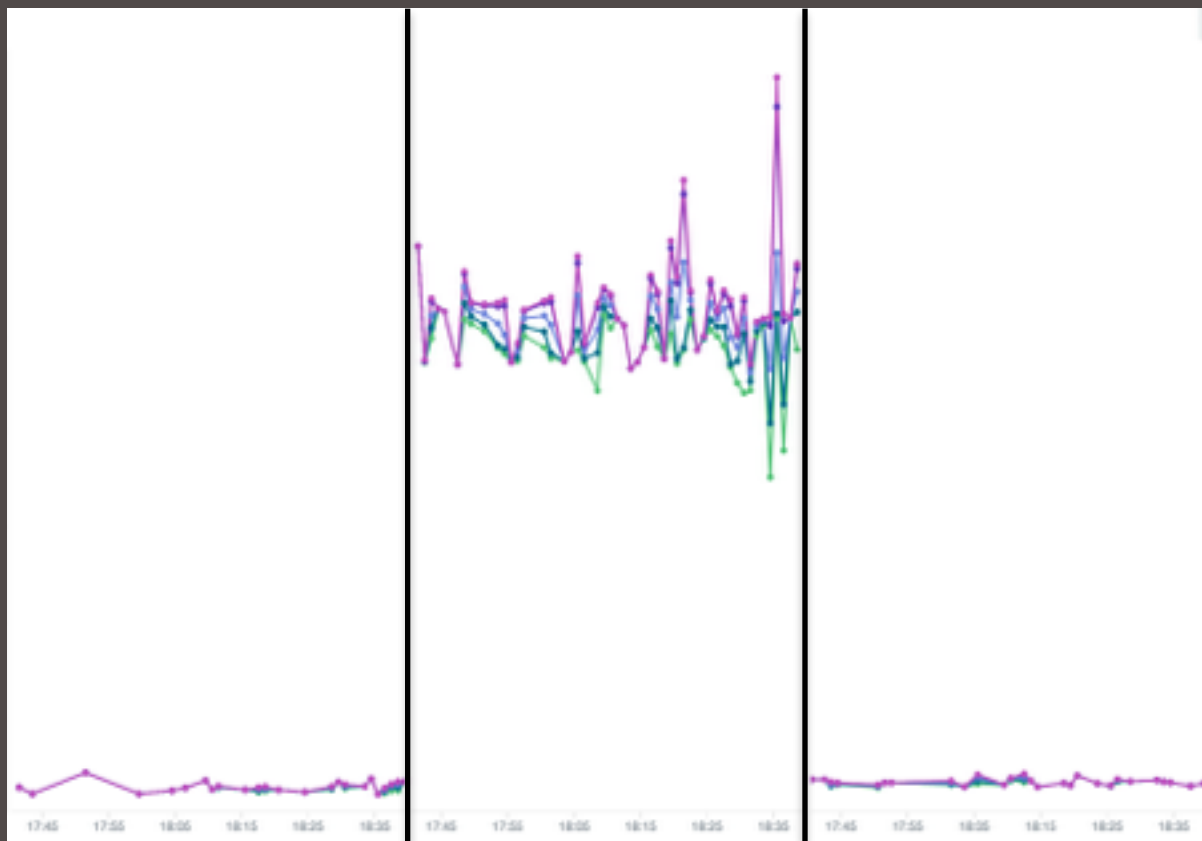


YTP: All CF requests

Time	request.method_name	request.params.yelp_order_id
July 12th 2015, 12:02:23.621	checkout.v2.orders.fulfillment_user_info	eb8b228a026960981bc5fb038760139c

2.5 s

0.15 s



Realistically:

Don't require the customer  
to report issues in  
the first place





Yelp / **elastalert**

Watch ▾

44



Easy & Flexible Alerting With ElasticSearch <https://elastalert.readthedocs.org>

300 commits

17 branches

34 releases

15 contributors



branch: **master** ▾

**elastalert** / +



Merge pull request [#158](#) from sophicware/feature/smtp\_ssl ...



**Qmando** authored 6 days ago


latest commit [e74eea9e6f](#)

```
es_host: elasticsearch-hostname
es_port: 14900
index: logstash-errors-%G.%V
```

```
type: frequency
num_events: 20
timeframe:
  minutes: 2
```

```
alert:
  - "modules.sensu_alert.SensuAlerter"
```

```
sensu:
  team: platform
  tip: "This alert indicates a large number of errors across the Platform
product. See <link to Kibana> for details."
  page: true
  status: 2 # CRITICAL
```



```
es_host: elasticsearch-hostname
es_port: 14900
index: logstash-errors-%G.%V
```

```
type: frequency
num_events: 20
timeframe:
  minutes: 2
```

```
alert:
  - "modules.sensu_alert.SensuAlerter"
```

```
sensu:
  team: platform
  tip: "This alert indicates a large number of errors across the Platform
product. See <link to Kibana> for details."
  page: true
  status: 2 # CRITICAL
```



```
es_host: elasticsearch-hostname
es_port: 14900
index: logstash-errors-%G.%V
```


```
type: frequency
num_events: 20
timeframe:
  minutes: 2
```

alert:

- "modules.sensu\_alert.SensuAlerter"

sensu:

```
team: platform
tip: "This alert indicates a large number of errors across
the Platform product. See <link to Kibana> for details."
page: true
status: 2 # CRITICAL
```





# Lessons



Logging is a superpower. Use it ~~wisely~~  
constantly.



But raw data is not enough! **Visualize  
and monitor actively.**



These approaches make a world of difference:

- Incident response from days to minutes
- Investigations from  $\infty$  to minutes



- Decoupling
- Defining
- Understanding Production
- Staying Agile



# Uncomfortable conversation:

“Customers had their orders interrupted. How are you preventing it going forward?”



Understandable response:

*“Deploy more carefully”*



Understandable response:  
“Expand oncall”





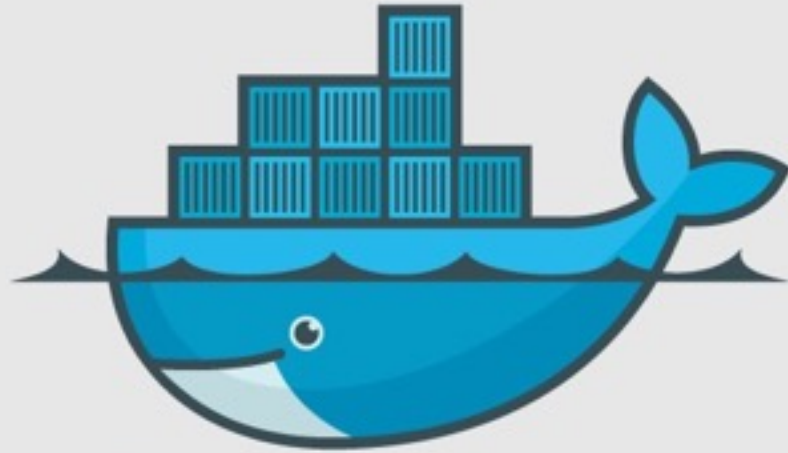
How do we ensure the team  
stays agile as our services  
grow in complexity?



# Pain point:

The testing environment is  
{broken, flaky, not like prod}





docker

Pain point:

Tests passed but  
production broke



Production monitoring is the  
**natural extension** of  
excellent pre-deploy testing.





Yelp / **elastalert**

Watch ▾

44



Easy & Flexible Alerting With ElasticSearch <https://elastalert.readthedocs.org>

300 commits

17 branches

34 releases

15 contributors



branch: **master** ▾

**elastalert** / +



Merge pull request [#158](#) from sophicware/feature/smtp\_ssl ...



**Qmando** authored 6 days ago

latest commit [e74eea9e6f](#)

Pain point:

No clue how much time we  
spend fixing production issues



Pain point:

Tough to argue what changes  
will make things more robust





**Outage Started On \***

mm/dd/yyyy, --:-- --


Example: 03/05/2013 11:30 AM

**Outage Finished On \***

mm/dd/yyyy, --:-- --

Example: 03/05/2013 11:30 AM

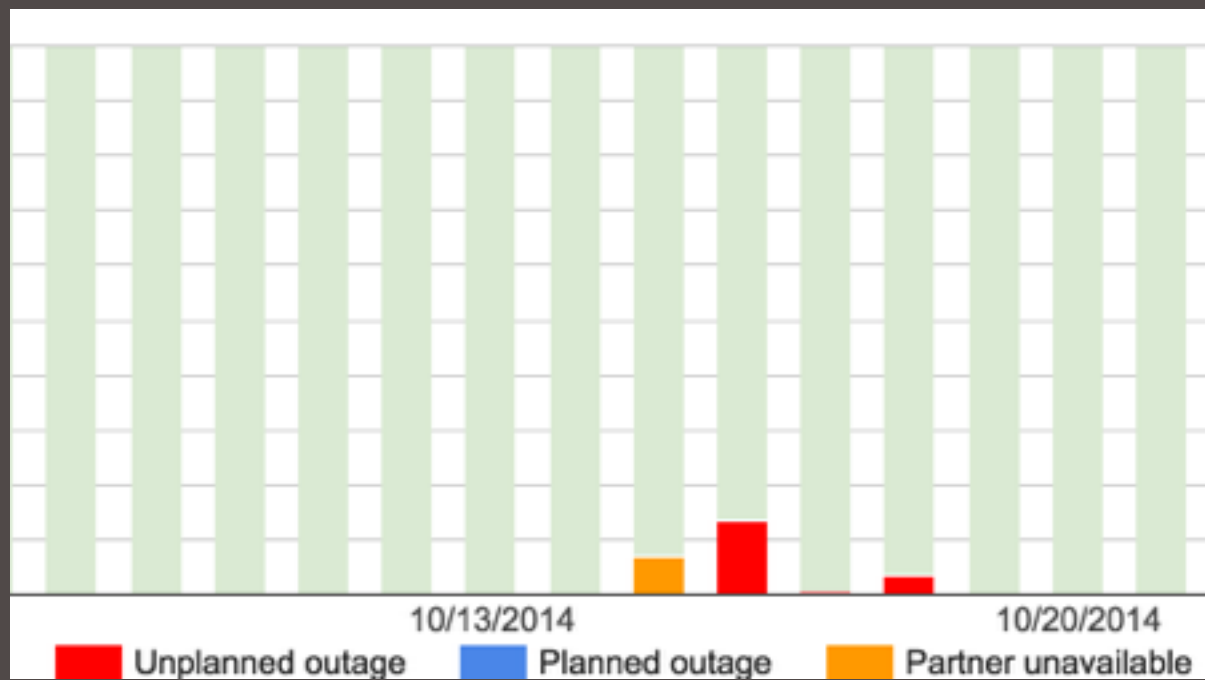
**Did this impact a specific partner, or all of them? \***

**Was this a planned outage? \***

- ☐ Unplanned
- ☐ Planned

**What is the nature of the outage? \***



And as with everything else, this **must**  
**eventually be automated**



# Lessons



Networks of services are **fundamentally harder to test**. Prepare accordingly.



Failure will happen. Focus on both  
identifying and recovering quickly.



Staying agile is easy if your application  
rarely fails and recovers automatically



# Wrap Up





Know your roots



# Be explicit



# Measure everything



# Scale via automation





Yelp/bravado

striglia/pyramid\_swagger

Yelp/elastalert



[http://engineeringblog.yelp.com/2015/03/  
using-services-to-break-down-  
monoliths.html](http://engineeringblog.yelp.com/2015/03/using-services-to-break-down-monoliths.html)



Our accumulated wisdom  
Yelp/service-principles



# Questions?

@scott\_triglia

scott.triglia@gmail.com

