

# High Availability at Braintree

Paul Gross

[paul.gross@braintreepayments.com](mailto:paul.gross@braintreepayments.com)

[twitter.com/pgr0ss](https://twitter.com/pgr0ss)

[github.com/pgr0ss](https://github.com/pgr0ss)


[pgrs.net](http://pgrs.net)

# Braintree

Braintree is a payment gateway

A payment gateway is software that allows merchants to process credit card payments from your website and/or application

# Our Merchants



# Why is uptime important?

5 billion dollars in annual processing

\$9,500 per minute for our merchants

# Uptime Percentages

<b>Availability %</b>	<b>Downtime per year</b>	<b>Downtime per month*</b>	<b>Downtime per week</b>
90% ("one nine")	36.5 days	72 hours	16.8 hours
99% ("two nines")	3.65 days	7.20 hours	1.68 hours
99.9% ("three nines")	8.76 hours	43.8 minutes	10.1 minutes
99.99% ("four nines")	52.56 minutes	4.32 minutes	1.01 minutes
99.999% ("five nines")	5.26 minutes	25.9 seconds	6.05 seconds
99.9999% ("six nines")	31.5 seconds	2.59 seconds	0.605 seconds
99.99999% ("seven nines")	3.15 seconds	0.259 seconds	0.0605 seconds

([https://en.wikipedia.org/wiki/High\\_availability](https://en.wikipedia.org/wiki/High_availability))

# 2 kinds of downtime

Planned

Unplanned

**Reduce our maintenance windows**

# Switched from MySQL to PostgreSQL

DDL migrations are extremely fast

Add indexes without locking tables

Transactional DDL

<http://www.pgrs.net/2011/03/25/migrating-from-mysql-to-postgresql-slides/>



# Deploy process

Add new tables and columns

Roll out new code (server by server)

Add indexes

```
namespace :db do
  task :migrate_pre => :environment do
    ActiveRecord::Migrator.migrate "db/migrate_pre"
  end

  task :migrate_post => :environment do
    ActiveRecord::Migrator.migrate "db/migrate_post"
  end
end
```

# Rails caches columns

Can't drop columns in a post migration

Need to tell Rails to forget the column

```
class User < ActiveRecord::Base
  deleted_columns :old_column
end
```

```
ActiveRecord::Base.class_eval do
  def self.deleted_columns(*column_names)
    @deleted_columns = column_names.map(&:to_s)
  end
end
```

```
ActiveRecord::Base.class_eval do
  def self.deleted_columns(*column_names)
    @deleted_columns = column_names.map(&:to_s)
  end

  def self.columns_with_removing_deleted
    columns_without_removing_deleted.reject do |c|
      @deleted_columns.include?(c.name)
    end
  end
  alias_method_chain :columns, :removing_deleted
end
```

We run multiple versions of the code at  
once

Fine for most features

Feature switches to turn on new features at once

# Limitations


Column renames

Database failover

Infrastructure changes



**Want a way to pause traffic**



# Broxy = Braintree Proxy

Python/Tornado (evented)

Accepts web requests

Feeds redis queue

Reads responses from redis

# Dispatchers

Lightweight rack adapter

Takes requests from redis

Processes through rails

Puts response back in redis

```
require "#{ENV['RAILS_ROOT']}/config/environment"  
app = Rack::Builder.new do  
  run ActionController::Dispatcher.new  
end
```

```
require "#{ENV['RAILS_ROOT']}/config/environment"
app = Rack::Builder.new do
  run ActionController::Dispatcher.new
end

loop do
  if request_data = redis.pull_request
    rack_request = {
      "PATH_INFO" => request_data["request"]["uri"]
      "rack.input" => StringIO.new(request_data["body"])
    }

    rack_response = app.call(rack_request)
  end
end
```

```

require "#{ENV['RAILS_ROOT']}/config/environment"
app = Rack::Builder.new do
  run ActionController::Dispatcher.new
end


loop do
  if request_data = redis.pull_request
    rack_request = {
      "PATH_INFO" => request_data["request"]["uri"]
      "rack.input" => StringIO.new(request_data["body"])
    }

    rack_response = app.call(rack_request)

    body = ""; rack_response[2].each { |part| body << part }

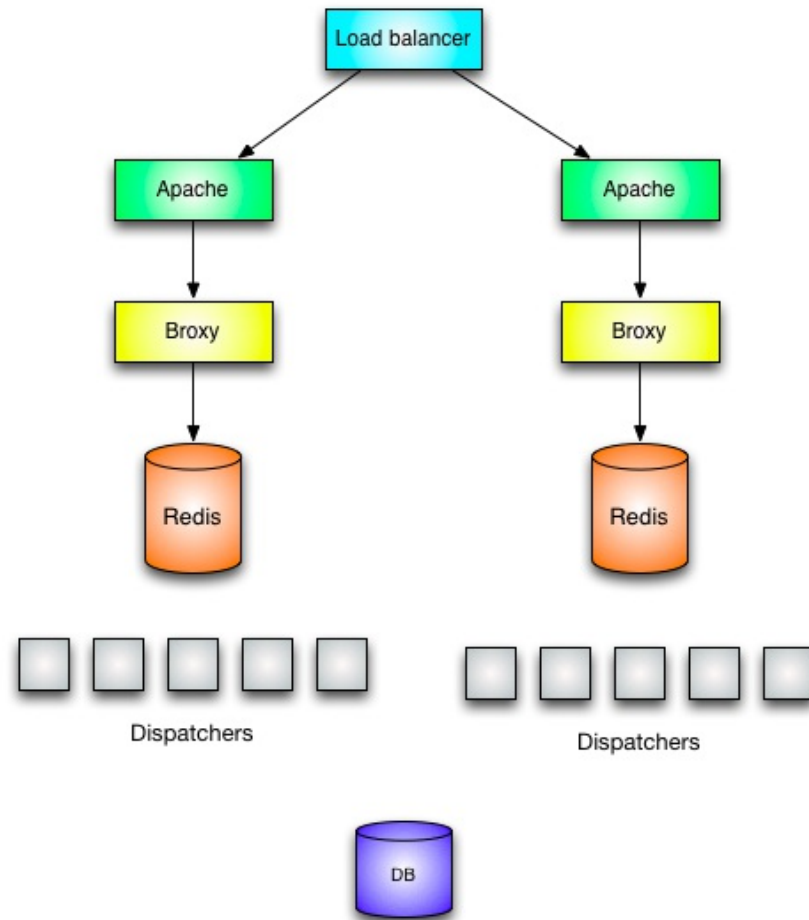
    @redis.push_response(
      "status" => rack_response[0].to_i,
      "body" => body
    )
  end
end

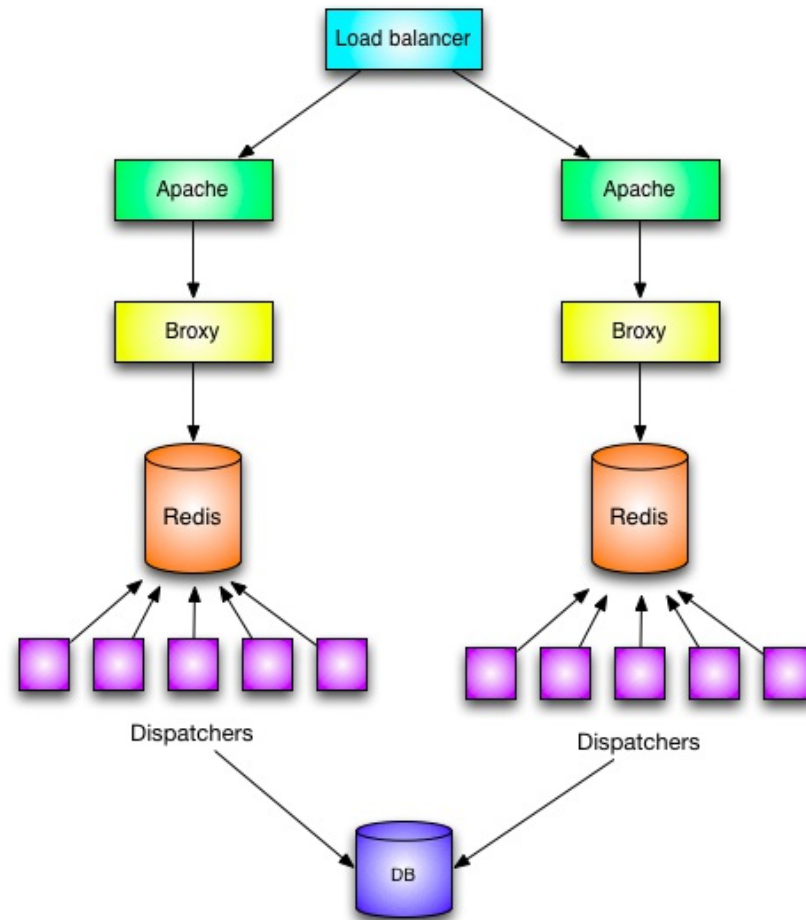
```





**Stop dispatchers to suspend traffic**





# Summary - reducing maintenance windows

Pre and post migrations

Rolling deploys

PostgreSQL for fast DDL

Broxy to pause traffic

# Unplanned failures

Servers will fail

Networks will go down

The unexpected will happen

We do our best to be resilient

**Server failure**

# Load balancers

Build our own


LVS/IPVS

Pacemaker

BigBrother

LitmusPaper

# Load balancing





# BigBrother

Ruby app

Runs on load balancers

Checks status of servers

Update IPVS rules

[https://github.com/braintree/big\\_brother](https://github.com/braintree/big_brother)

# LitmusPaper

Ruby app

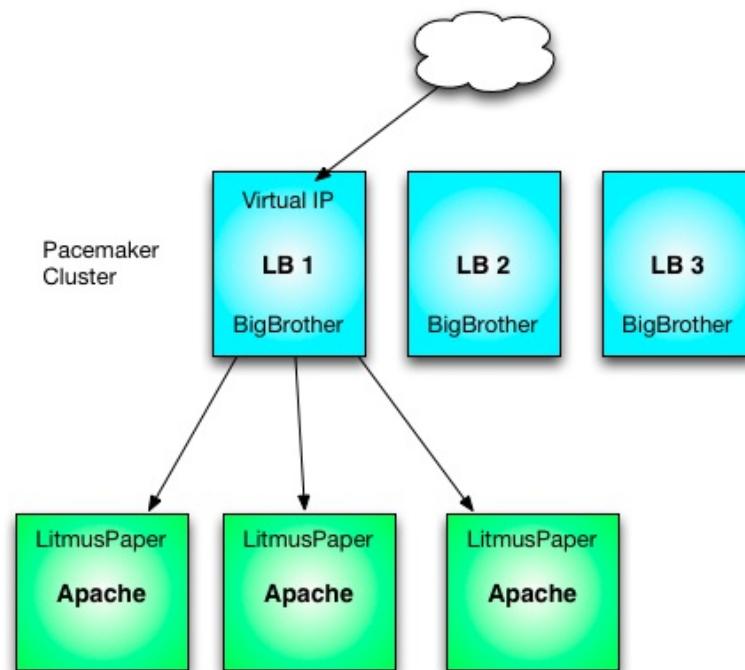
Runs on backend servers

Queried by BigBrother via HTTP

Returns a health level

[https://github.com/braintree/litmus\\_paper](https://github.com/braintree/litmus_paper)

# Load balancing



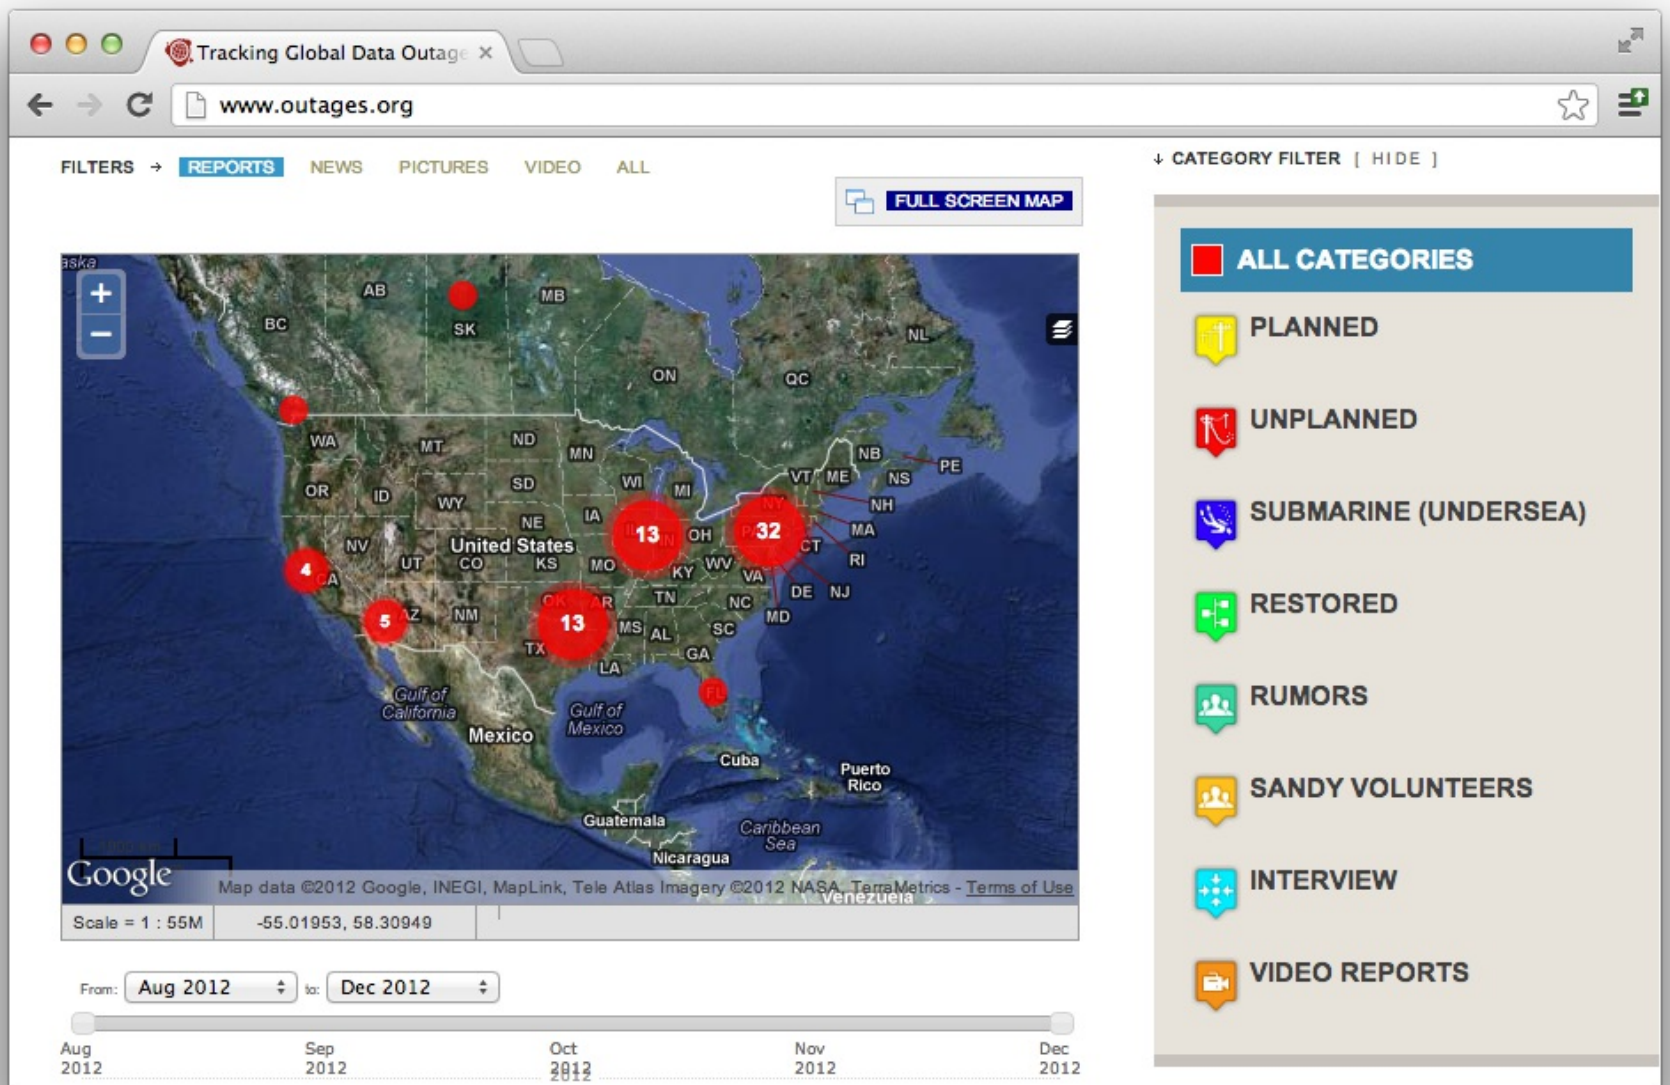
# Stateful services

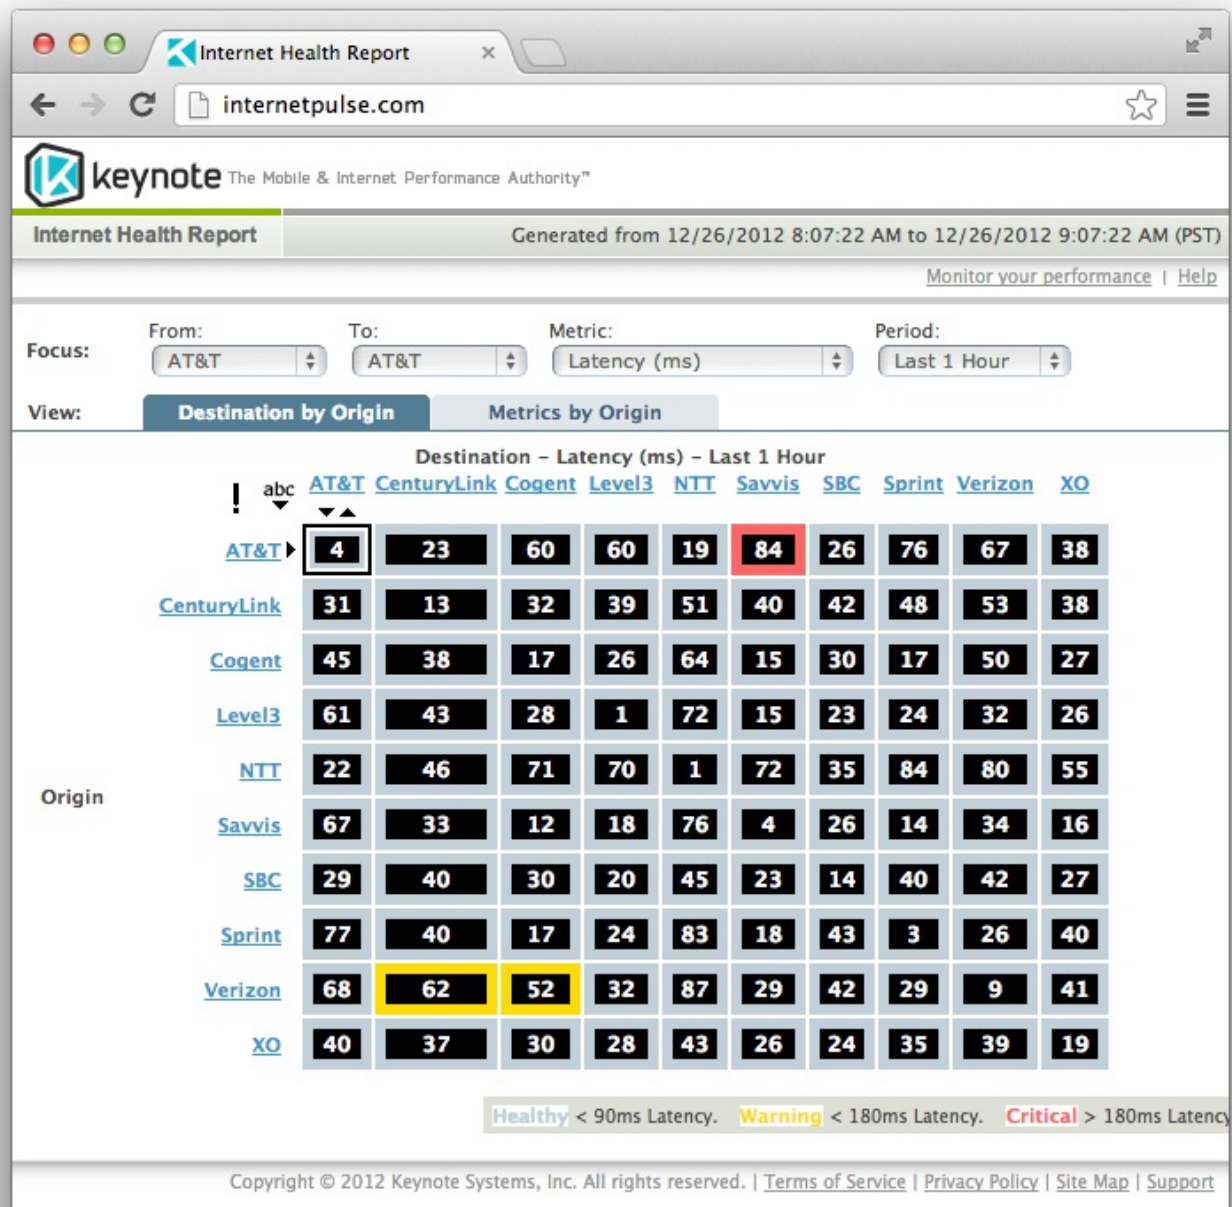
Load balancers and PostgreSQL clusters

Pacemaker manages failover

Virtual IP follows the new primary

# Network failures





# Networking - inbound

BGP routes traffic through multiple ISPs and data centers

We use Pingdom and a handful of globally distributed servers to test connectivity



# Networking - outbound

We connect to many processing networks

ISP outages are usually partial

Sometimes, we can't reach every endpoint on all of our ISPs

Needed a way to choose an ISP per processing network

# Processor proxies

Instead of connecting directly, connect through proxies

One proxy per TCP endpoint and uplink ISP

Load balance over these proxies

Allows us to route around ISP connection issues

# Mallory

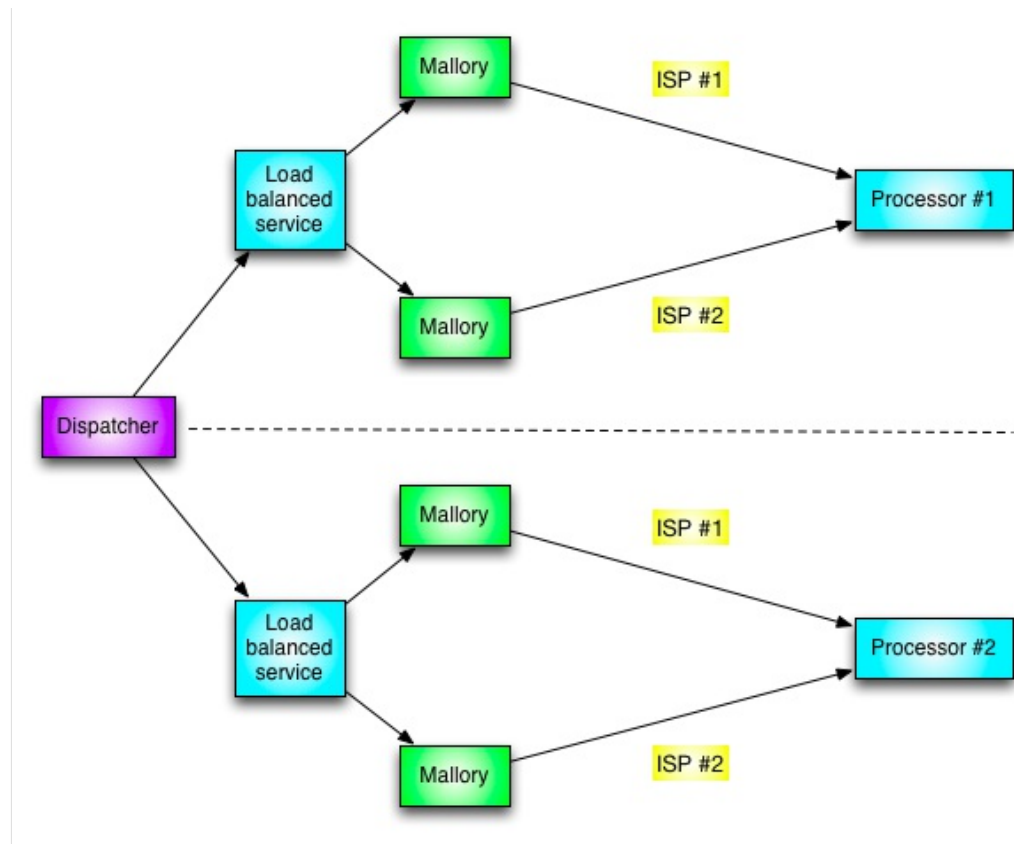
Python/Tornado (evented)

Proxies requests

SSL verification

Acts like LitmusPaper

<https://github.com/braintree/mallory>



# Connection failures

Let the service heal (unbalance or pacemaker)

Retry request

# Automate everything

Reduces human errors

Gives confidence that task will work

Speeds up processes

Less fiddling around in production

Capistrano

# Summary - unplanned failures

Load balancing

Redundancy across ISPs

Let the system heal and retry

Automation

# Questions?

Paul Gross

[paul.gross@braintreepayments.com](mailto:paul.gross@braintreepayments.com)

[twitter.com/pgr0ss](https://twitter.com/pgr0ss)

[github.com/pgr0ss](https://github.com/pgr0ss)

[pgrs.net](http://pgrs.net)