Architecting for Failure

Michael Brunton-Spall
@bruntonspall

QCon London 2012

The inevitability of failure

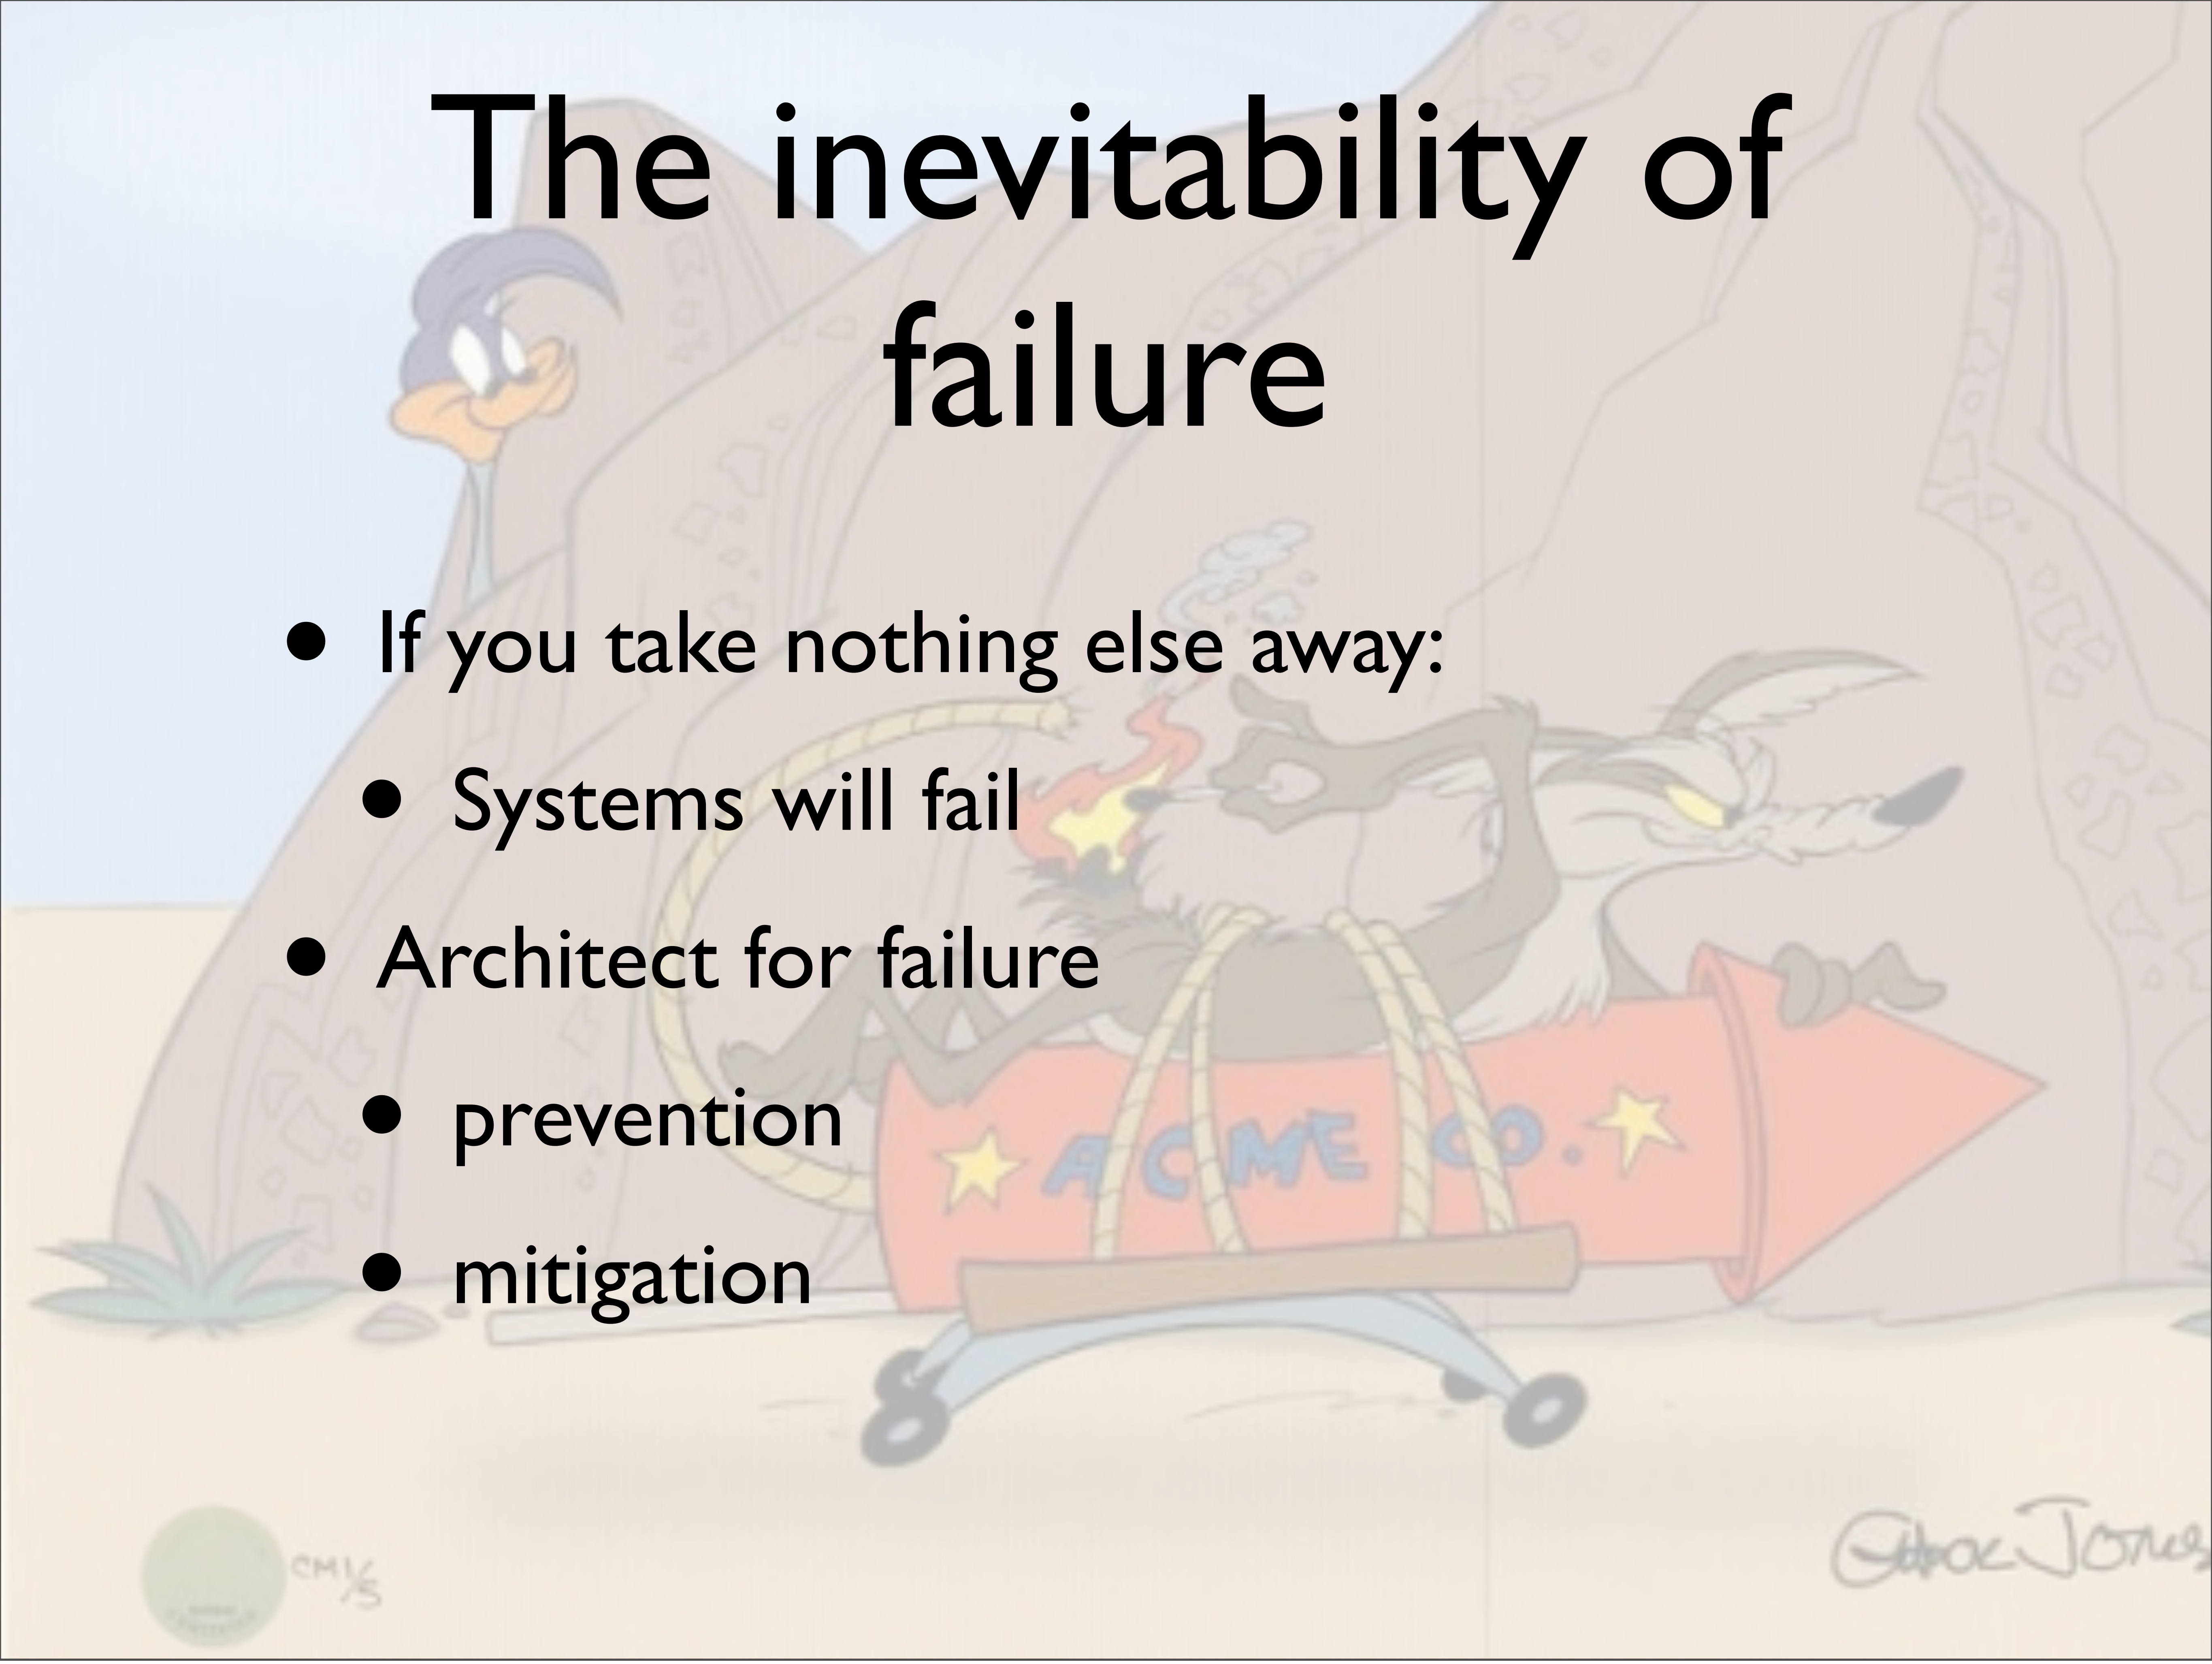

- If you take nothing else away:
  - Systems will fail
- Architect for failure
  - prevention
  - mitigation

guardian.co.uk
circa 2008

Basic Architecture

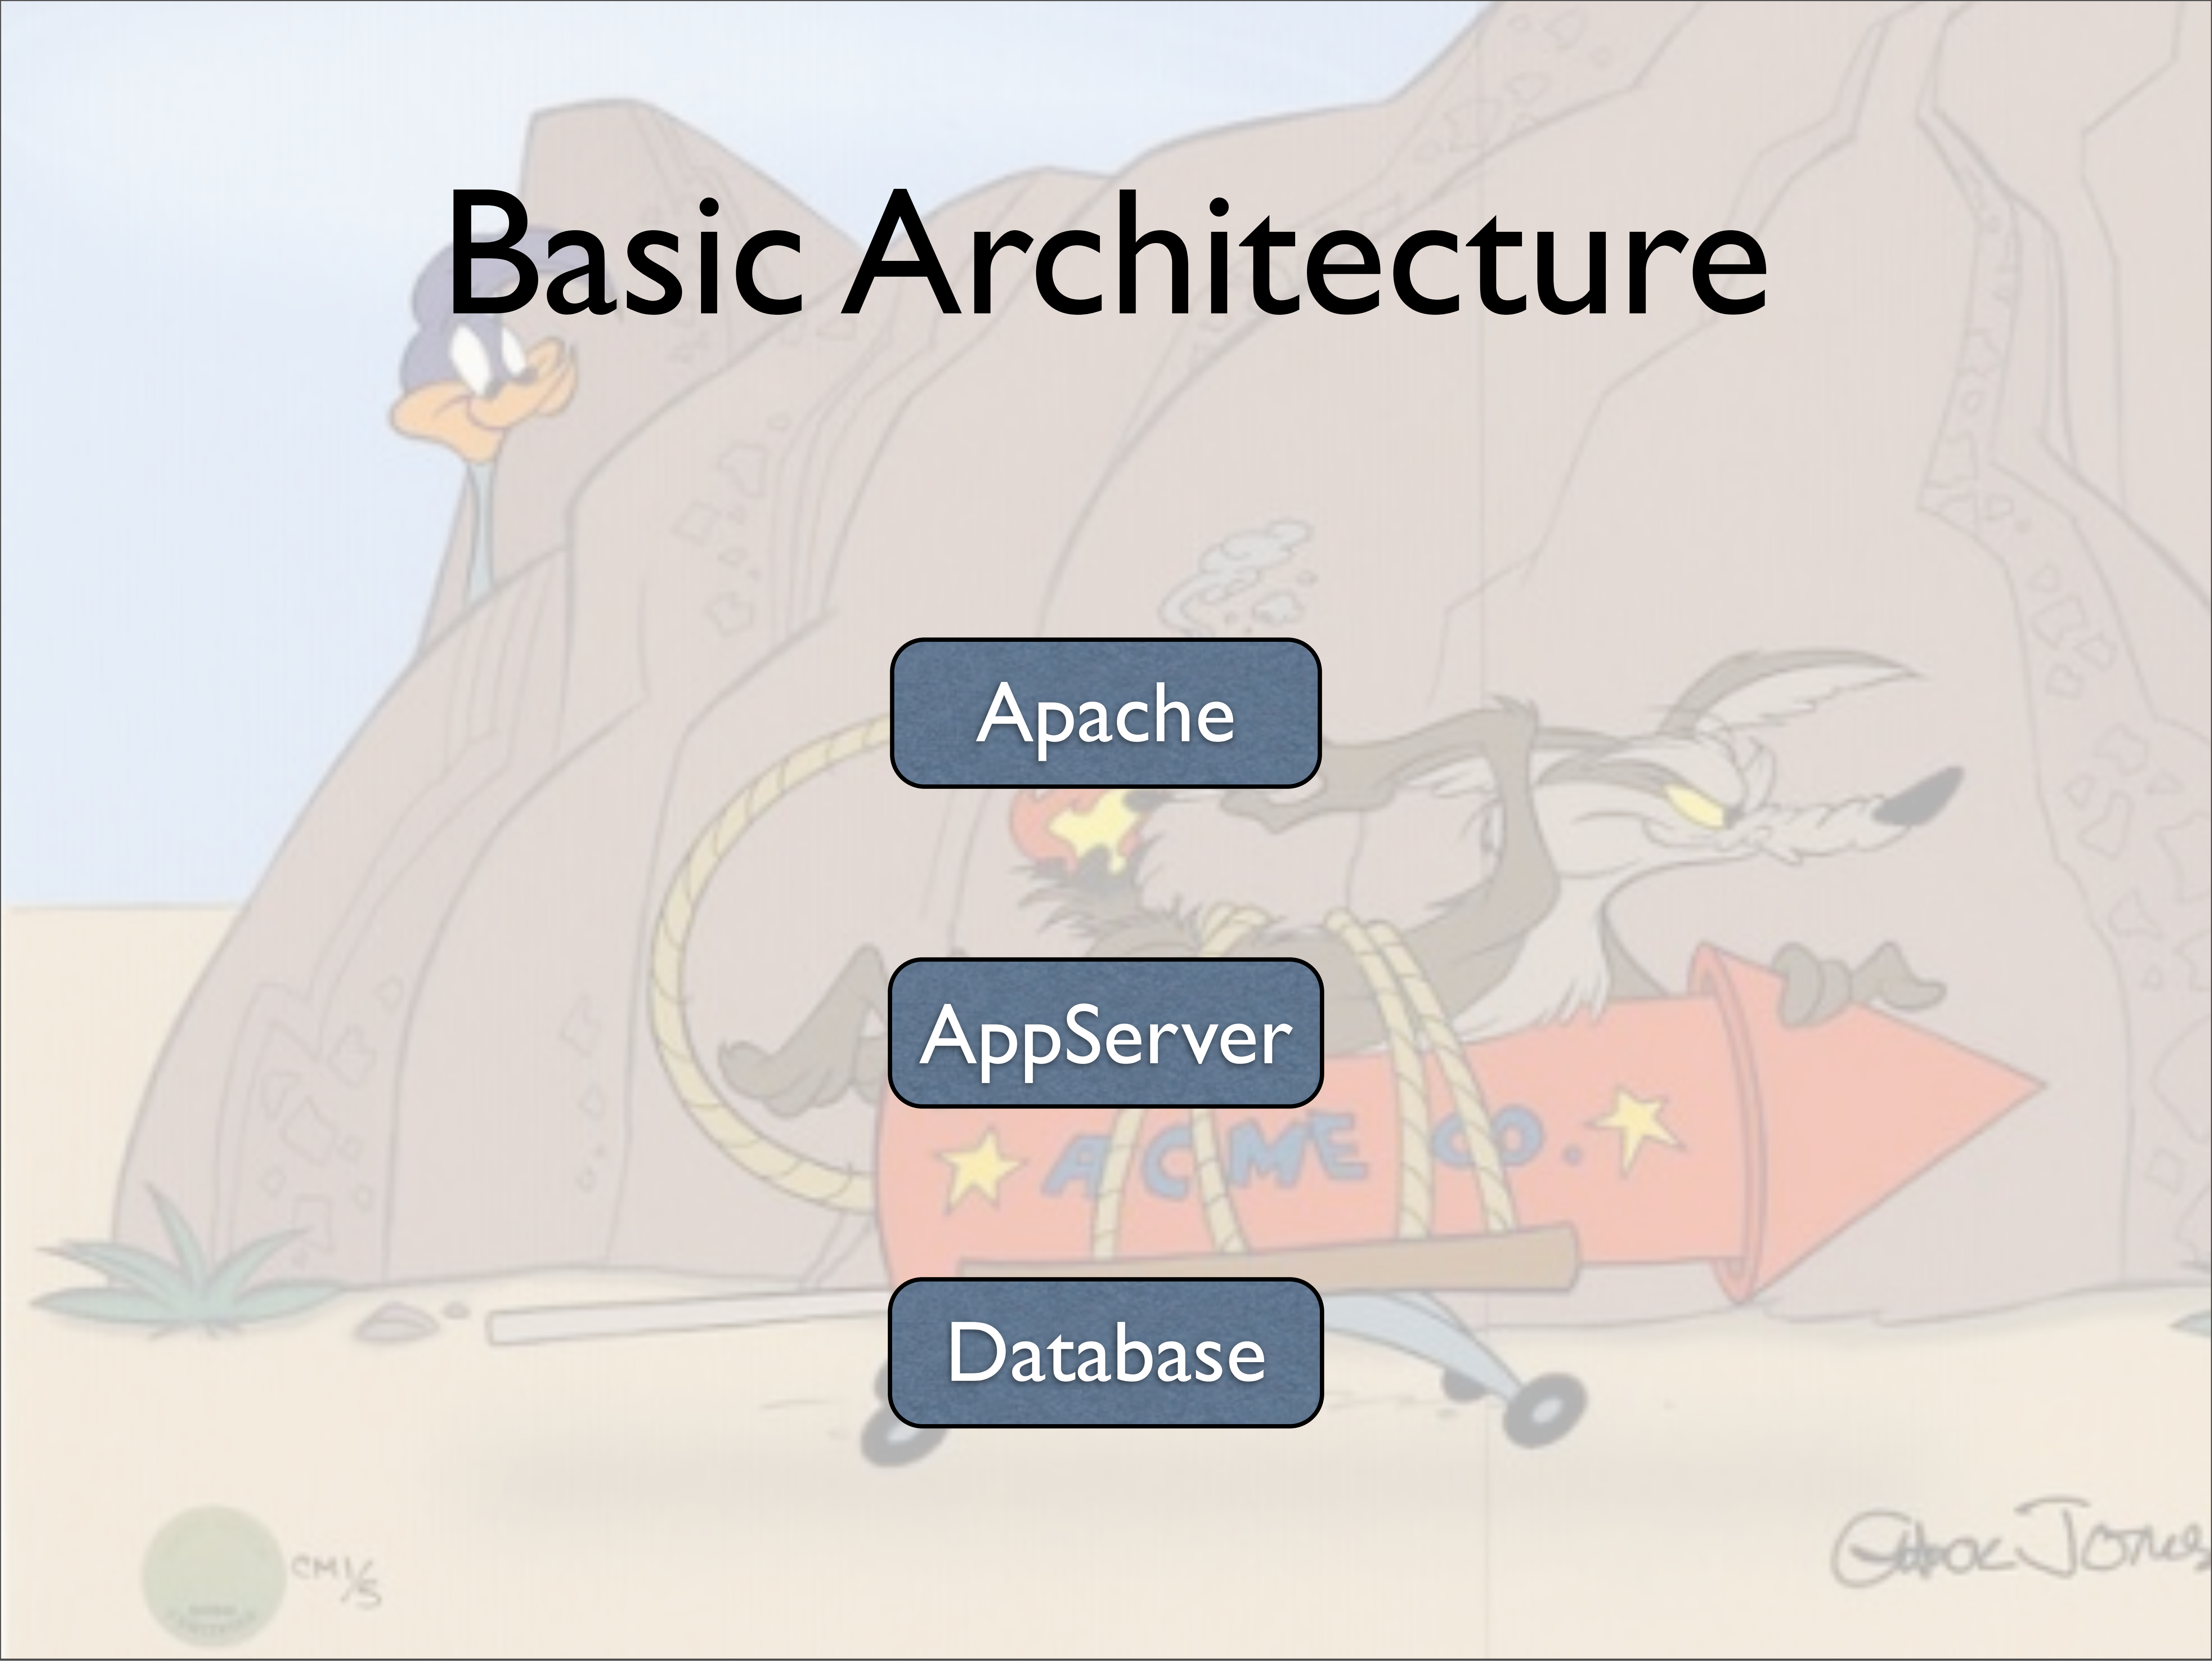

Apache

AppServer

Database

Basic Architecture

- This is your basic J2EE stack
- Lets apply scaling basics

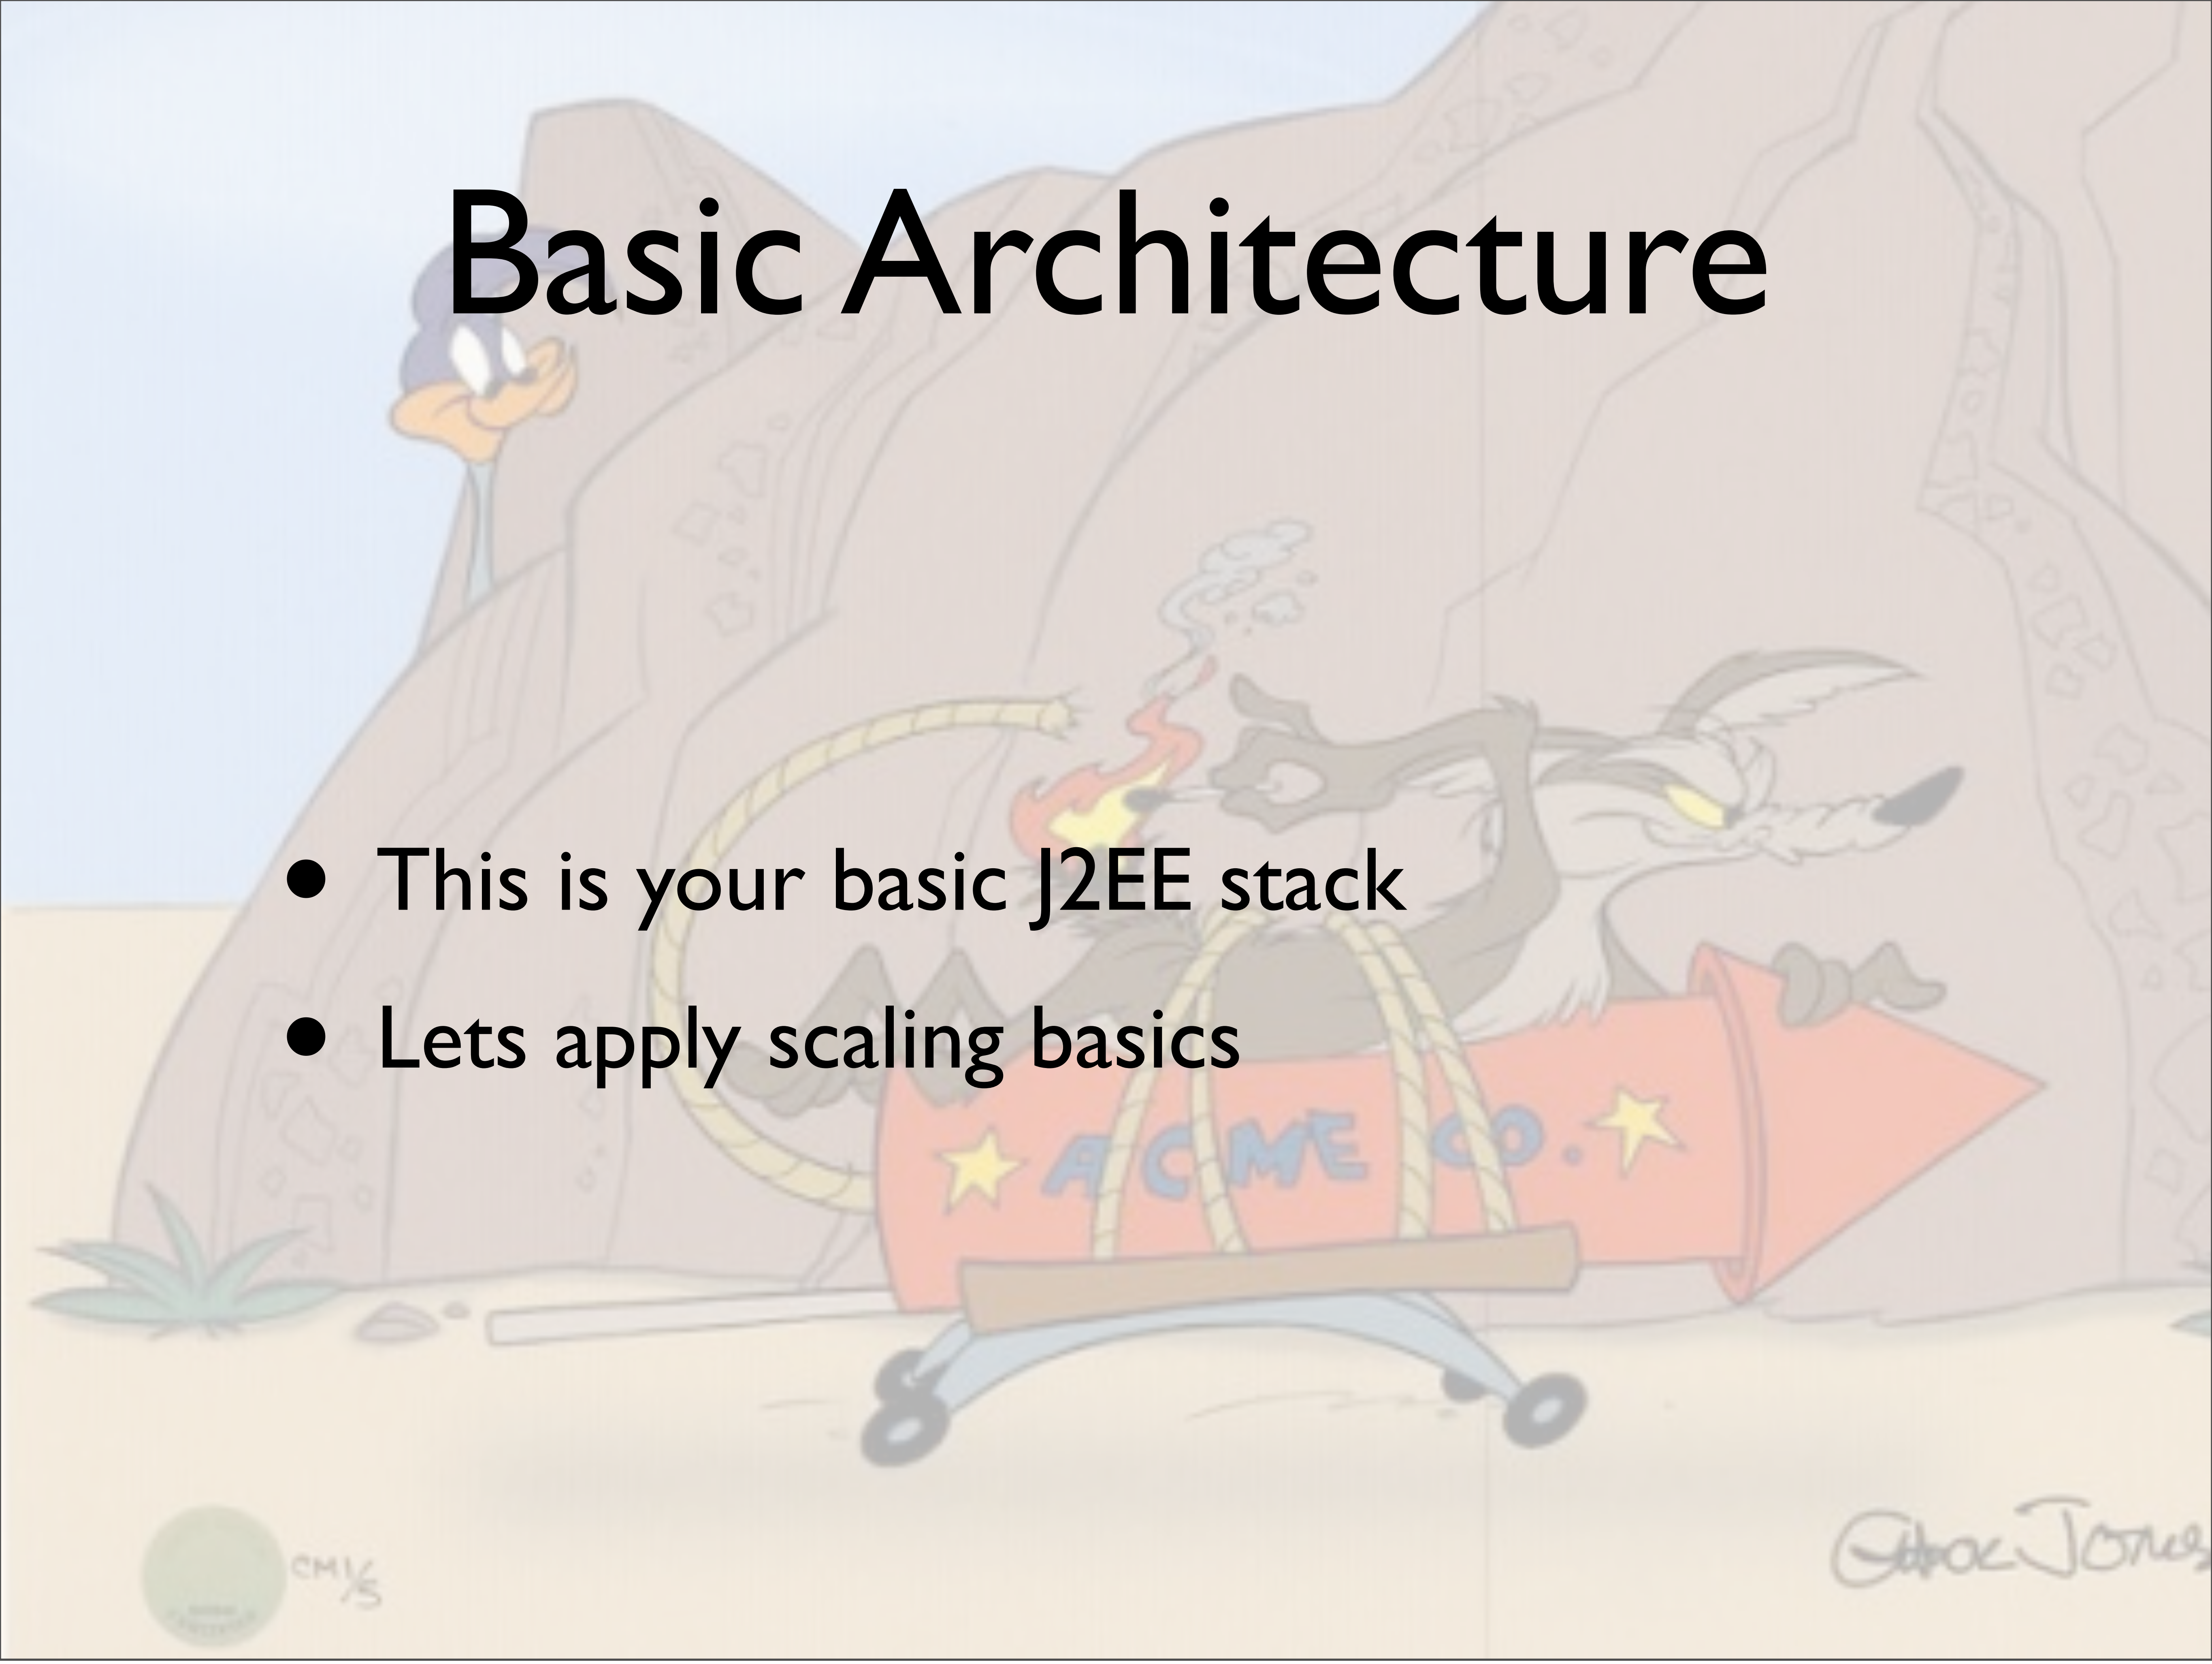

Scaling

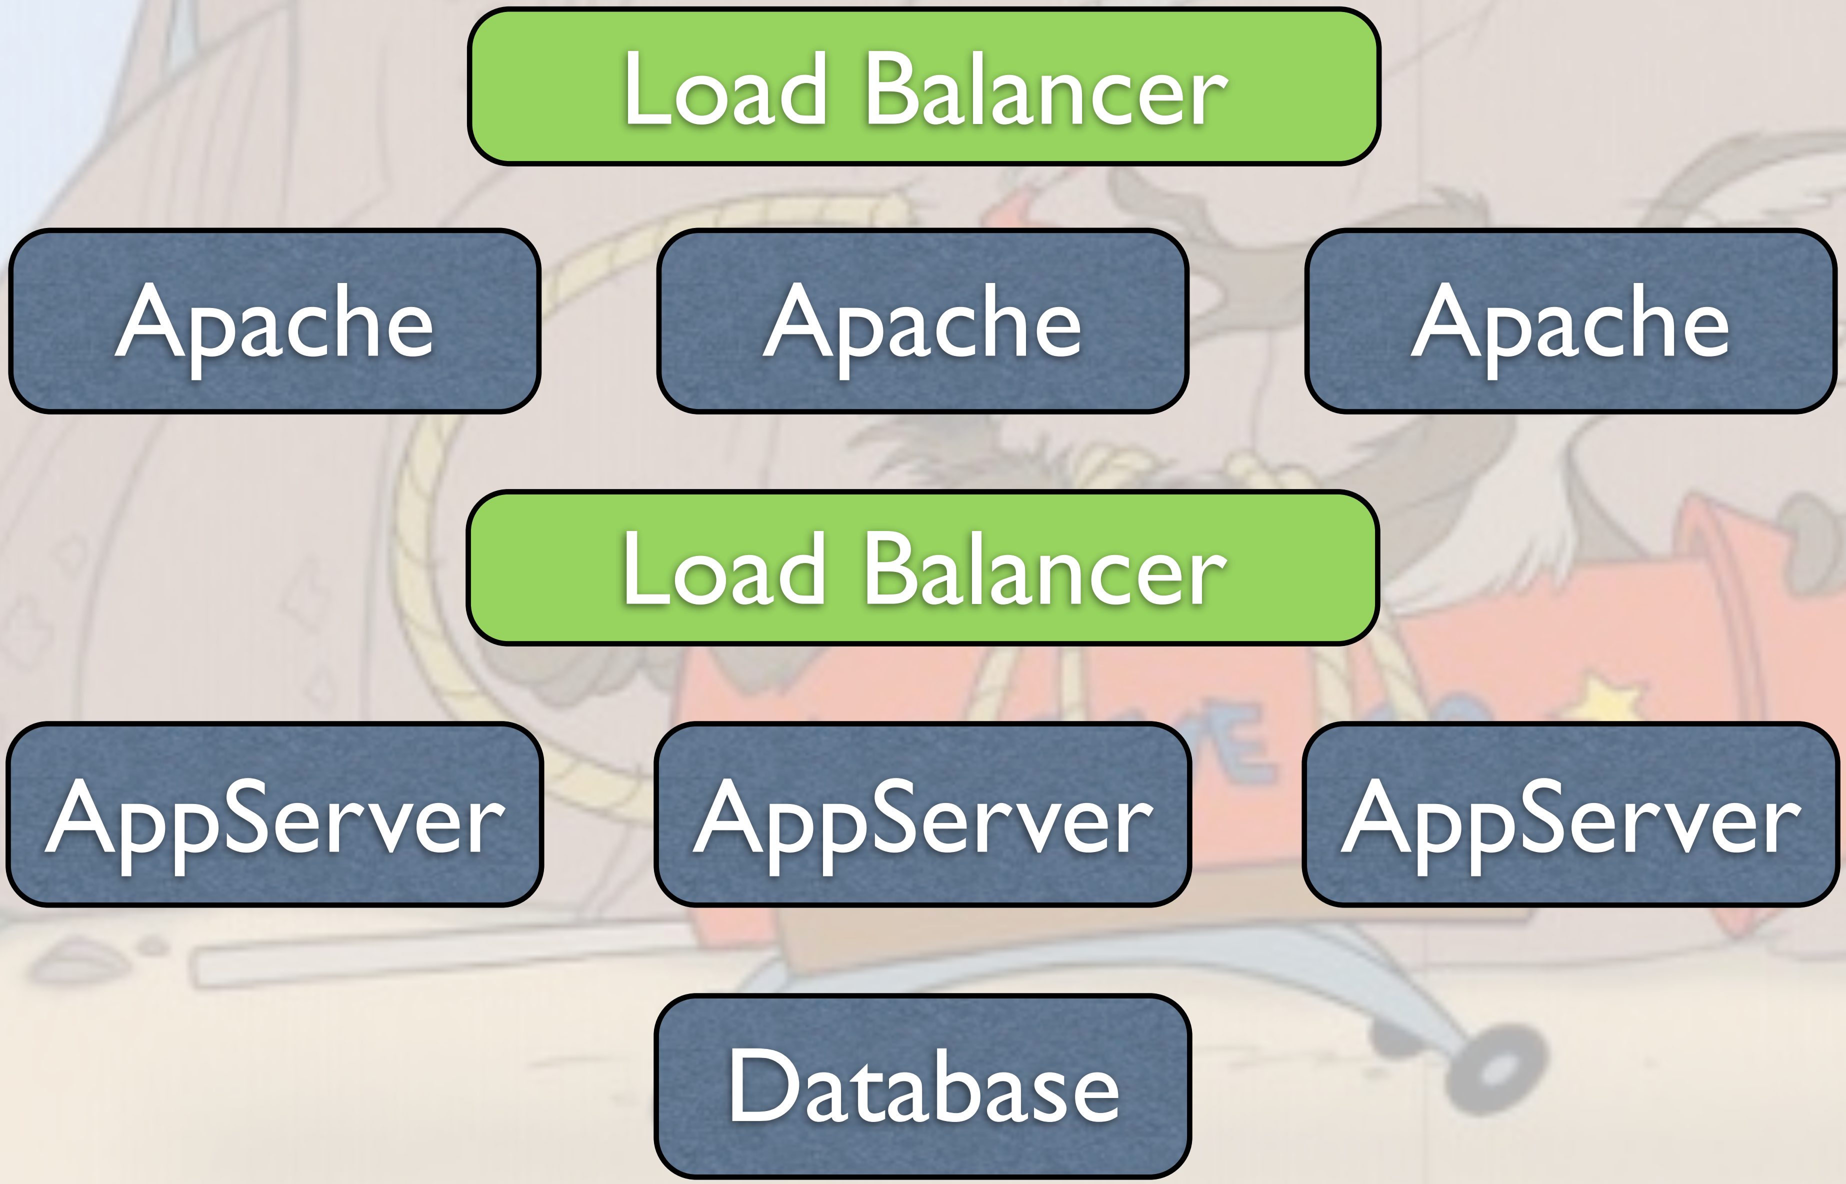

Load Balancer

Apache Apache Apache

Load Balancer

AppServer AppServer AppServer

Database

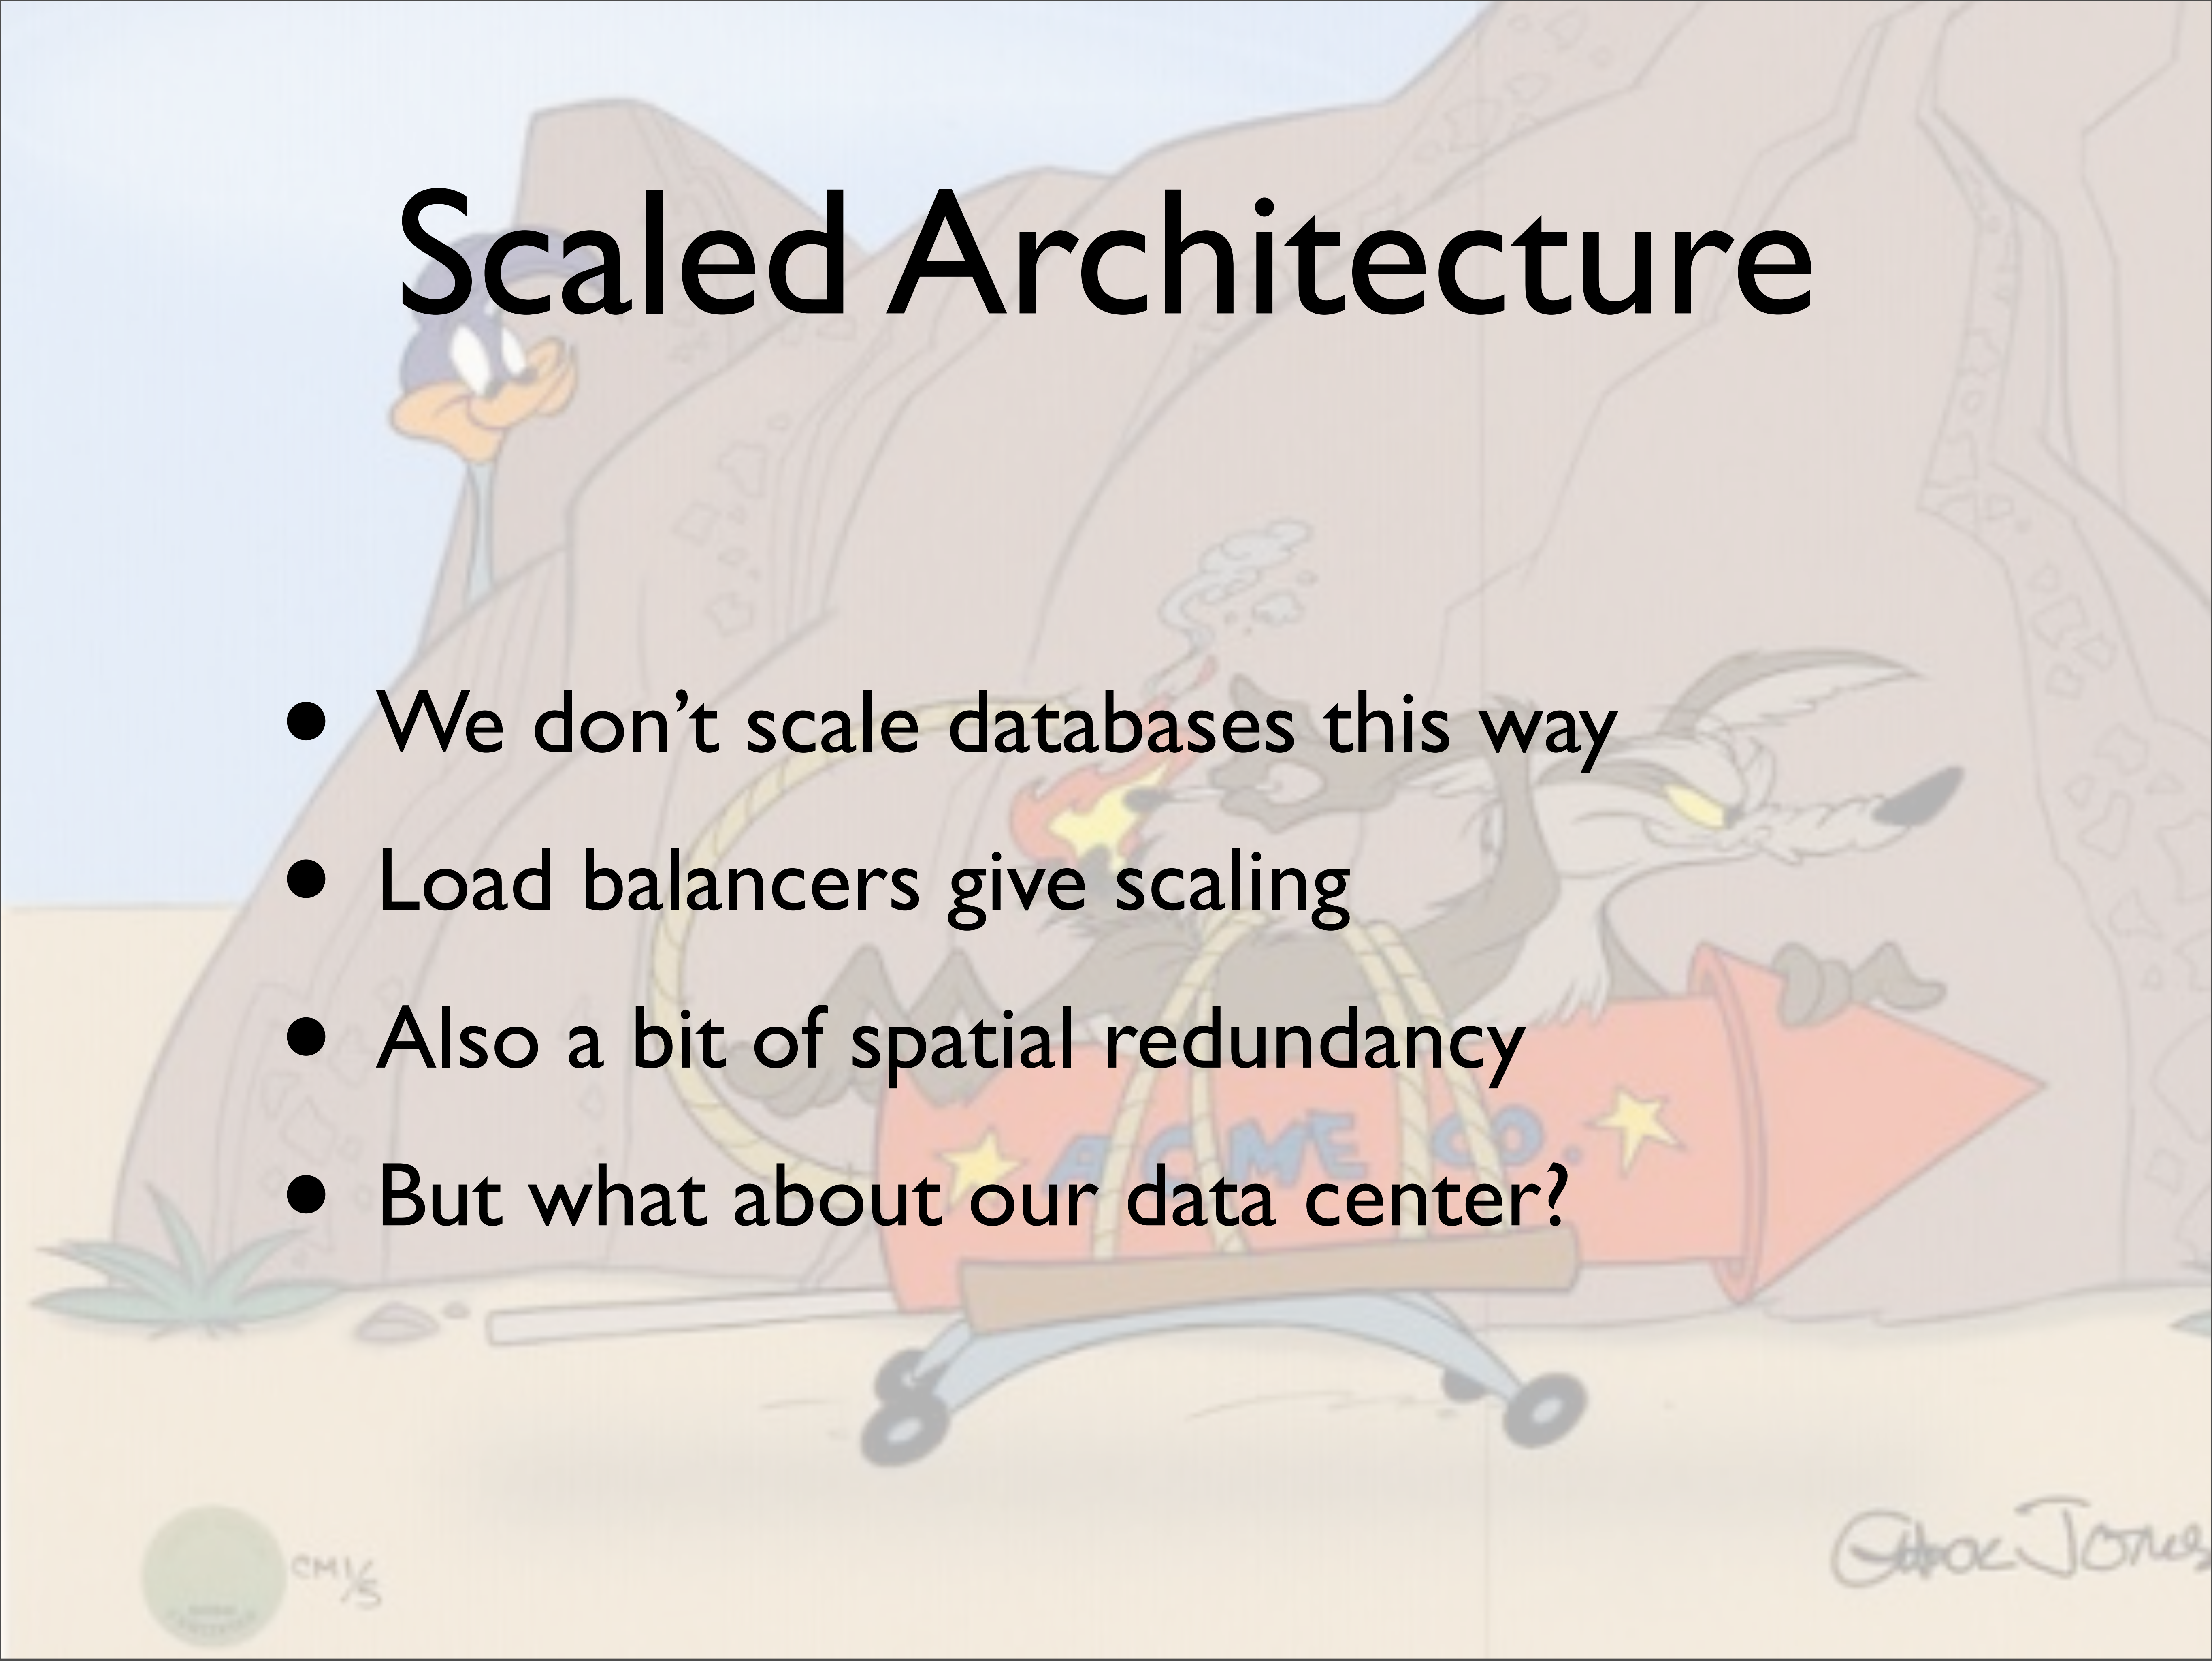

Scaled Architecture

- We don’t scale databases this way
- Load balancers give scaling
- Also a bit of spatial redundancy
- But what about our data center?

Redundancy

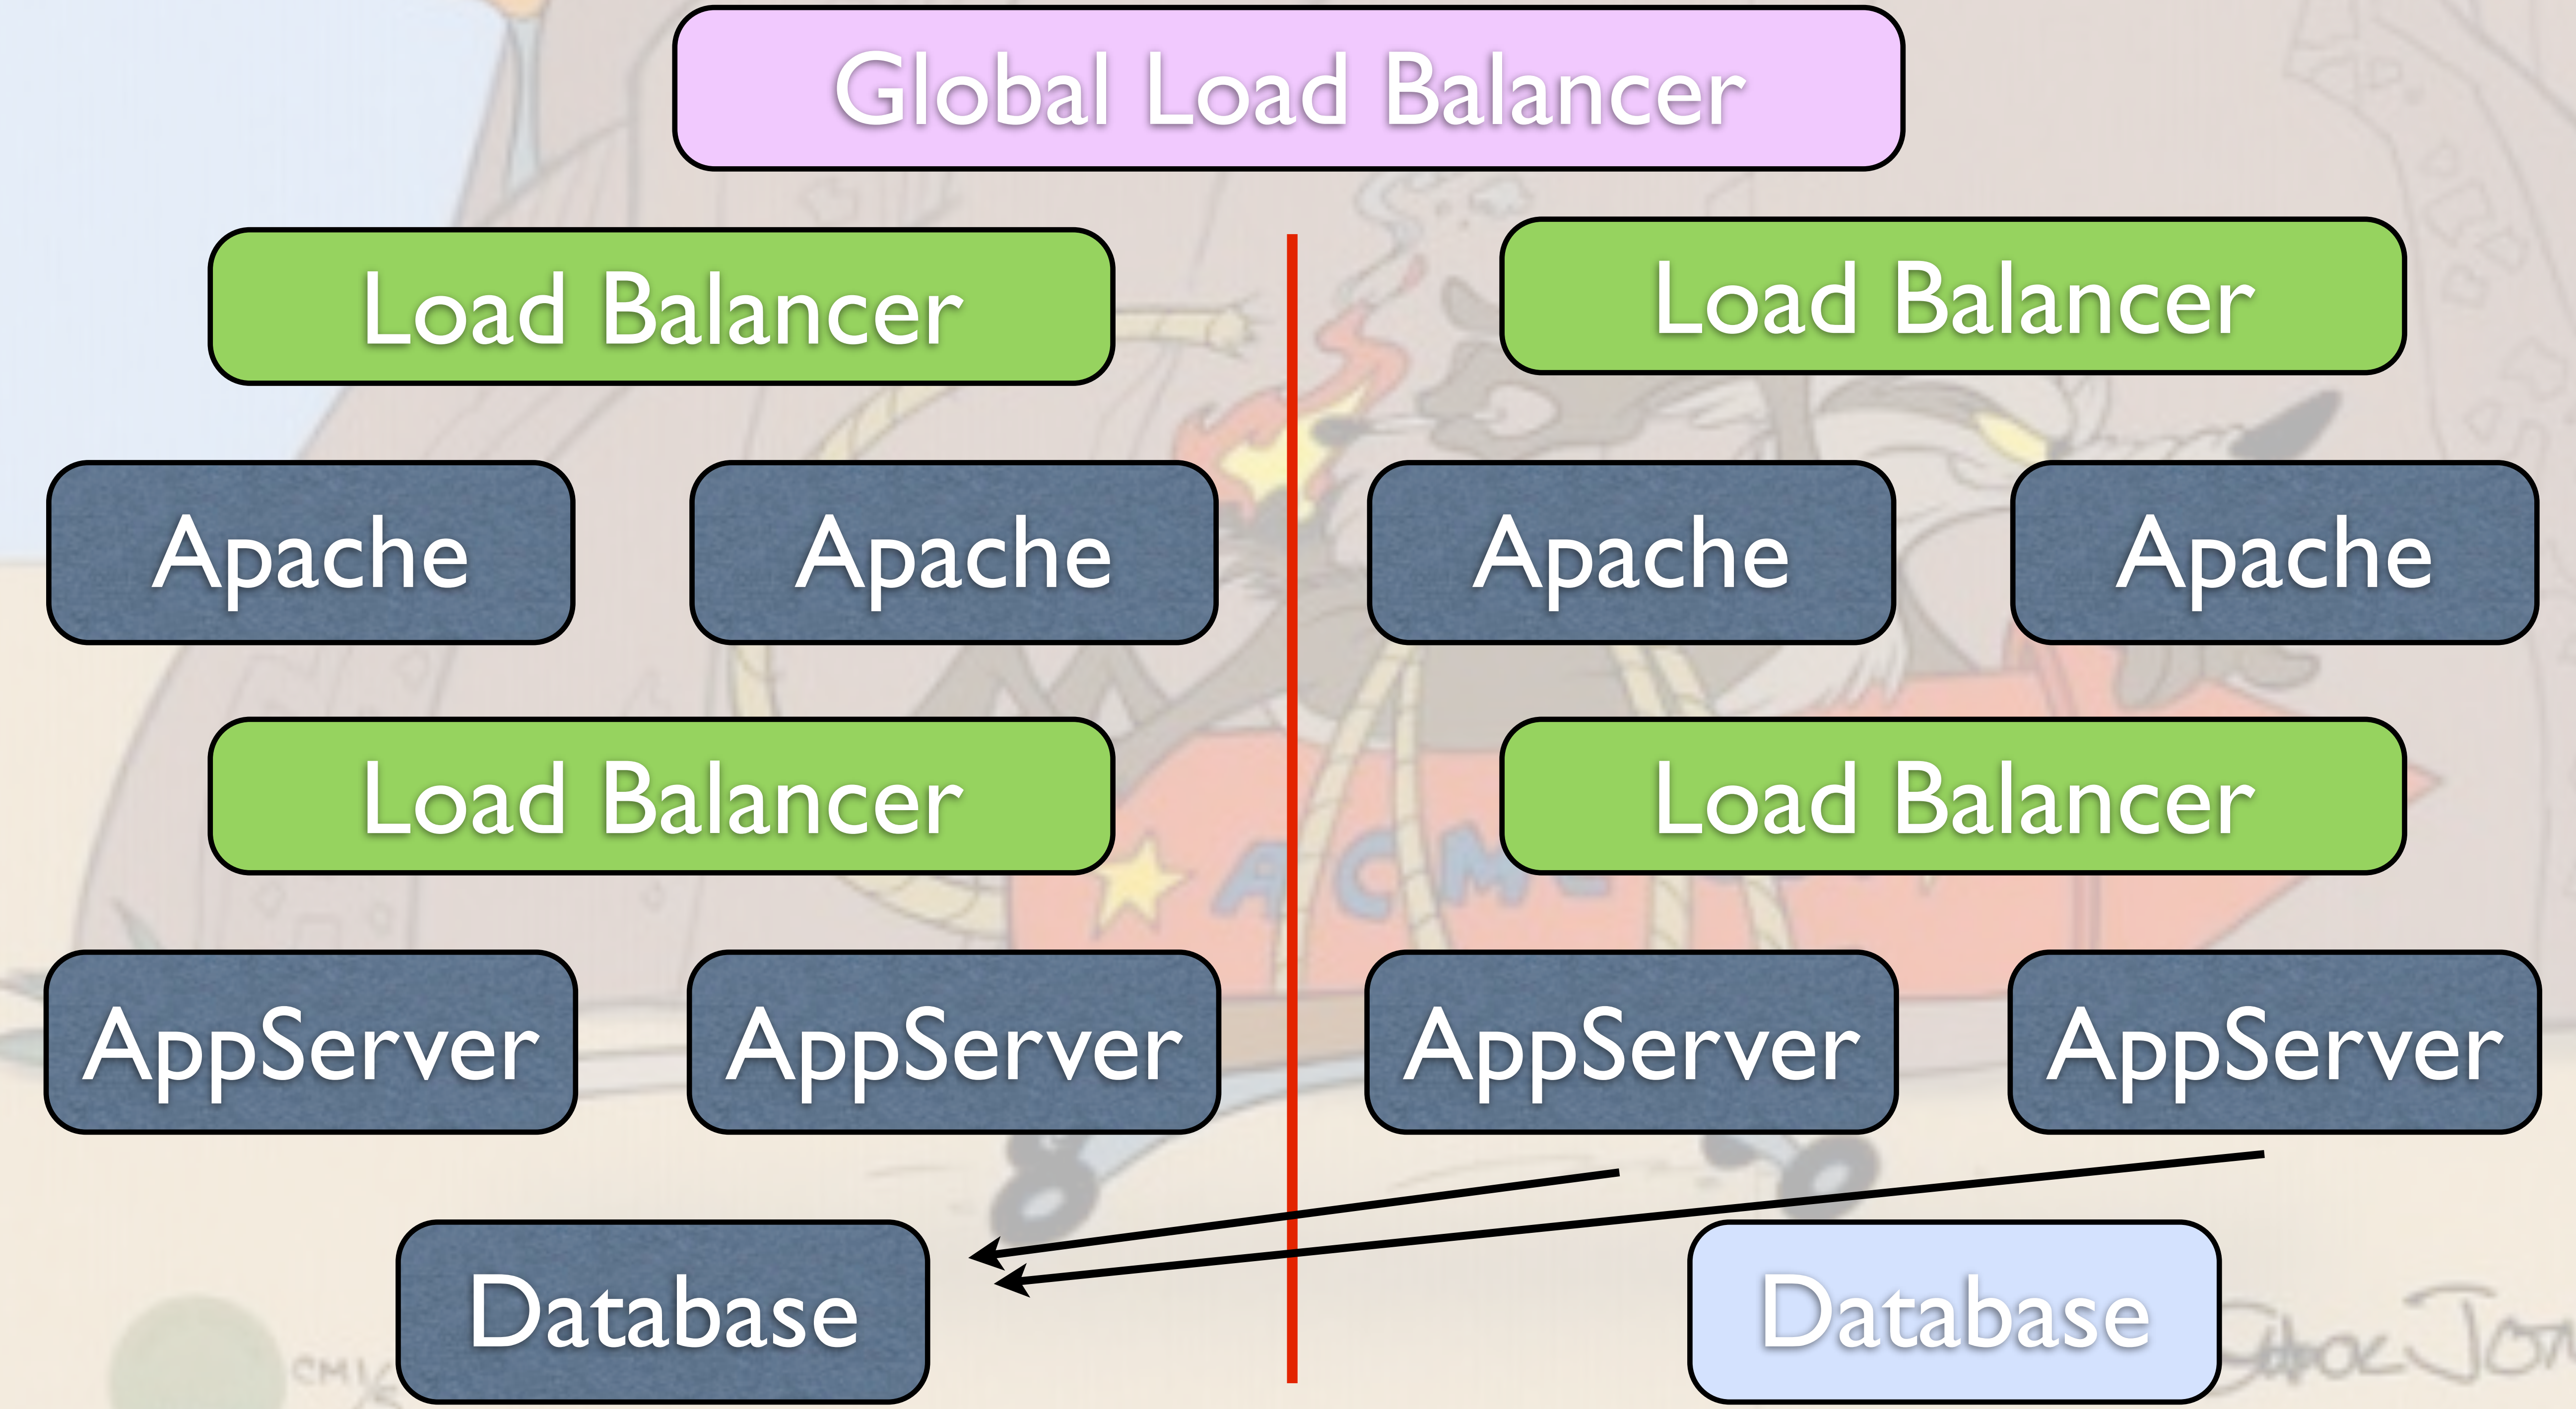

Global Load Balancer

Load Balancer

Load Balancer

Apache

Apache

Apache

Apache

Load Balancer

Load Balancer

AppServer

AppServer

AppServer

AppServer

Database

Database

Redundant Architecture

- Real spatial redundancy
- Global load balancing via DNS
- Twin datacenters
  - Redundant power
  - Redundant internet connectivity
- Database in Active/Passive

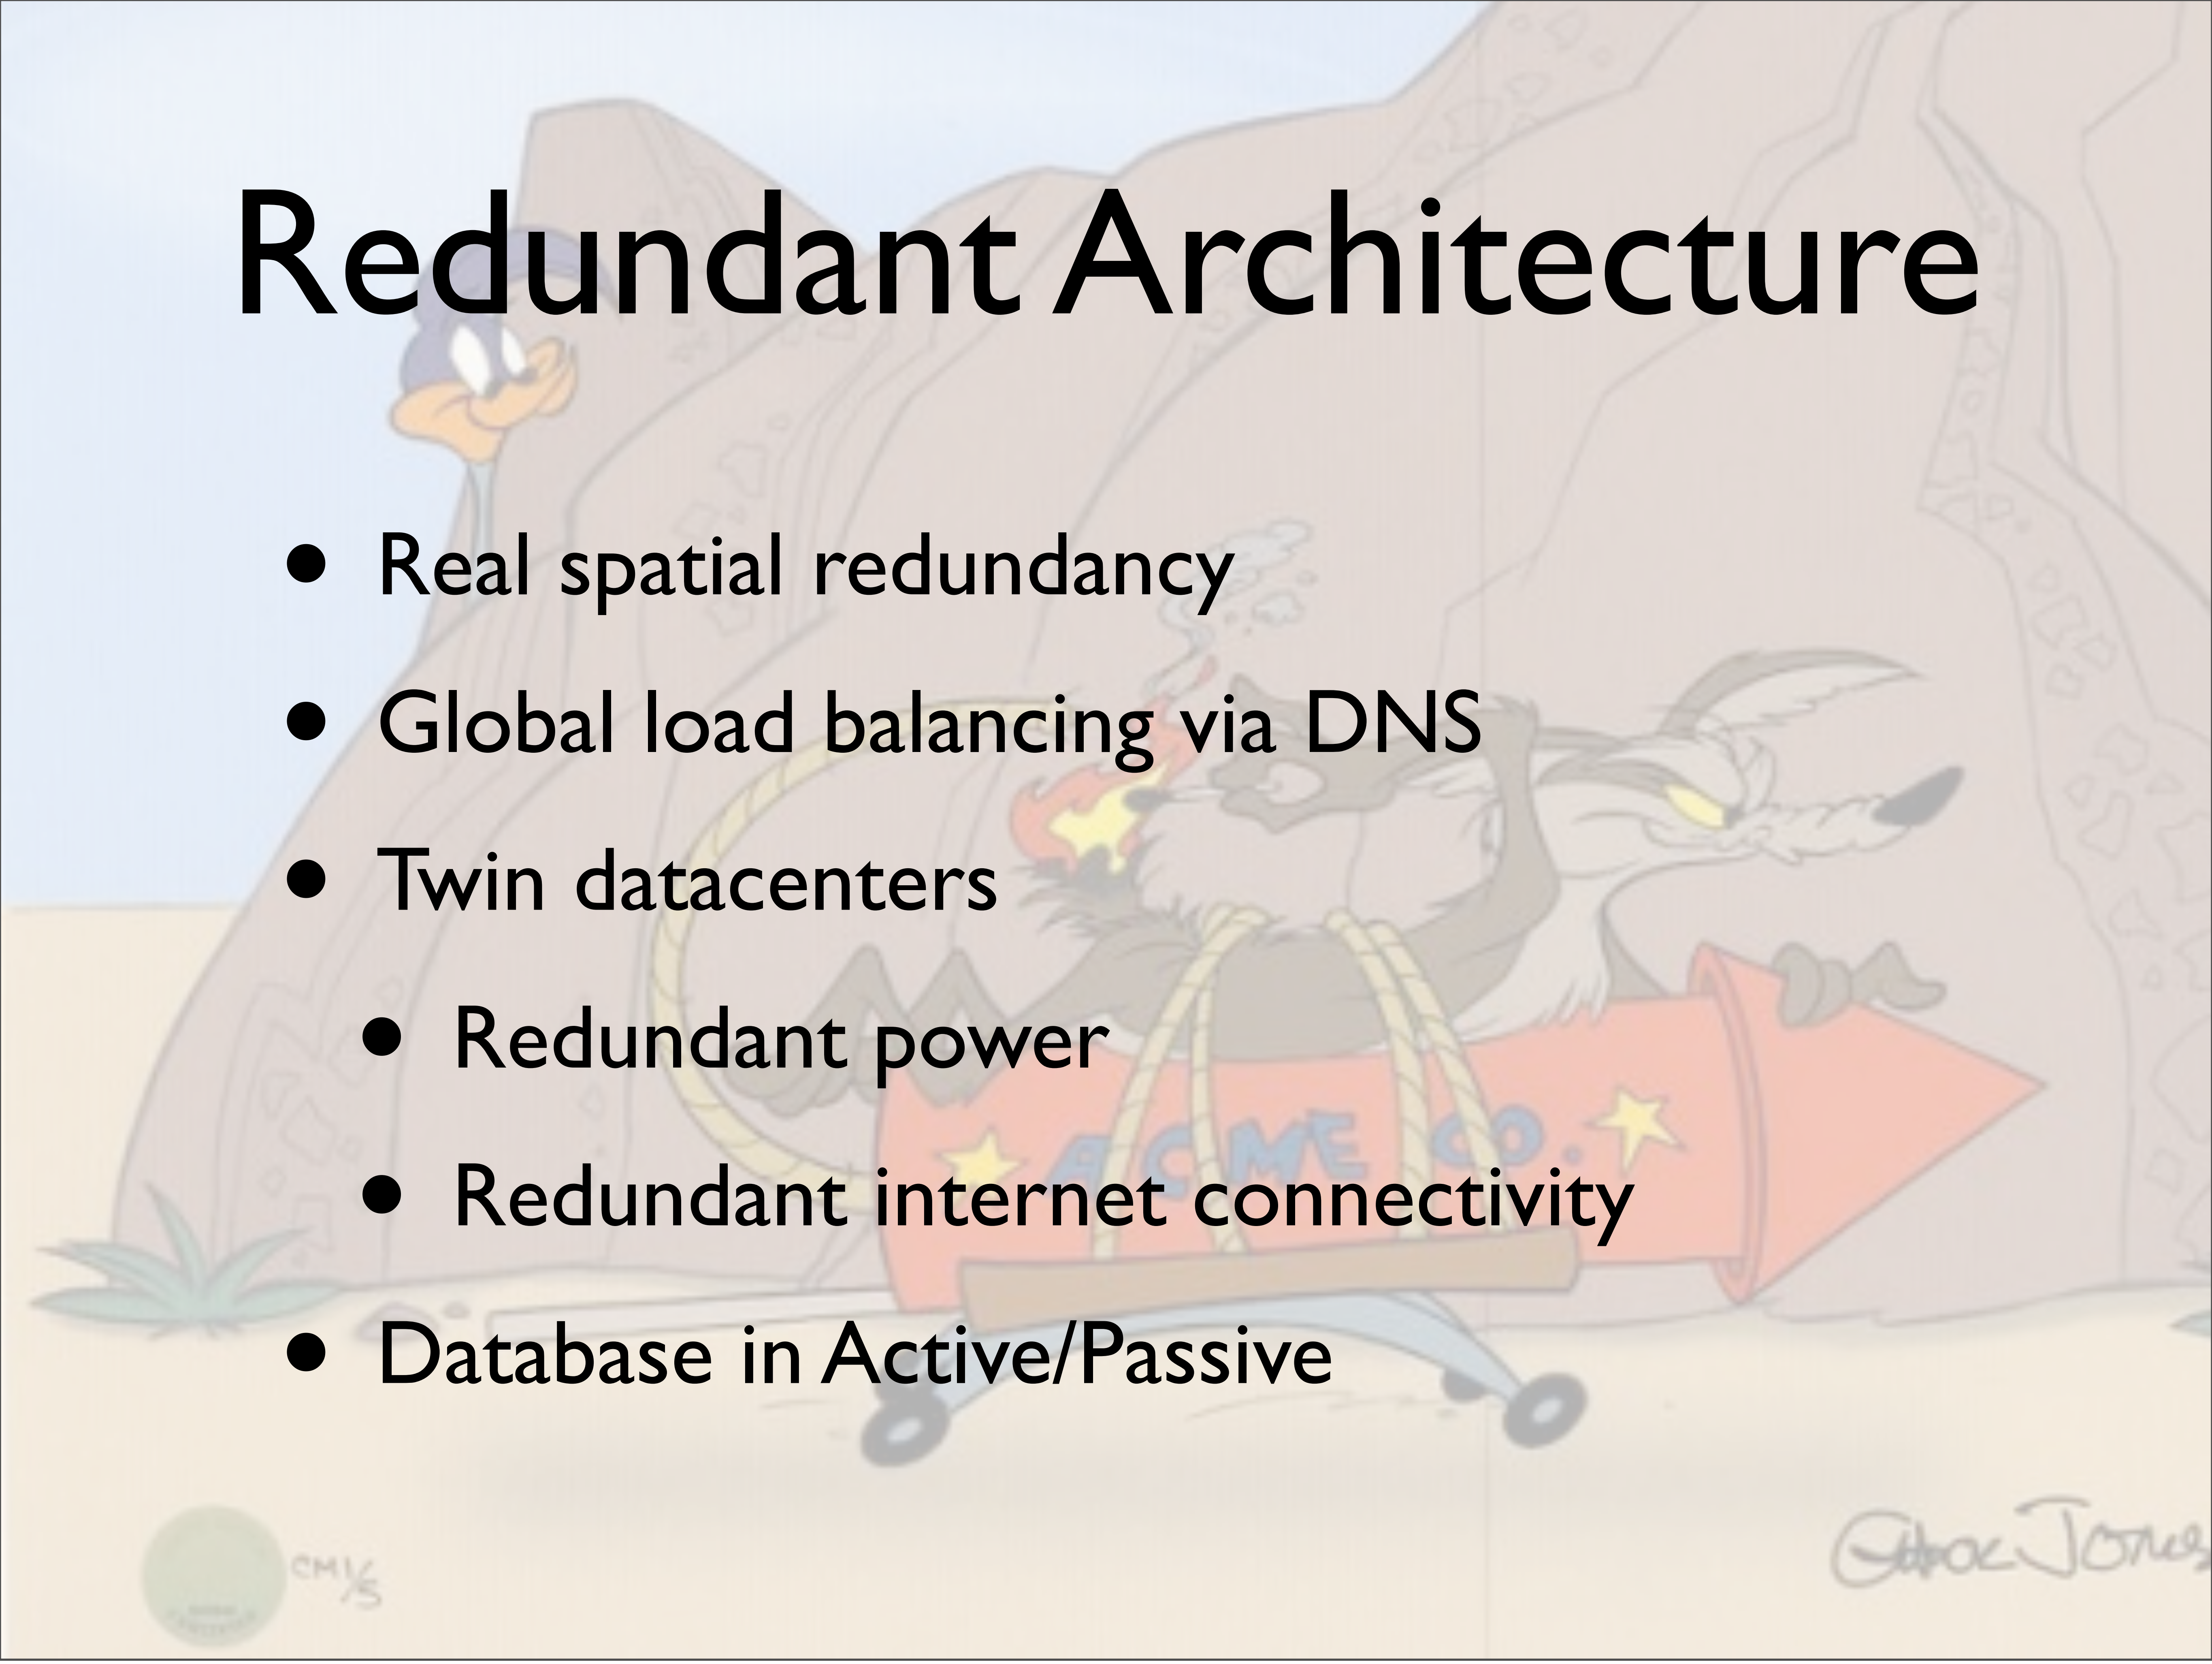

Success stories

- Serves 3.5m unique daily browsers
- Over 1.6m unique pieces of content
- supports hundreds of editorial staff
- create articles, audio, video, galleries, interactives, micro-sites

Drawbacks

- Monolithic system
  - Understands everything
    - football leagues
    - financial markets
    - mortgage applications
    - content!

The microapps revolution circa 2011

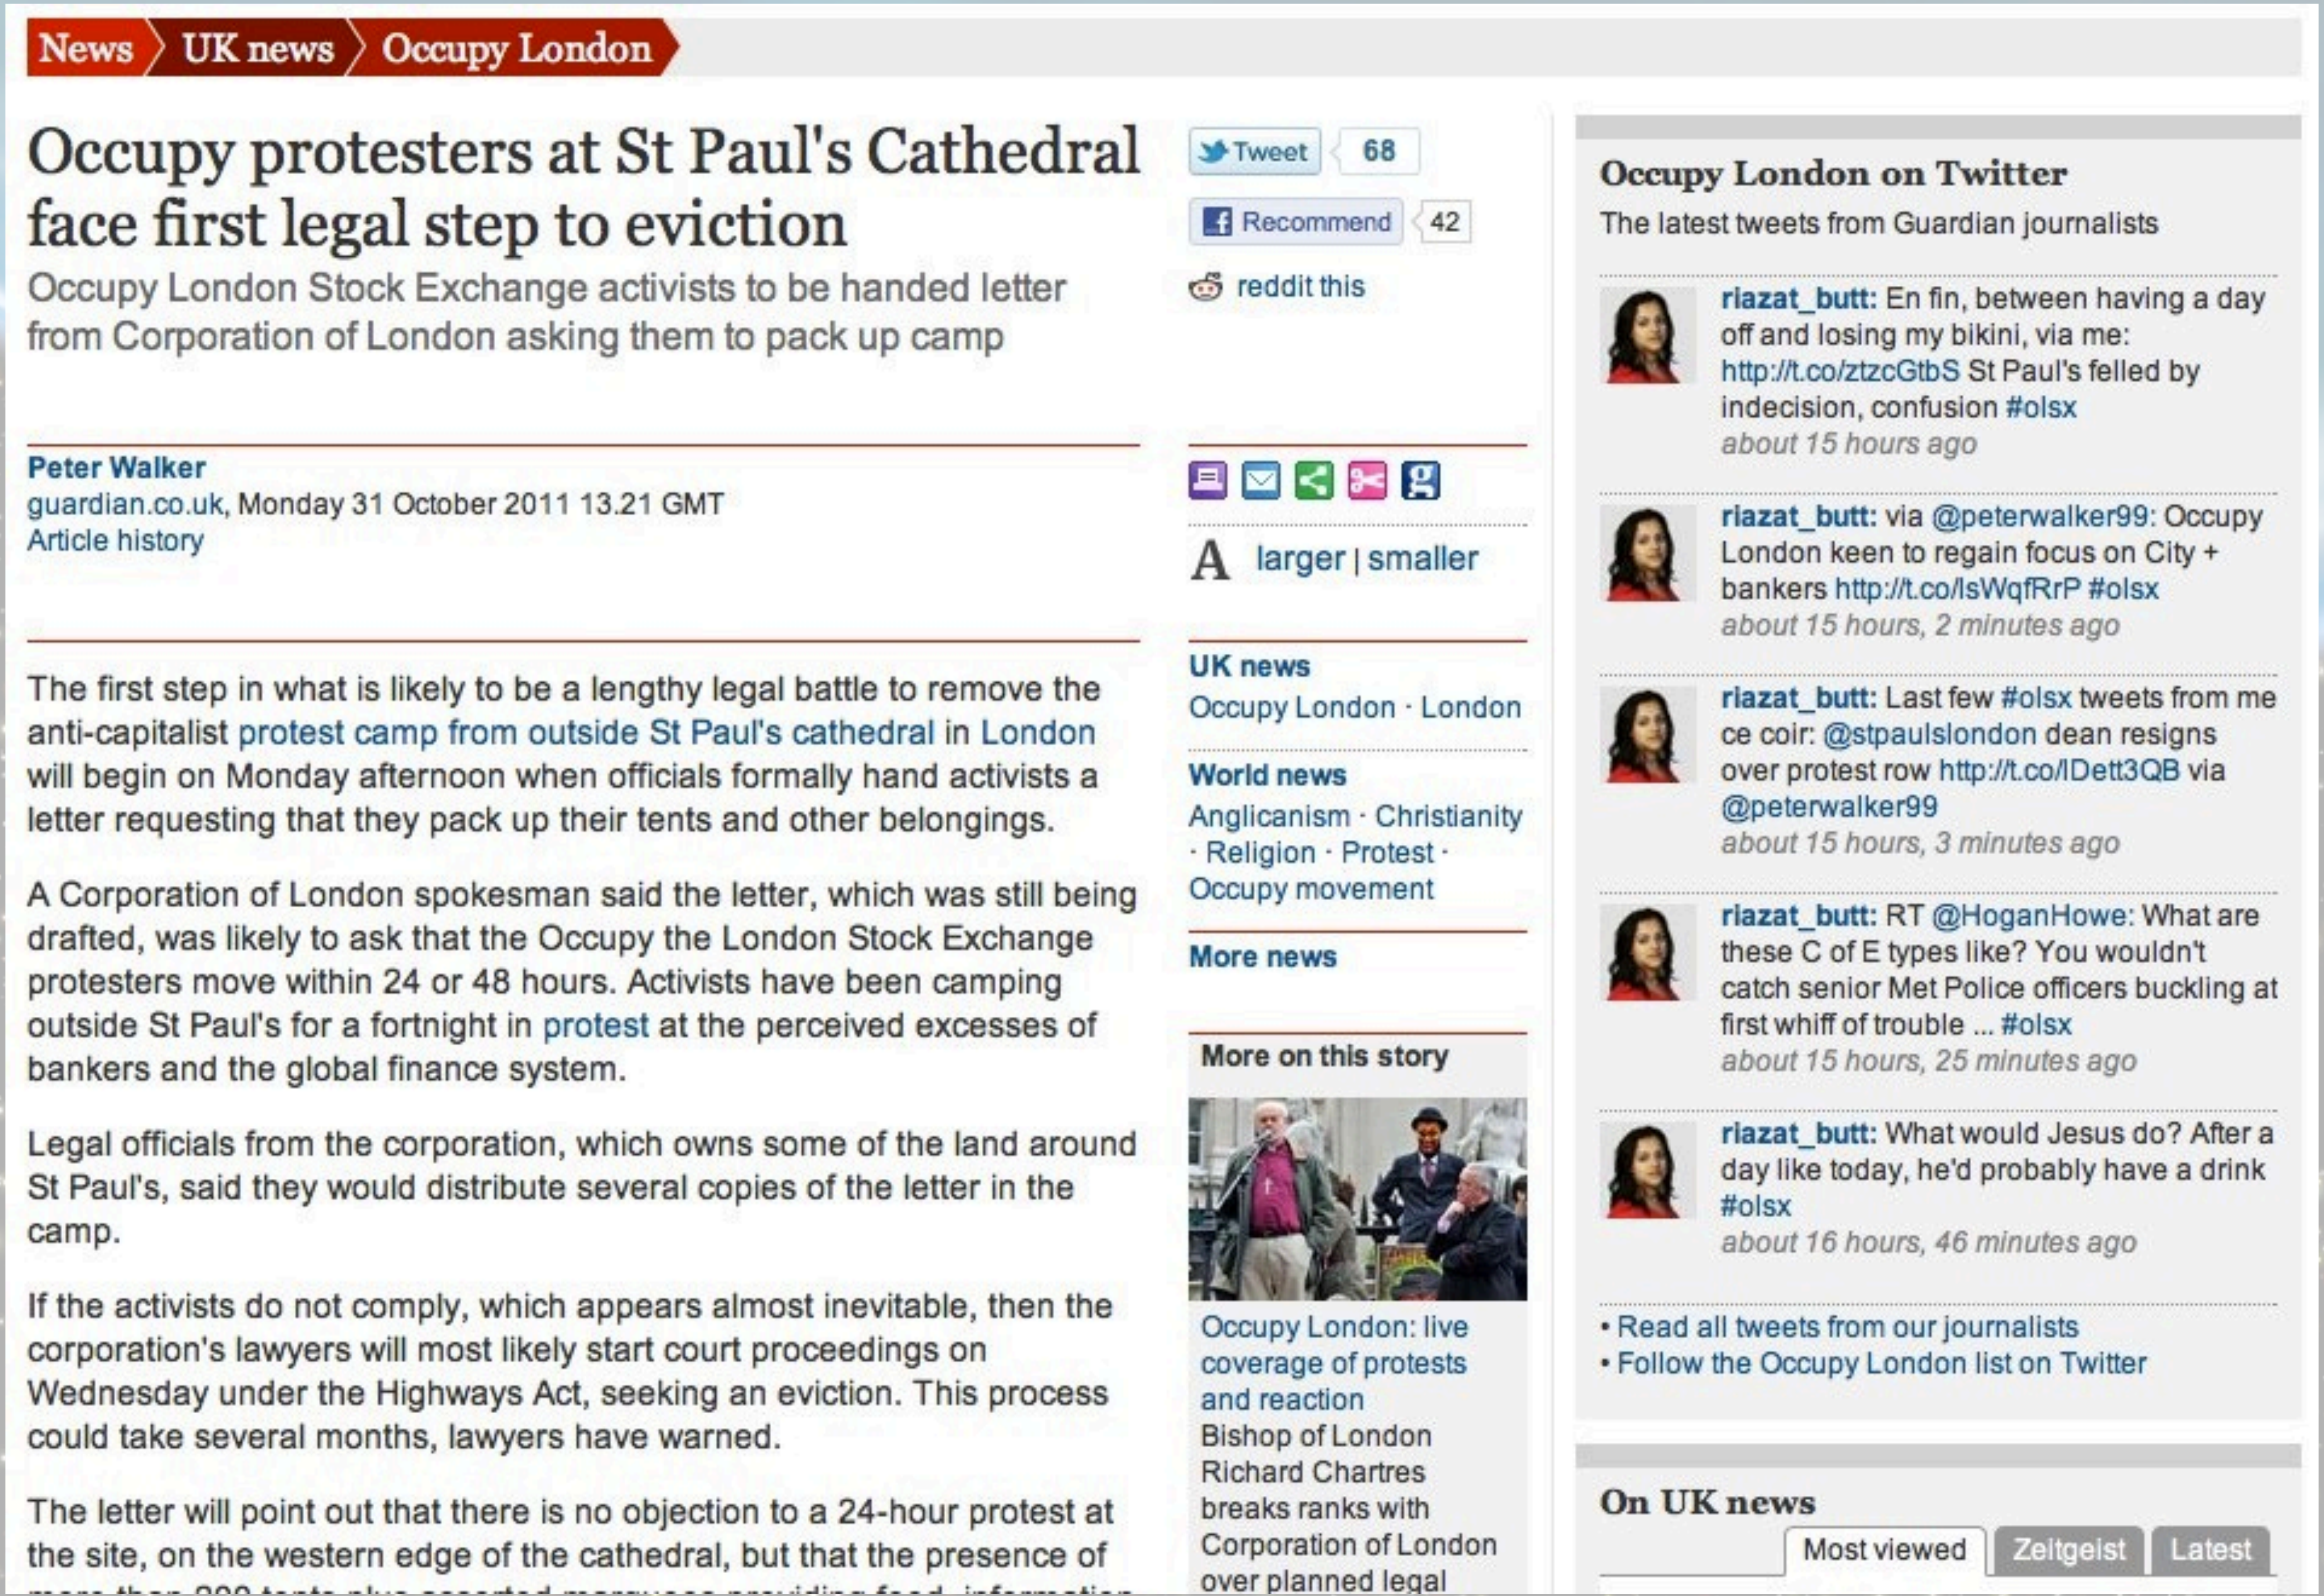

News › UK news › Occupy London

Occupy protesters at St Paul's Cathedral face first legal step to eviction

Occupy London Stock Exchange activists to be handed letter from Corporation of London asking them to pack up camp

Peter Walker
guardian.co.uk, Monday 31 October 2011 13.21 GMT
Article history

The first step in what is likely to be a lengthy legal battle to remove the anti-capitalist protest camp from outside St Paul's cathedral in London will begin on Monday afternoon when officials formally hand activists a letter requesting that they pack up their tents and other belongings.

A Corporation of London spokesman said the letter, which was still being drafted, was likely to ask that the Occupy the London Stock Exchange protesters move within 24 or 48 hours. Activists have been camping outside St Paul's for a fortnight in protest at the perceived excesses of bankers and the global finance system.

Legal officials from the corporation, which owns some of the land around St Paul's, said they would distribute several copies of the letter in the camp.

If the activists do not comply, which appears almost inevitable, then the corporation's lawyers will most likely start court proceedings on Wednesday under the Highways Act, seeking an eviction. This process could take several months, lawyers have warned.

The letter will point out that there is no objection to a 24-hour protest at the site, on the western edge of the cathedral, but that the presence of more than 200 tents plus assorted marquees providing food, information

UK news
Occupy London · London

World news
Anglicanism · Christianity · Religion · Protest · Occupy movement

More news

More on this story

Occupy London: live coverage of protests and reaction
Bishop of London Richard Chartres breaks ranks with Corporation of London over planned legal

Occupy London on Twitter

The latest tweets from Guardian journalists

riazat_butt: En fin, between having a day off and losing my bikini, via me: http://t.co/ztzcGtbS St Paul's felled by indecision, confusion #olsx
about 15 hours ago

riazat_butt: via @peterwalker99: Occupy London keen to regain focus on City + bankers http://t.co/lsWqfRrP #olsx
about 15 hours, 2 minutes ago

riazat_butt: Last few #olsx tweets from me ce coir: @stpaulslondon dean resigns over protest row http://t.co/lDett3QB via @peterwalker99
about 15 hours, 3 minutes ago

riazat_butt: RT @HoganHowe: What are these C of E types like? You wouldn't catch senior Met Police officers buckling at first whiff of trouble ... #olsx
about 15 hours, 25 minutes ago

riazat_butt: What would Jesus do? After a day like today, he'd probably have a drink #olsx
about 16 hours, 46 minutes ago

- Read all tweets from our journalists
- Follow the Occupy London list on Twitter

On UK news

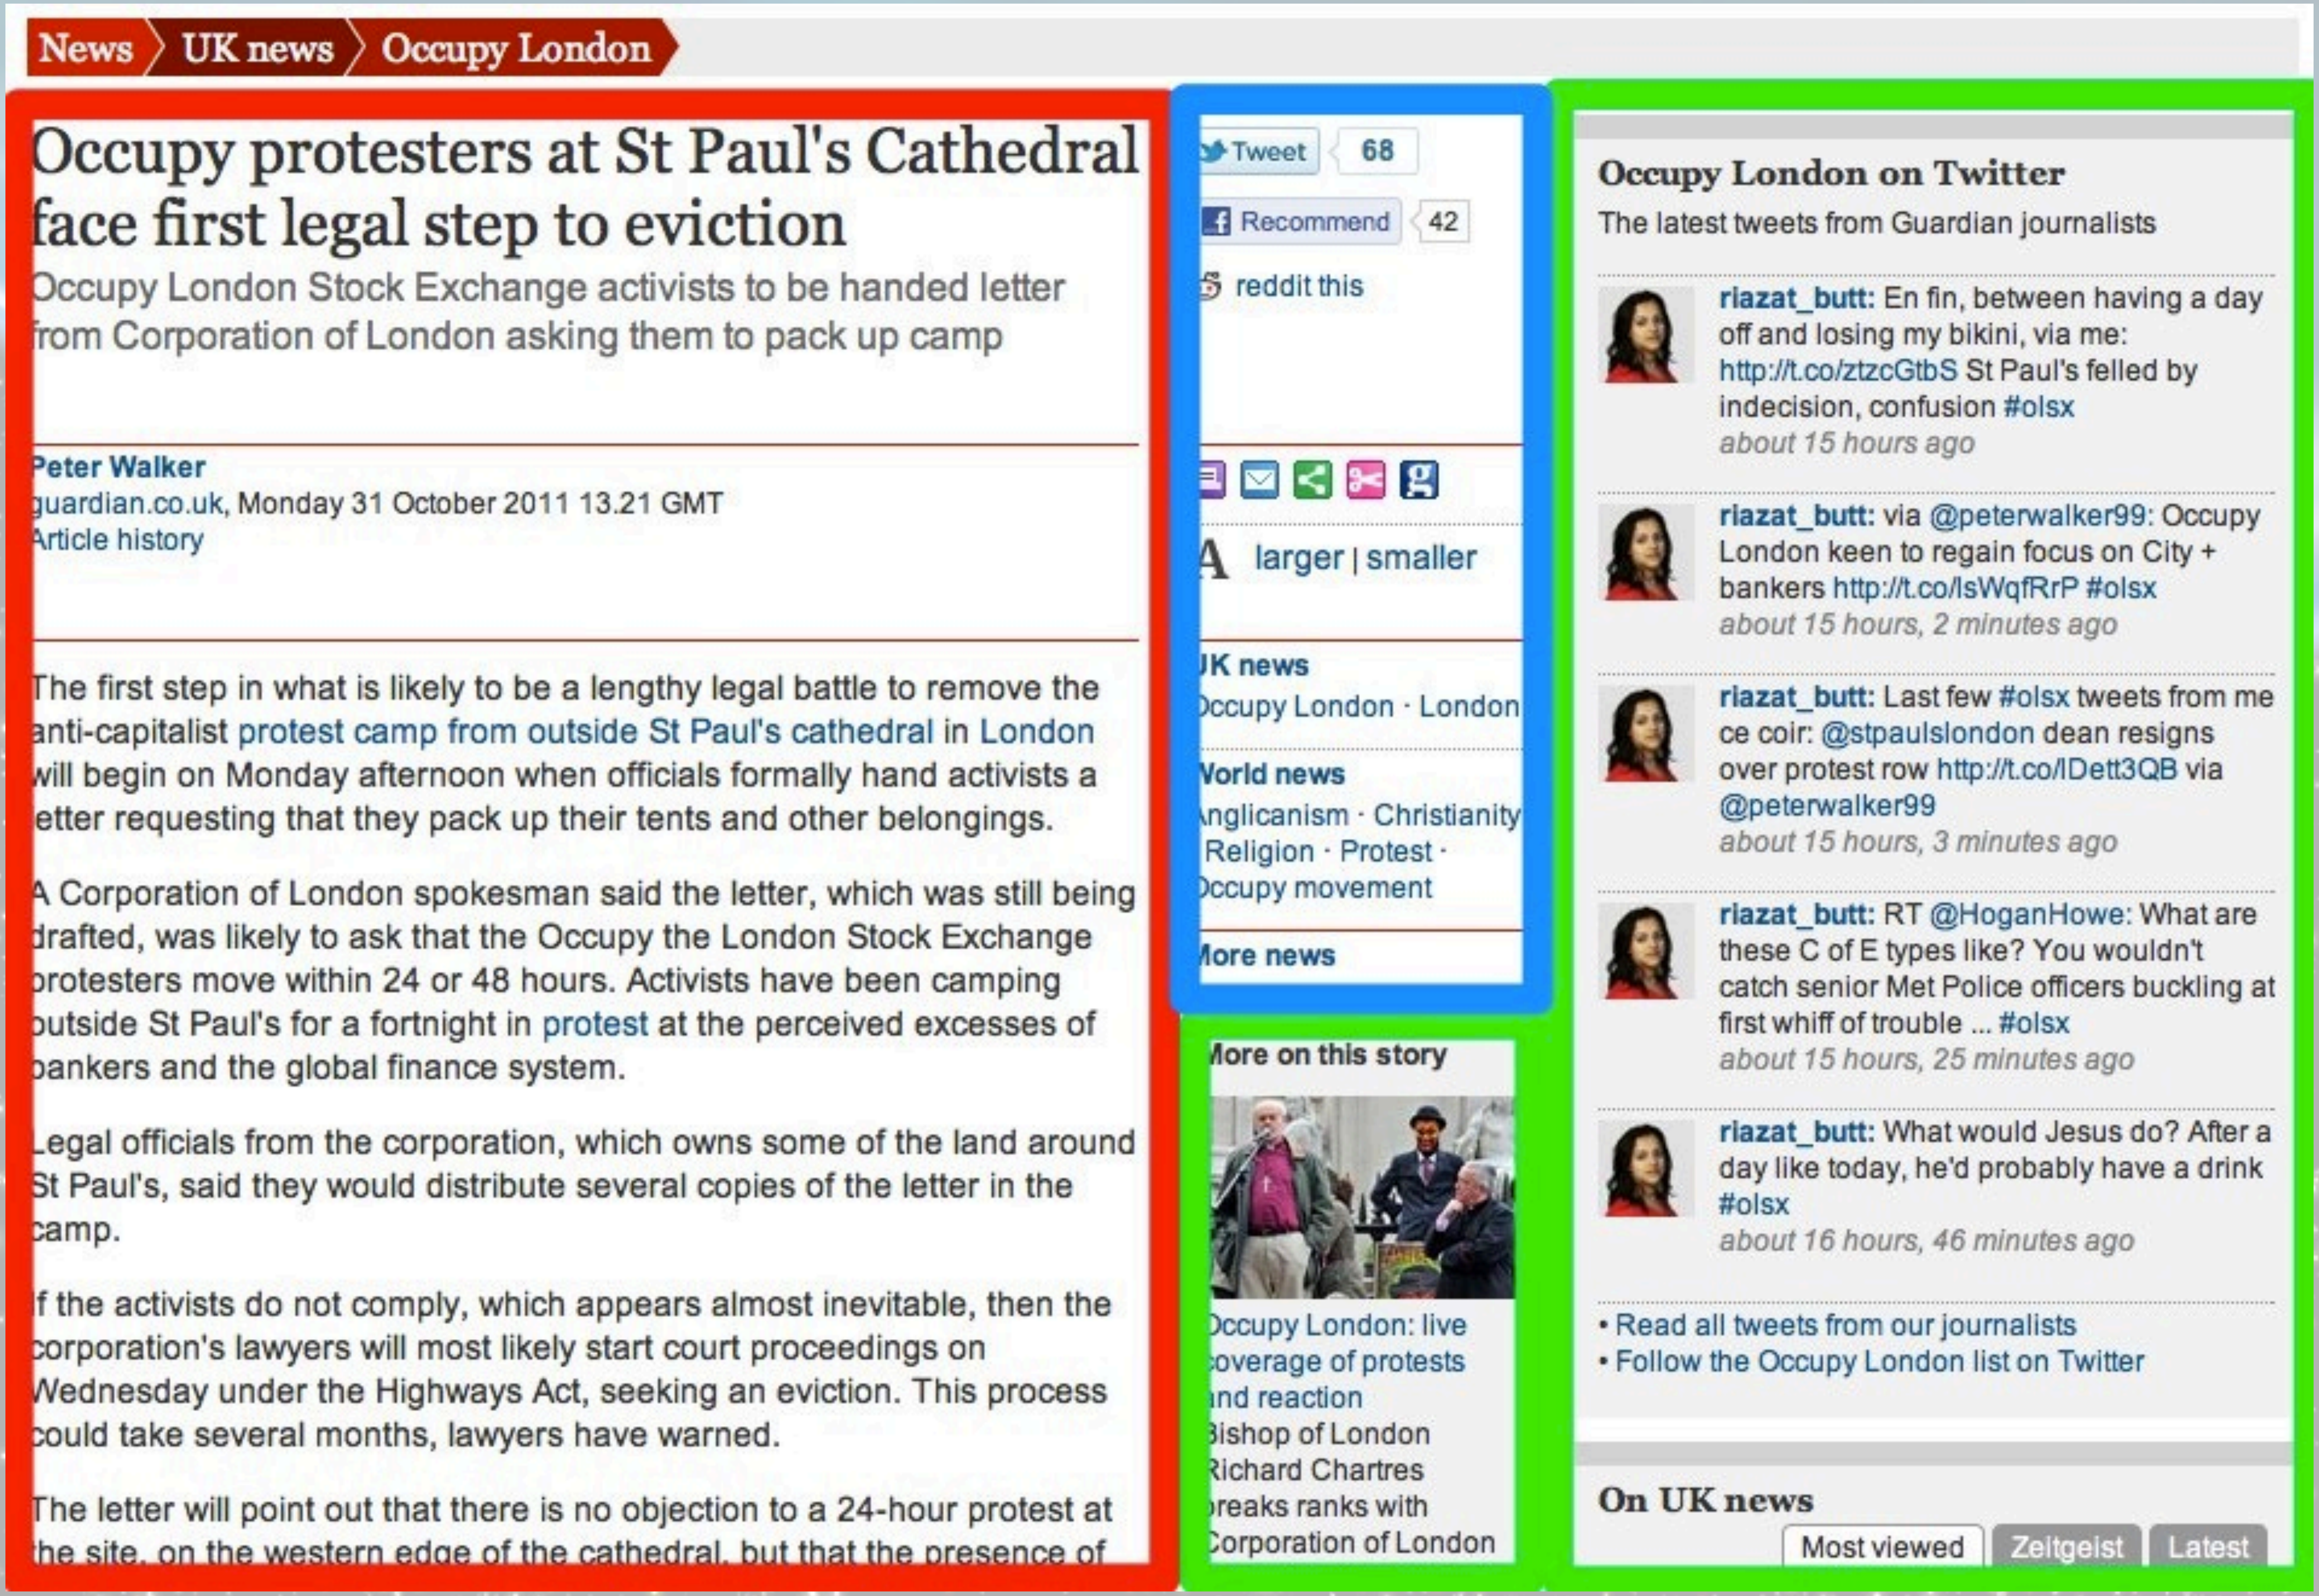

News › UK news › Occupy London

Occupy protesters at St Paul's Cathedral face first legal step to eviction

Occupy London Stock Exchange activists to be handed letter from Corporation of London asking them to pack up camp

Peter Walker
guardian.co.uk, Monday 31 October 2011 13.21 GMT
Article history

The first step in what is likely to be a lengthy legal battle to remove the anti-capitalist protest camp from outside St Paul's cathedral in London will begin on Monday afternoon when officials formally hand activists a letter requesting that they pack up their tents and other belongings.

A Corporation of London spokesman said the letter, which was still being drafted, was likely to ask that the Occupy the London Stock Exchange protesters move within 24 or 48 hours. Activists have been camping outside St Paul's for a fortnight in protest at the perceived excesses of bankers and the global finance system.

Legal officials from the corporation, which owns some of the land around St Paul's, said they would distribute several copies of the letter in the camp.

If the activists do not comply, which appears almost inevitable, then the corporation's lawyers will most likely start court proceedings on Wednesday under the Highways Act, seeking an eviction. This process could take several months, lawyers have warned.

The letter will point out that there is no objection to a 24-hour protest at the site, on the western edge of the cathedral, but that the presence of

UK news
Occupy London · London

World news
Anglicanism · Christianity · Religion · Protest · Occupy movement

More news

More on this story

Occupy London: live coverage of protests and reaction
Bishop of London Richard Chartres breaks ranks with Corporation of London

Occupy London on Twitter

The latest tweets from Guardian journalists

riazat_butt: En fin, between having a day off and losing my bikini, via me: http://t.co/ztzcGtbS St Paul's felled by indecision, confusion #olsx
about 15 hours ago

riazat_butt: via @peterwalker99: Occupy London keen to regain focus on City + bankers http://t.co/lsWqfRrP #olsx
about 15 hours, 2 minutes ago

riazat_butt: Last few #olsx tweets from me ce coir: @stpaulslondon dean resigns over protest row http://t.co/lDett3QB via @peterwalker99
about 15 hours, 3 minutes ago

riazat_butt: RT @HoganHowe: What are these C of E types like? You wouldn't catch senior Met Police officers buckling at first whiff of trouble ... #olsx
about 15 hours, 25 minutes ago

riazat_butt: What would Jesus do? After a day like today, he'd probably have a drink #olsx
about 16 hours, 46 minutes ago

- Read all tweets from our journalists
- Follow the Occupy London list on Twitter

On UK news

Microapps

- Separation of Systems
- SSI-like technology
- HTTP

AppEngine, Python, Ruby, EC2 - Oh My

- Proliferation of systems, languages and frameworks
- Faster development
- Increased innovation
- Hack Days!
- Built on content API

Microapps Architecture

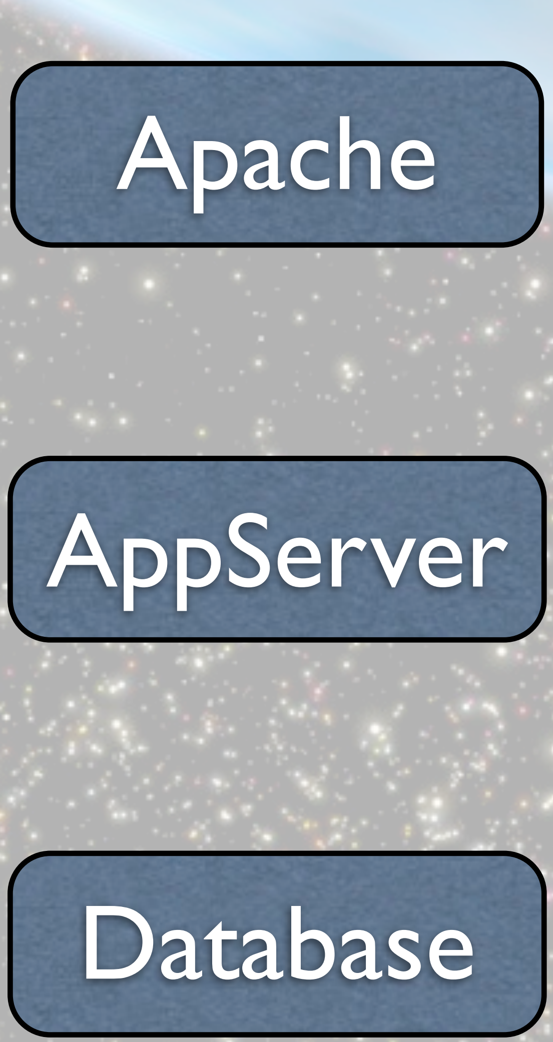

Apache

AppServer

Database

Microapps Architecture

Apache

AppServer

Database

EC2

AppEngine

Microapps Architecture

Apache	EC2
AppServer	AppEngine
Database	Content API (EC2)

Microapps Architecture

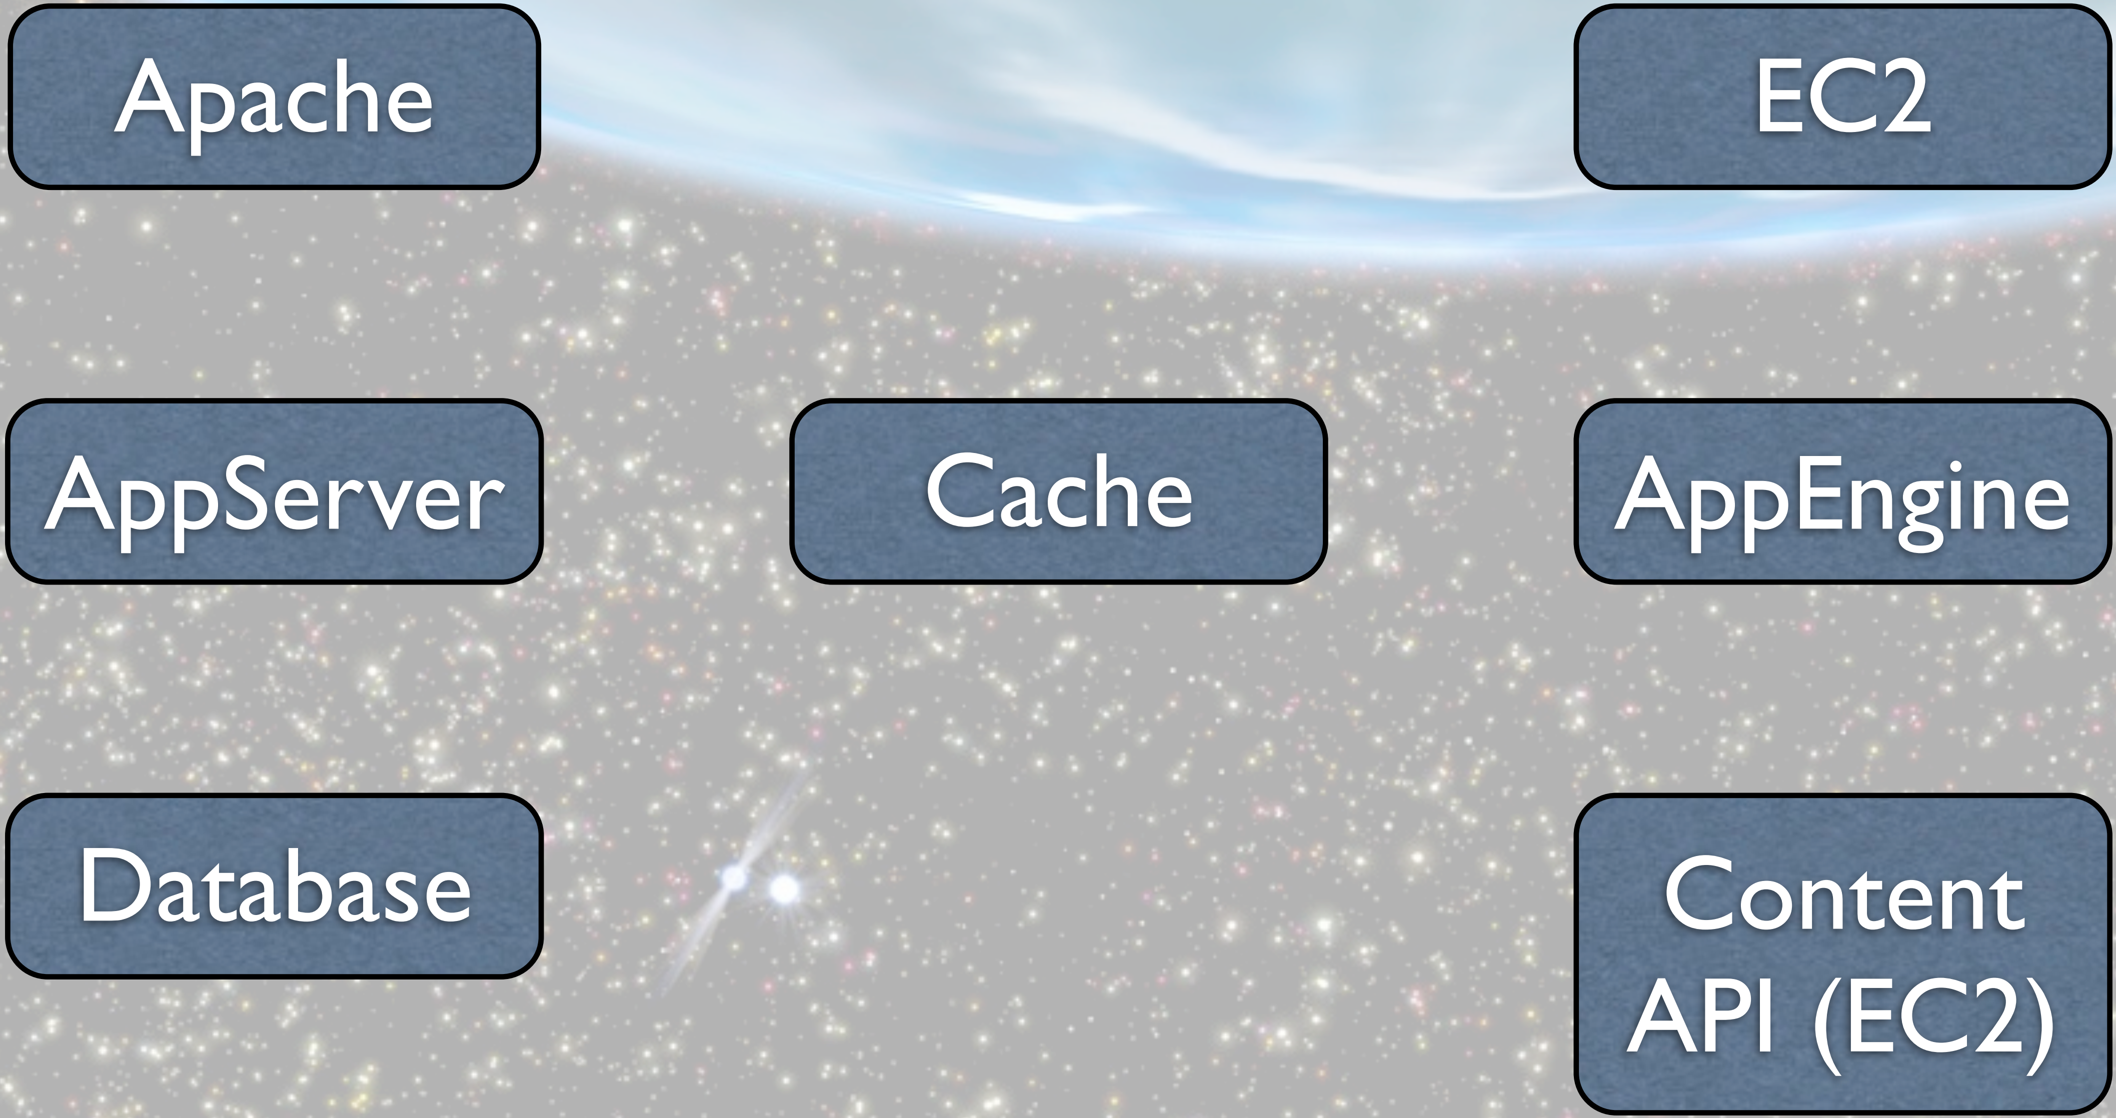

Apache

EC2

AppServer

Cache

AppEngine

Database

Content API (EC2)

The cost of diversification

- Support
- Maintenance
- Decided to settle on JVM stack primarily

Benefits

- Lots of small simple applications
- Can code, release, test in isolation
- Cache
  - max-age
  - stale-if-error

Drawbacks

- Increased architectural complexity
  - Need a big cache
  - Context

The biggest problem

- Microapp latency affects CMS latency
- Failure is not a problem
- Slow is a problem
  - stale-while-revalidate?

Emergency Mode

Emergency Mode

- Dynamic pages are expensive
- ‘Peaky’ traffic
- Often small subset of functionality
- Trade off dynamism for speed

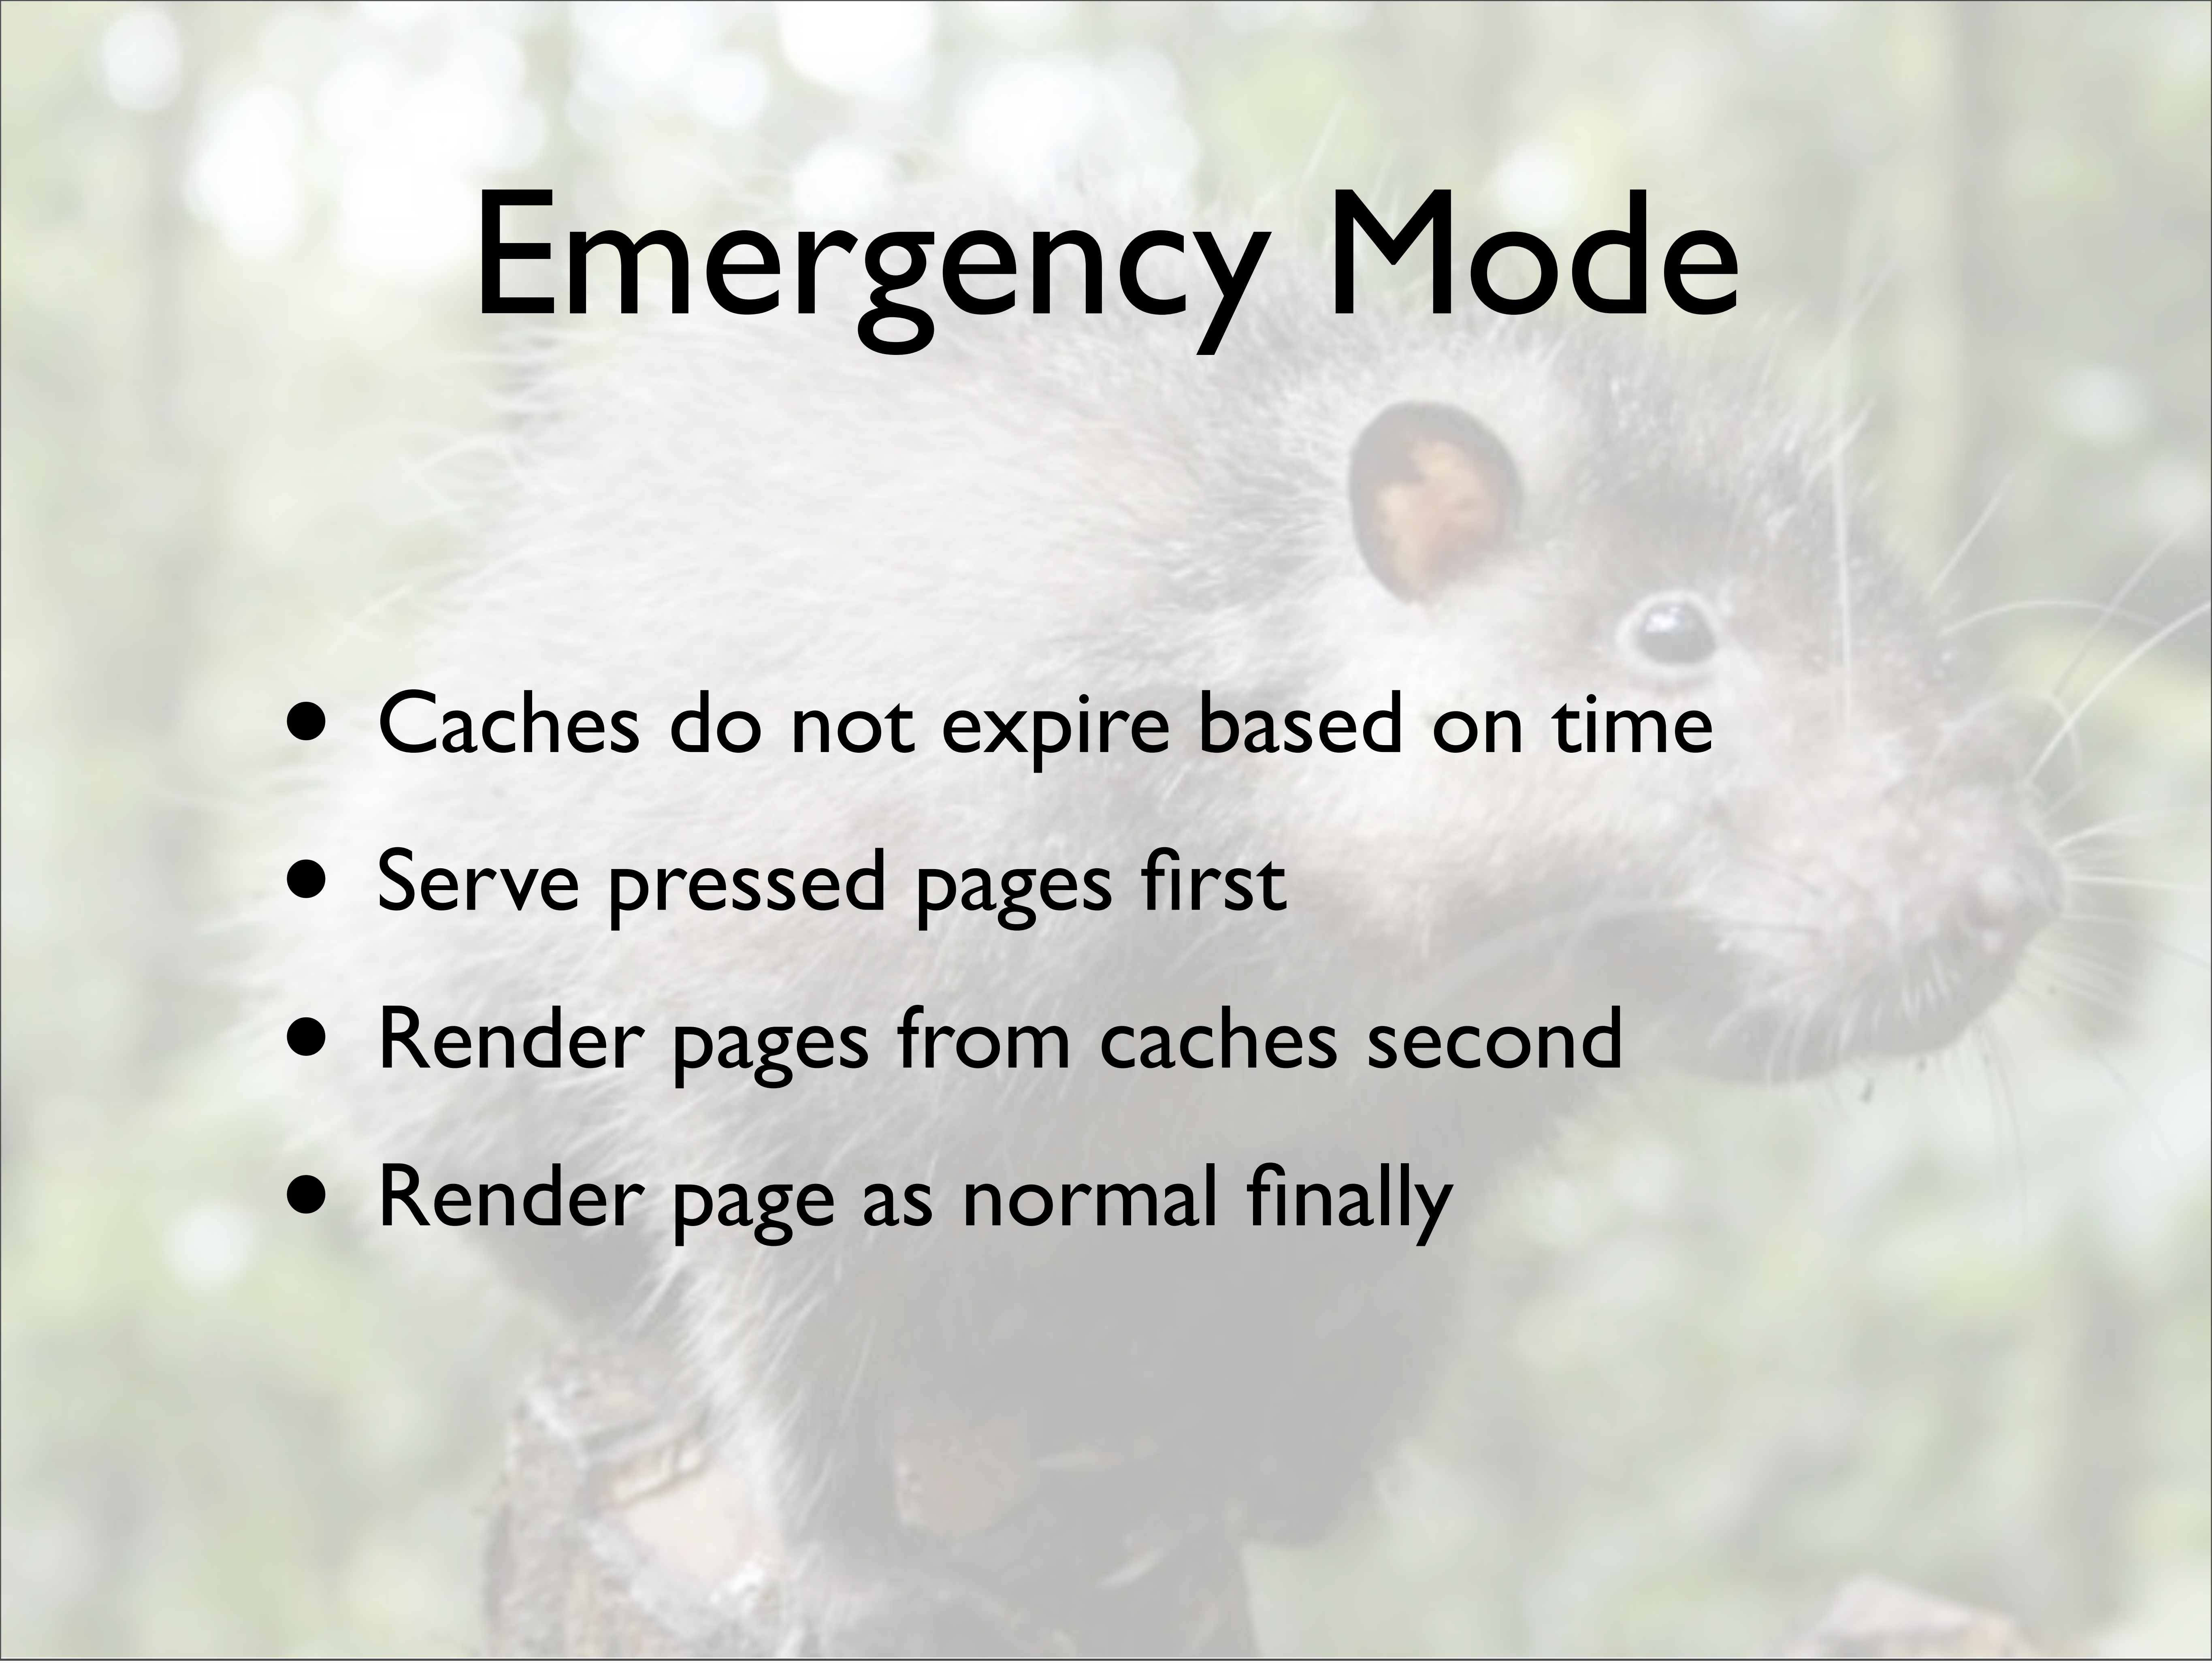

Emergency Mode

- Caches do not expire based on time
- Serve pressed pages first
- Render pages from caches second
- Render page as normal finally

Page Pressing

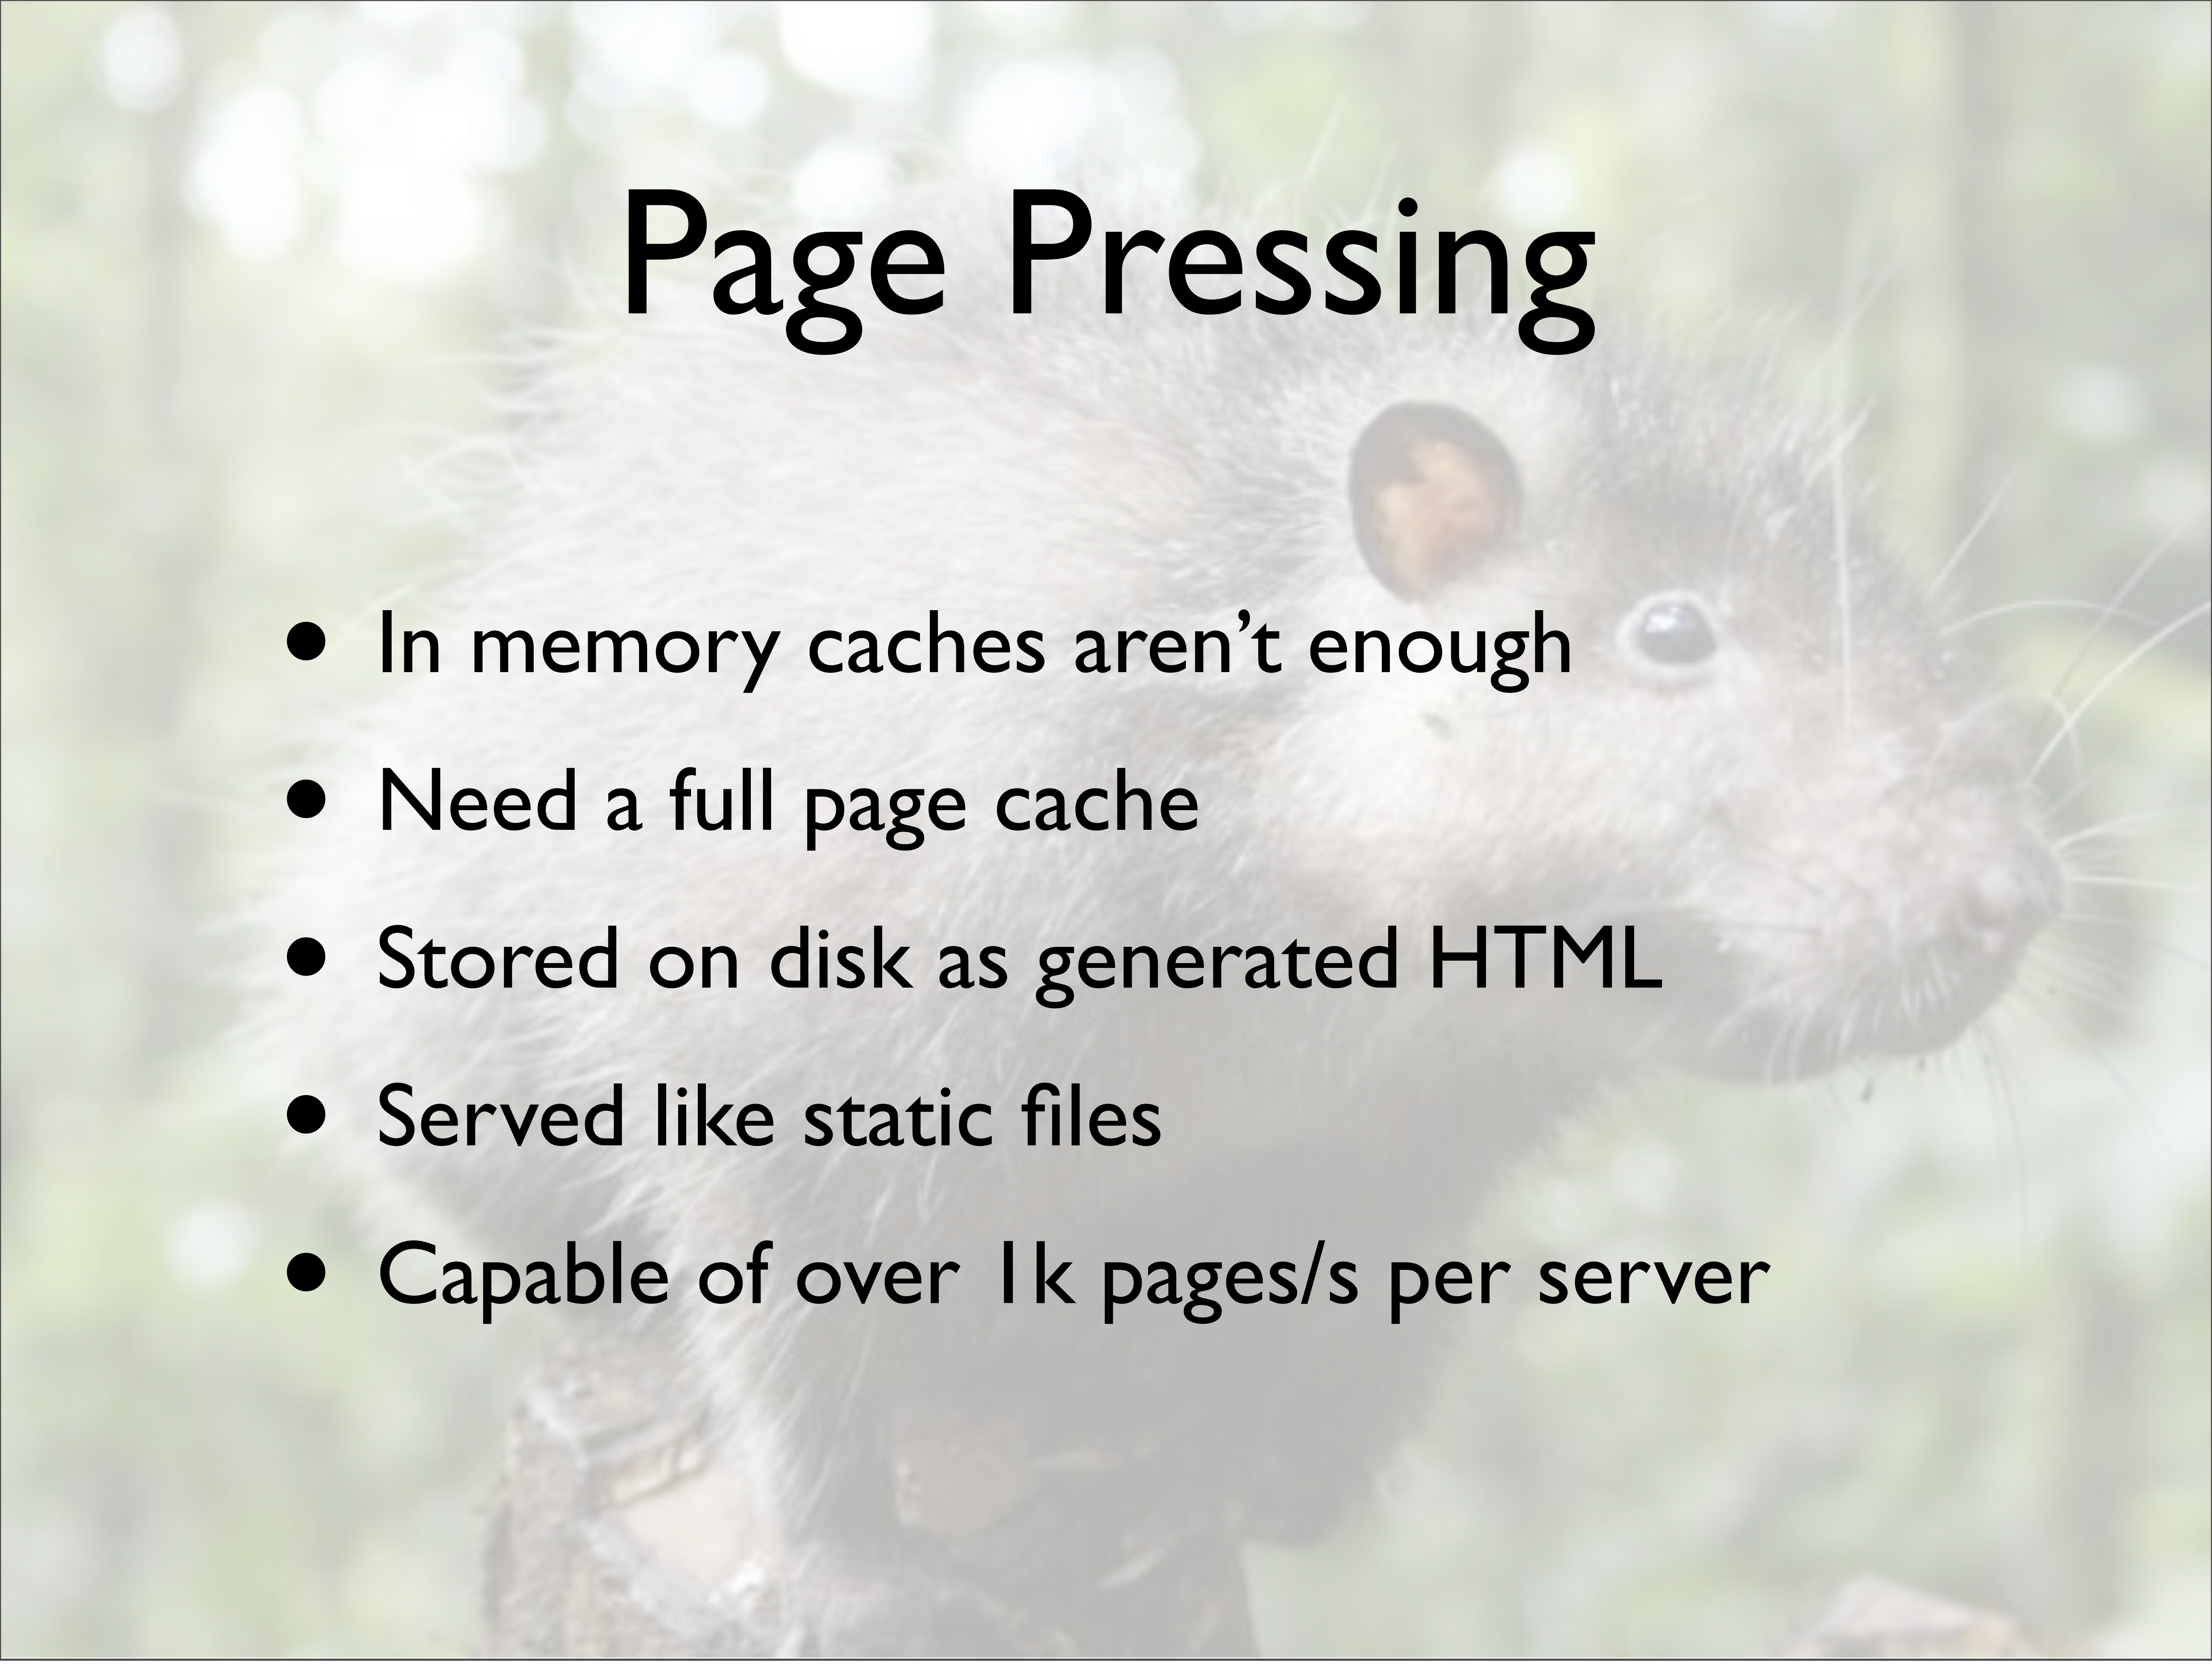

- In memory caches aren’t enough
- Need a full page cache
- Stored on disk as generated HTML
- Served like static files
- Capable of over 1k pages/s per server

Really cache everything

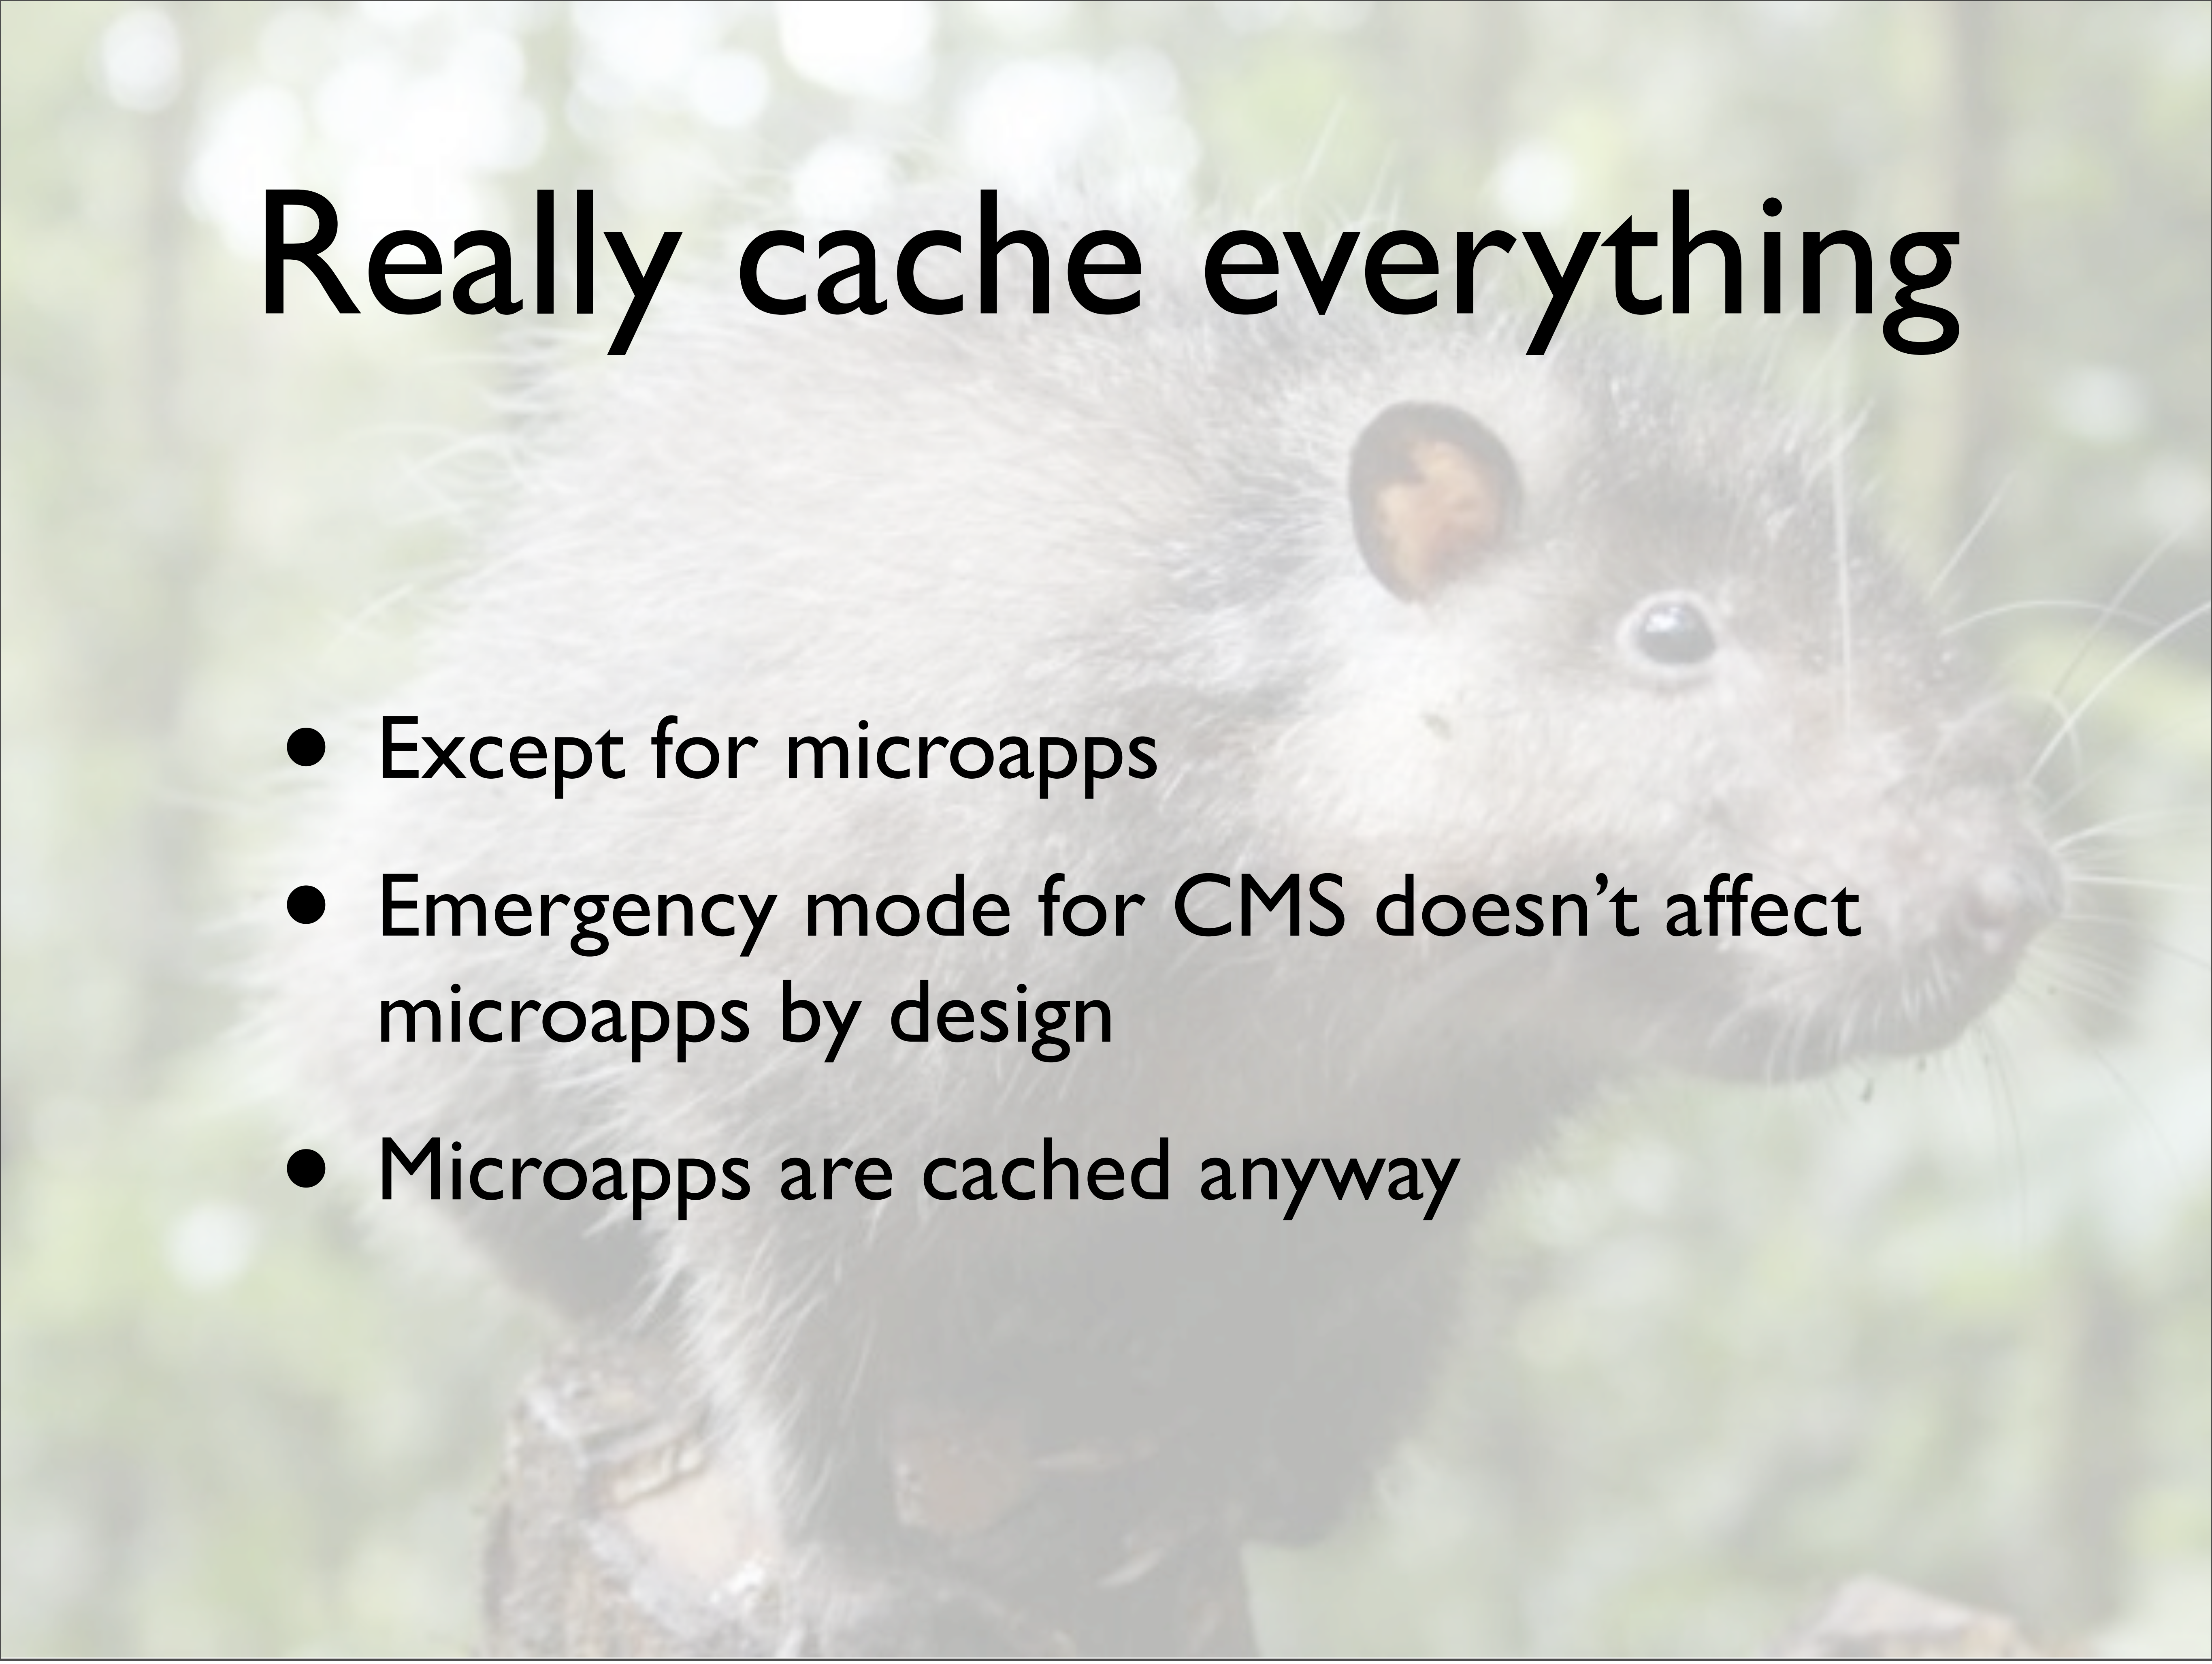

- Except for microapps
- Emergency mode for CMS doesn’t affect microapps by design
- Microapps are cached anyway

Gotta cache 'em all

- 1.6 million pieces of content
- http://www.guardian.co.uk/uk/budget
- http://www.guardian.co.uk/travel/france+travel/skiing
- http://www.guardian.co.uk/theguardian/2012/mar/02
- http://www.guardian.co.uk/technology/apple?page=2

Cache what’s important

- Content - when modified
- Navigation - Every 2 weeks
- Automatic but important - Every 2 weeks
- Automatic (eg tag combiners) - Never
- Can force a page press

Monitoring

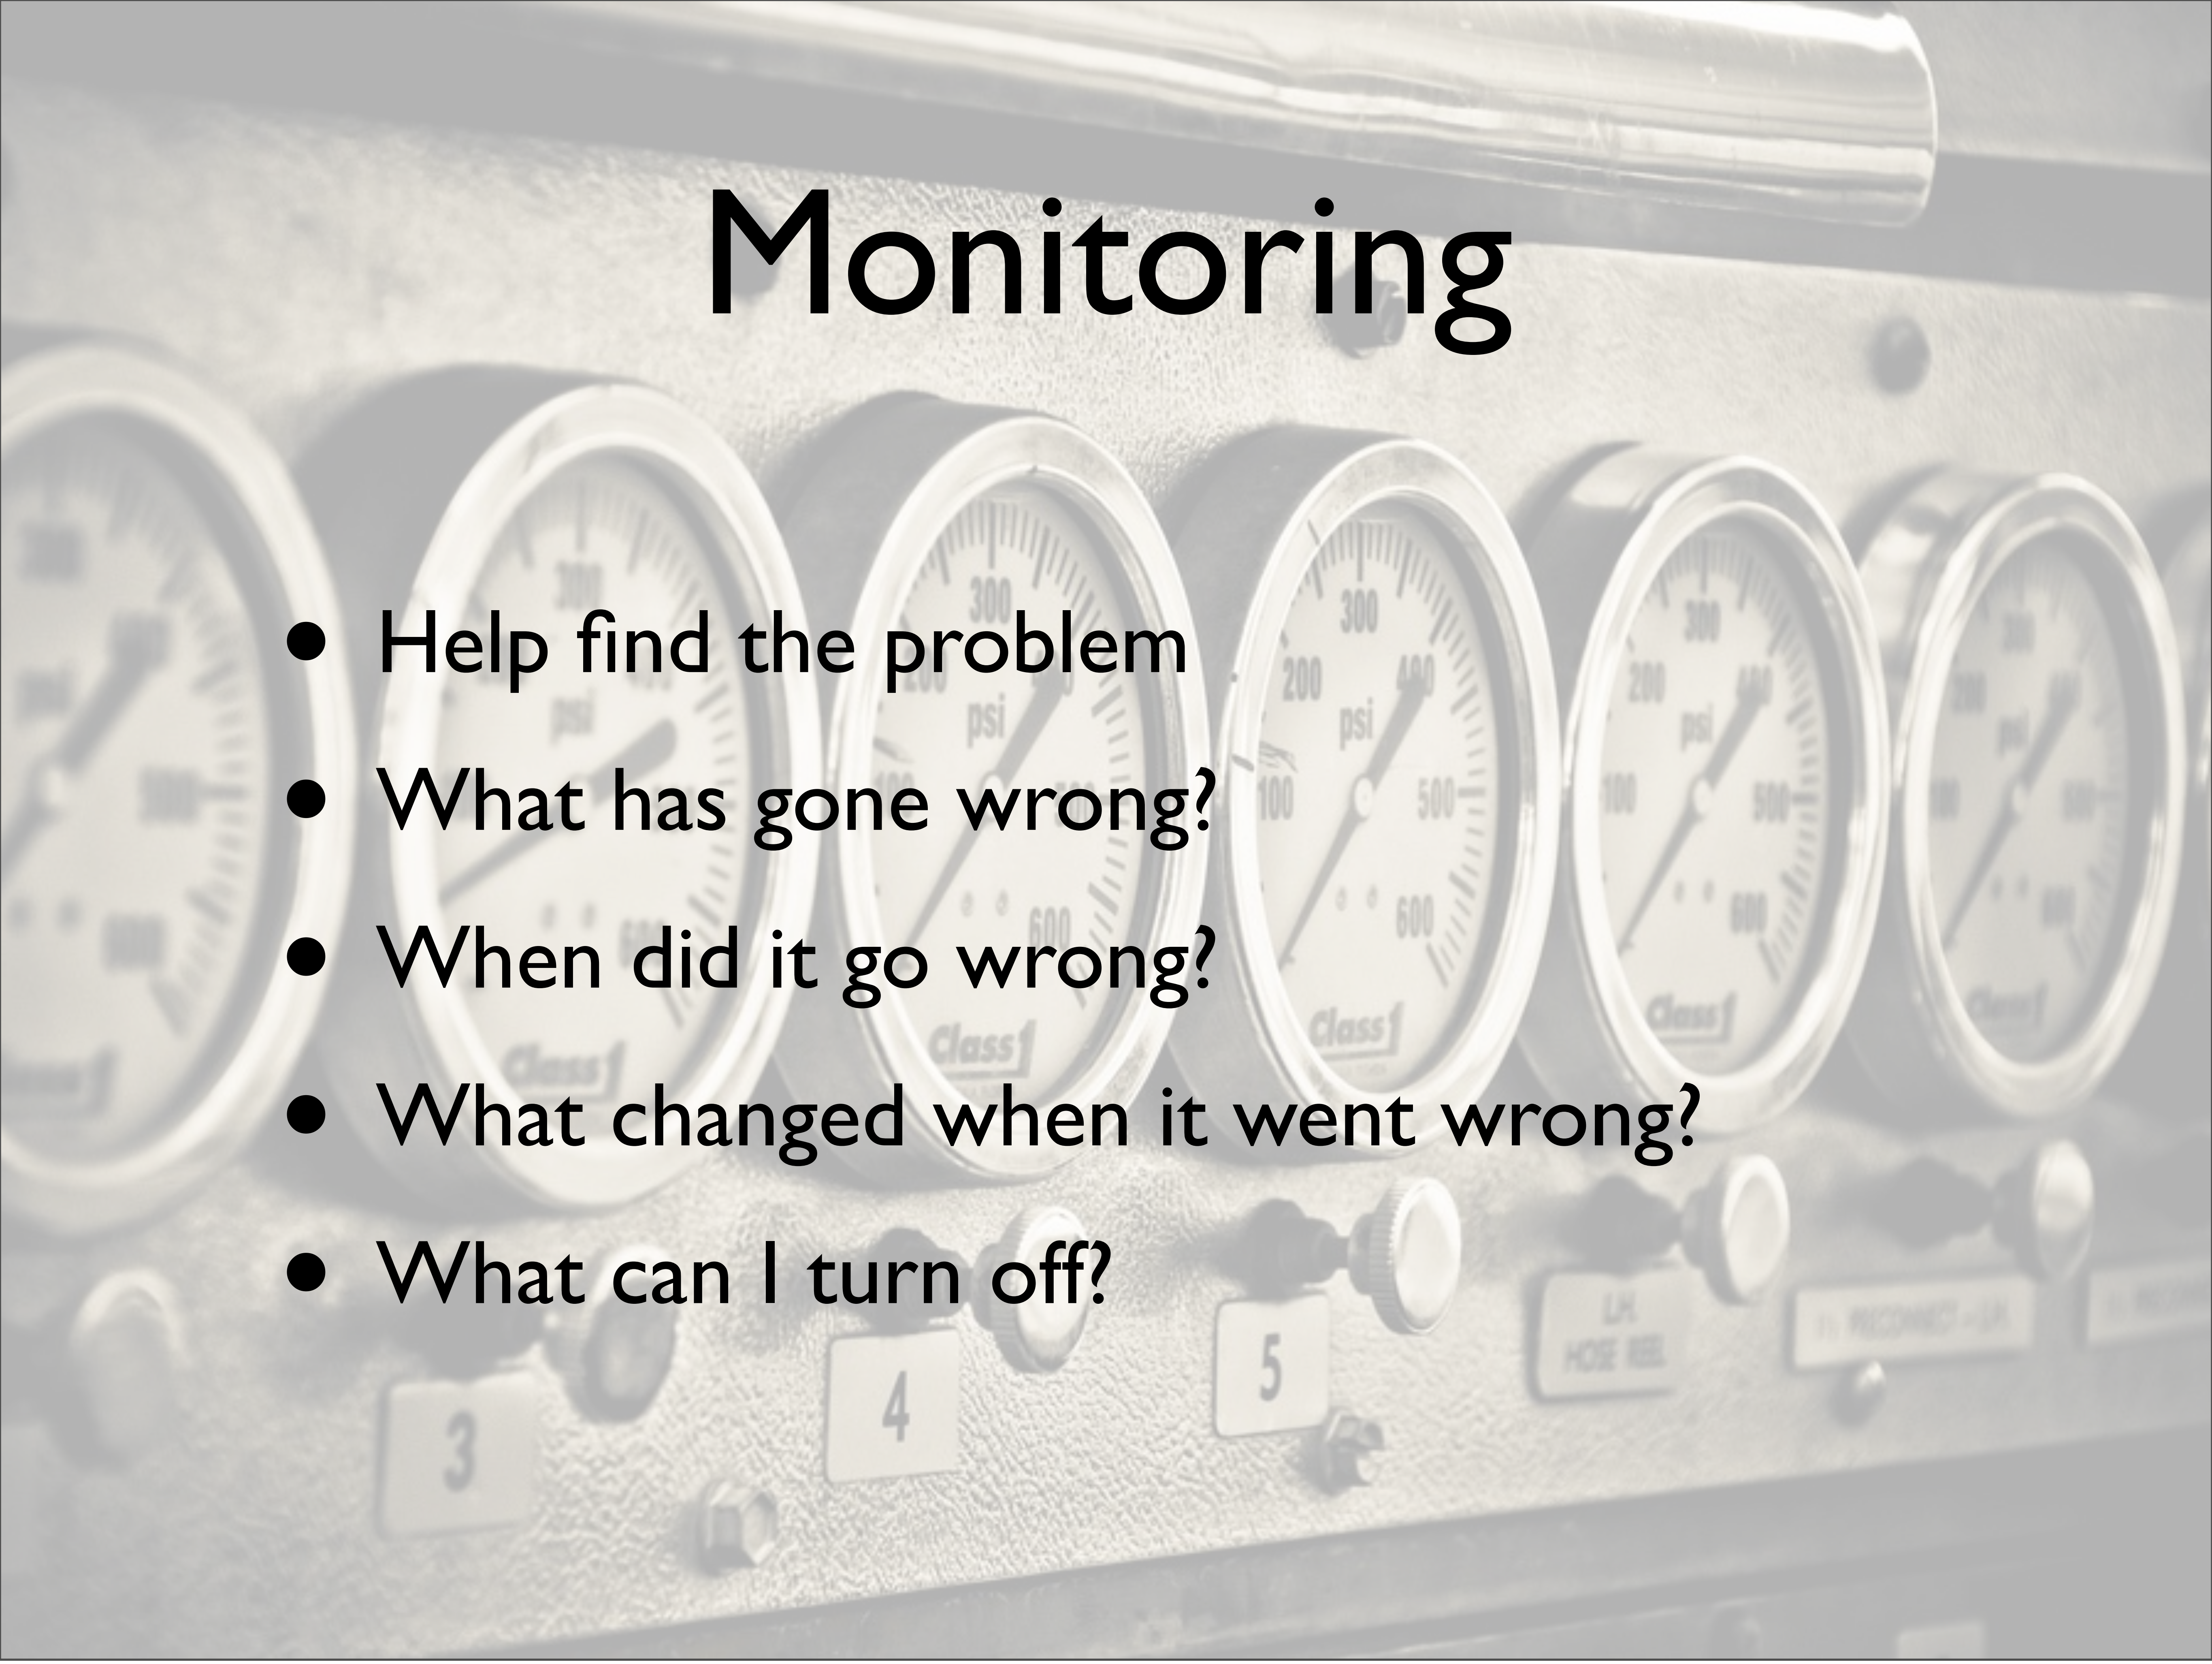

- Help find the problem
- What has gone wrong?
- When did it go wrong?
- What changed when it went wrong?
- What can I turn off?

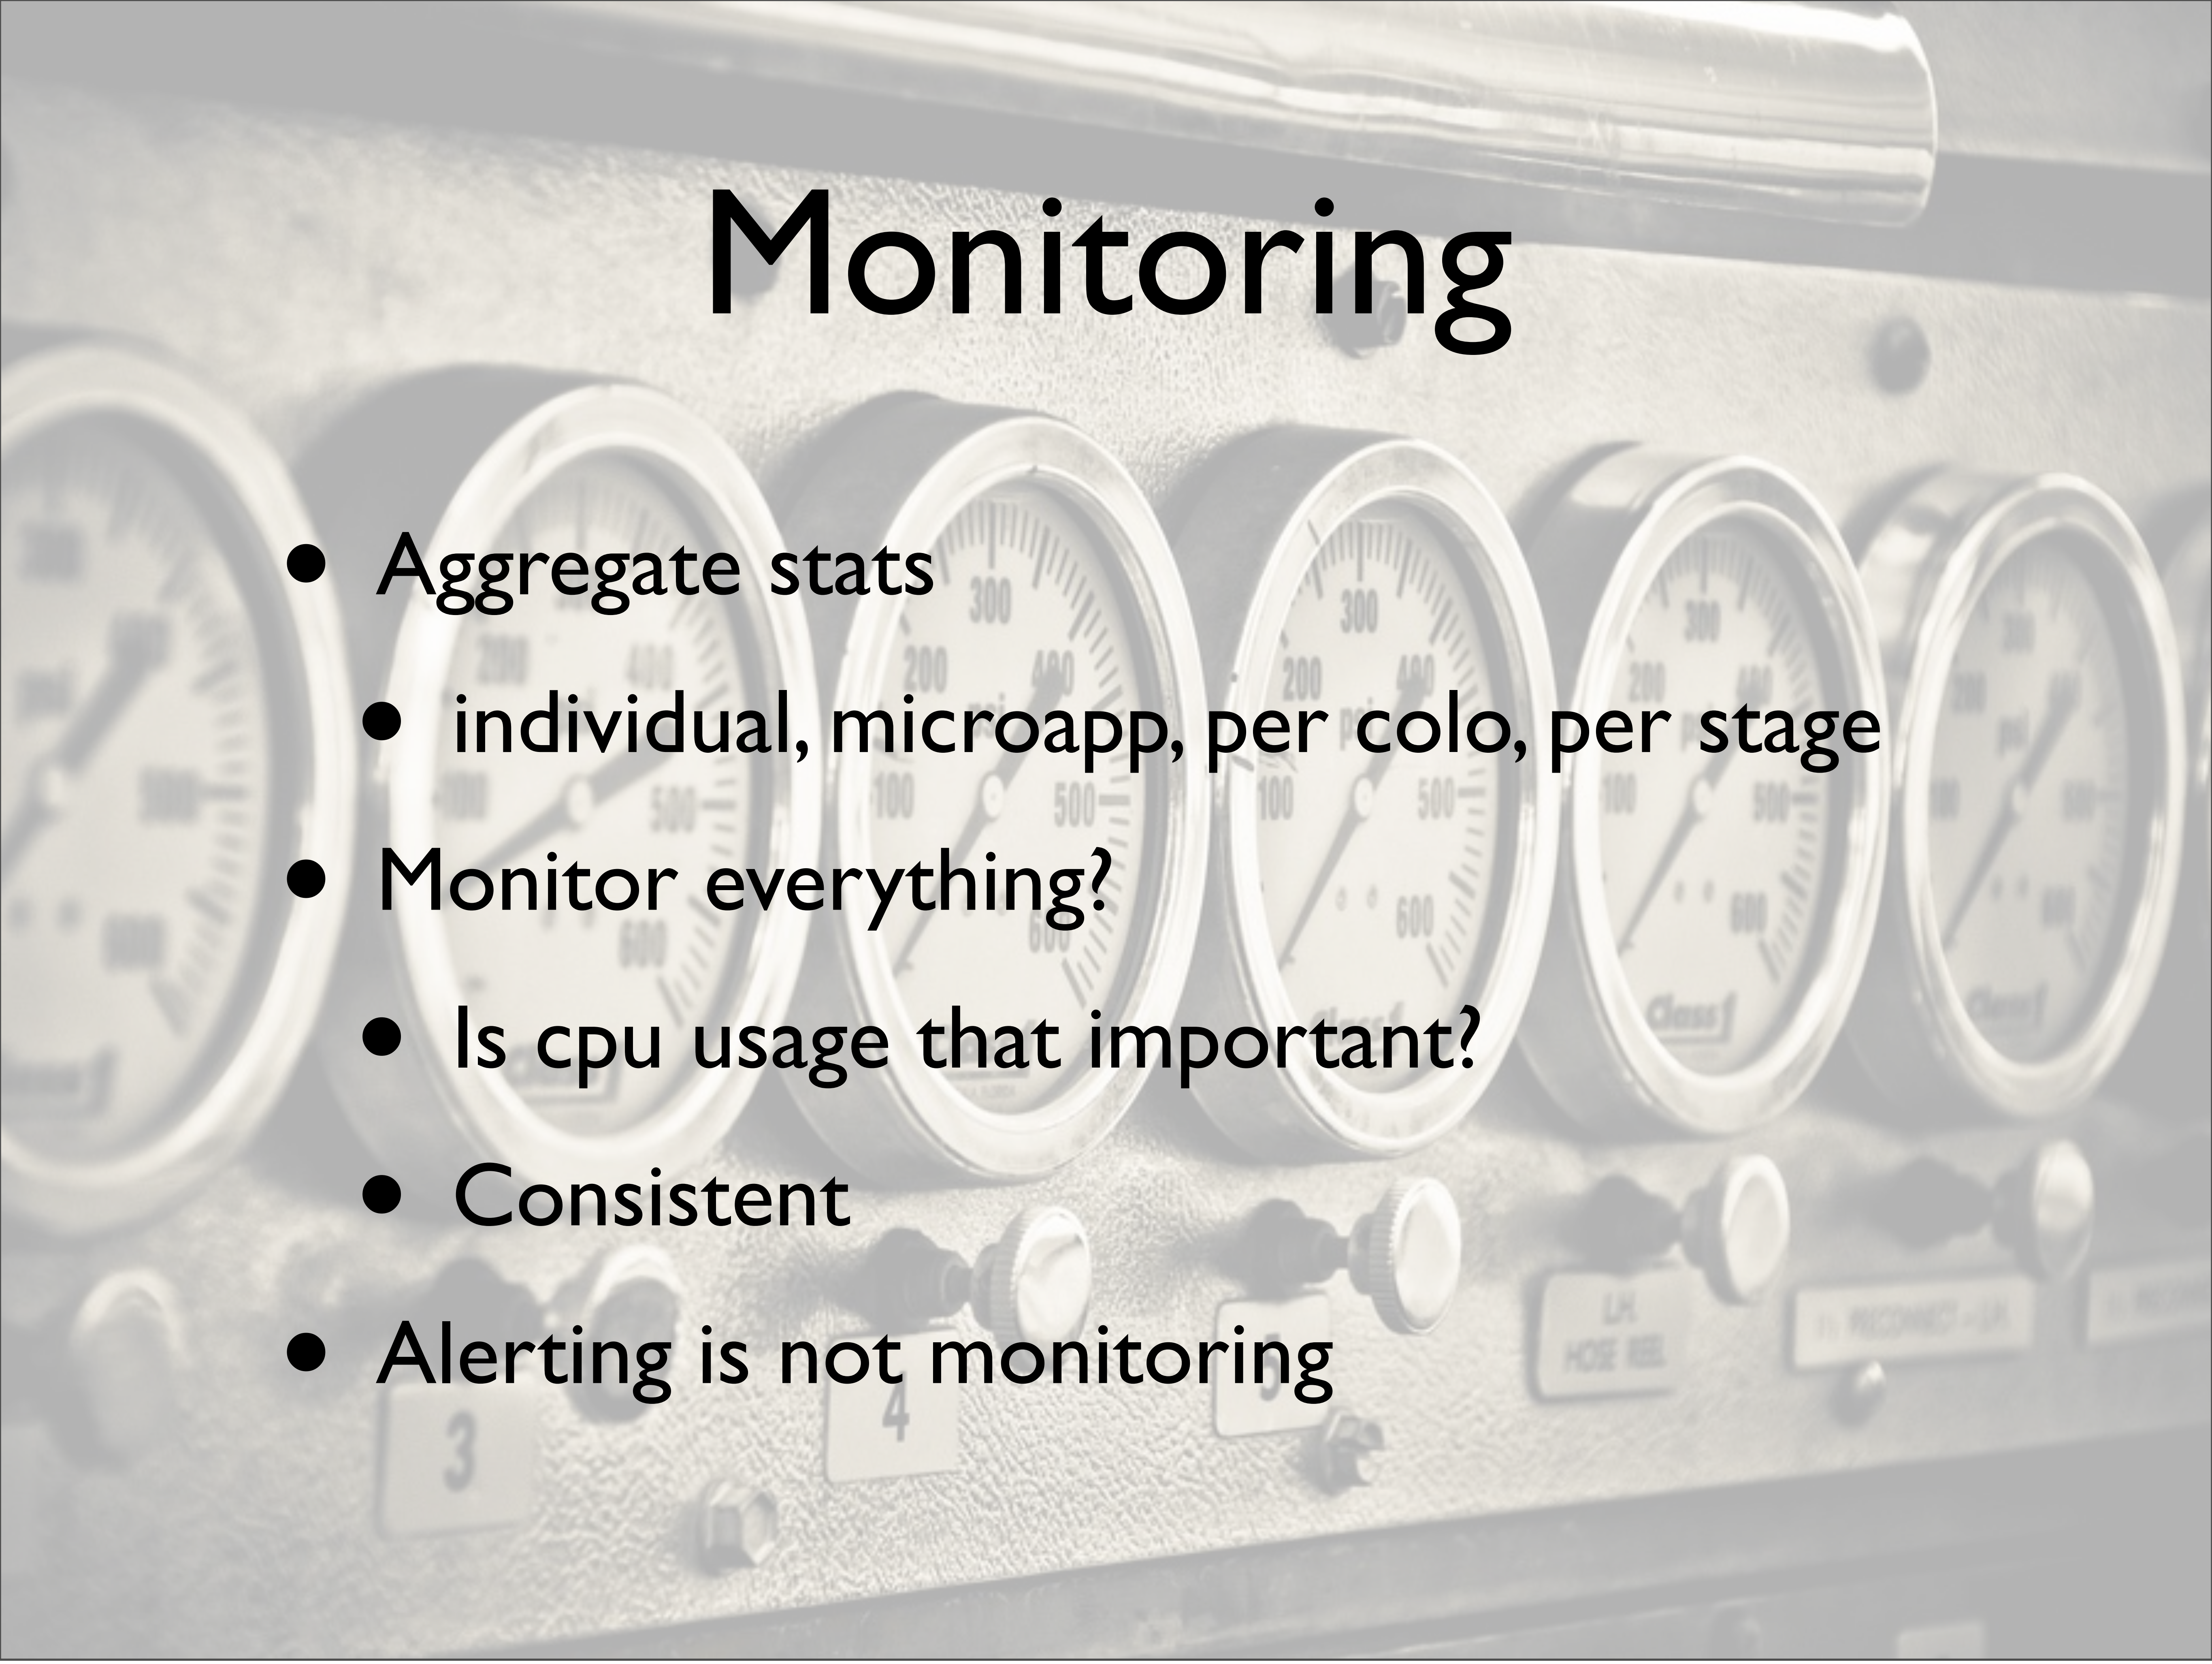

Monitoring

- Aggregate stats
  - individual, microapp, per colo, per stage
- Monitor everything?
  - Is cpu usage that important?
  - Consistent
- Alerting is not monitoring

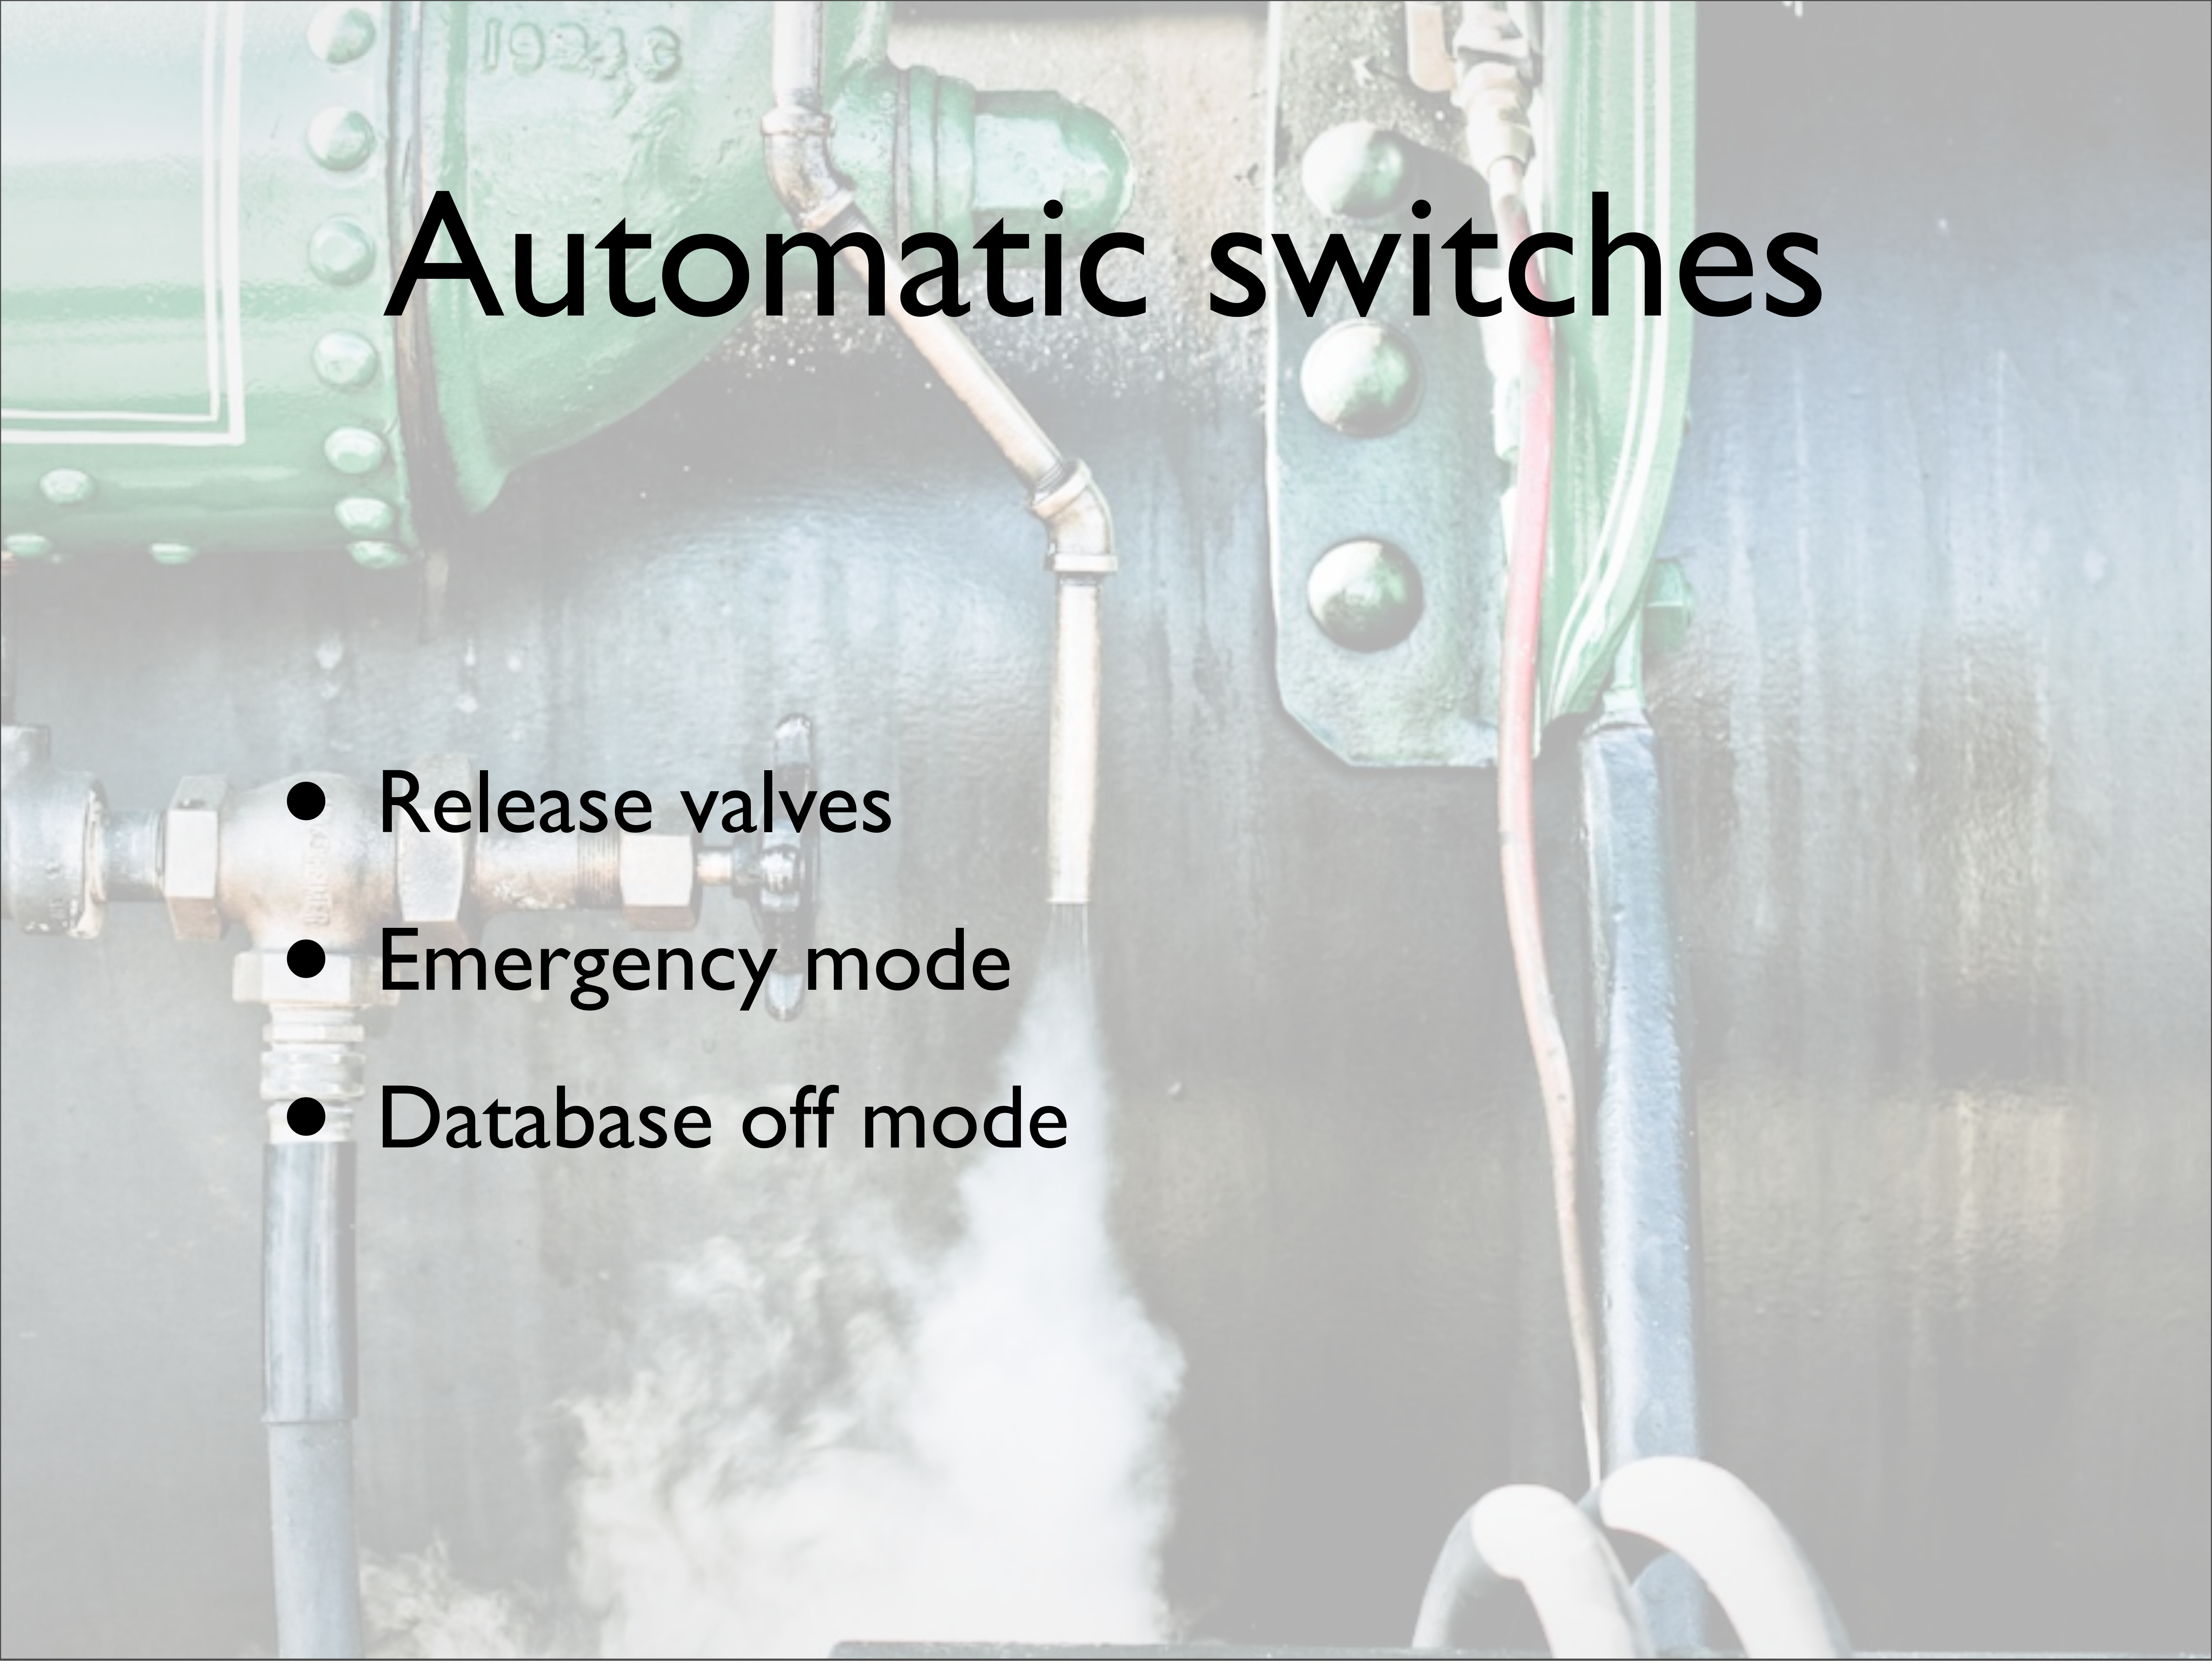

Automatic switches

- Release valves
- Emergency mode
- Database off mode

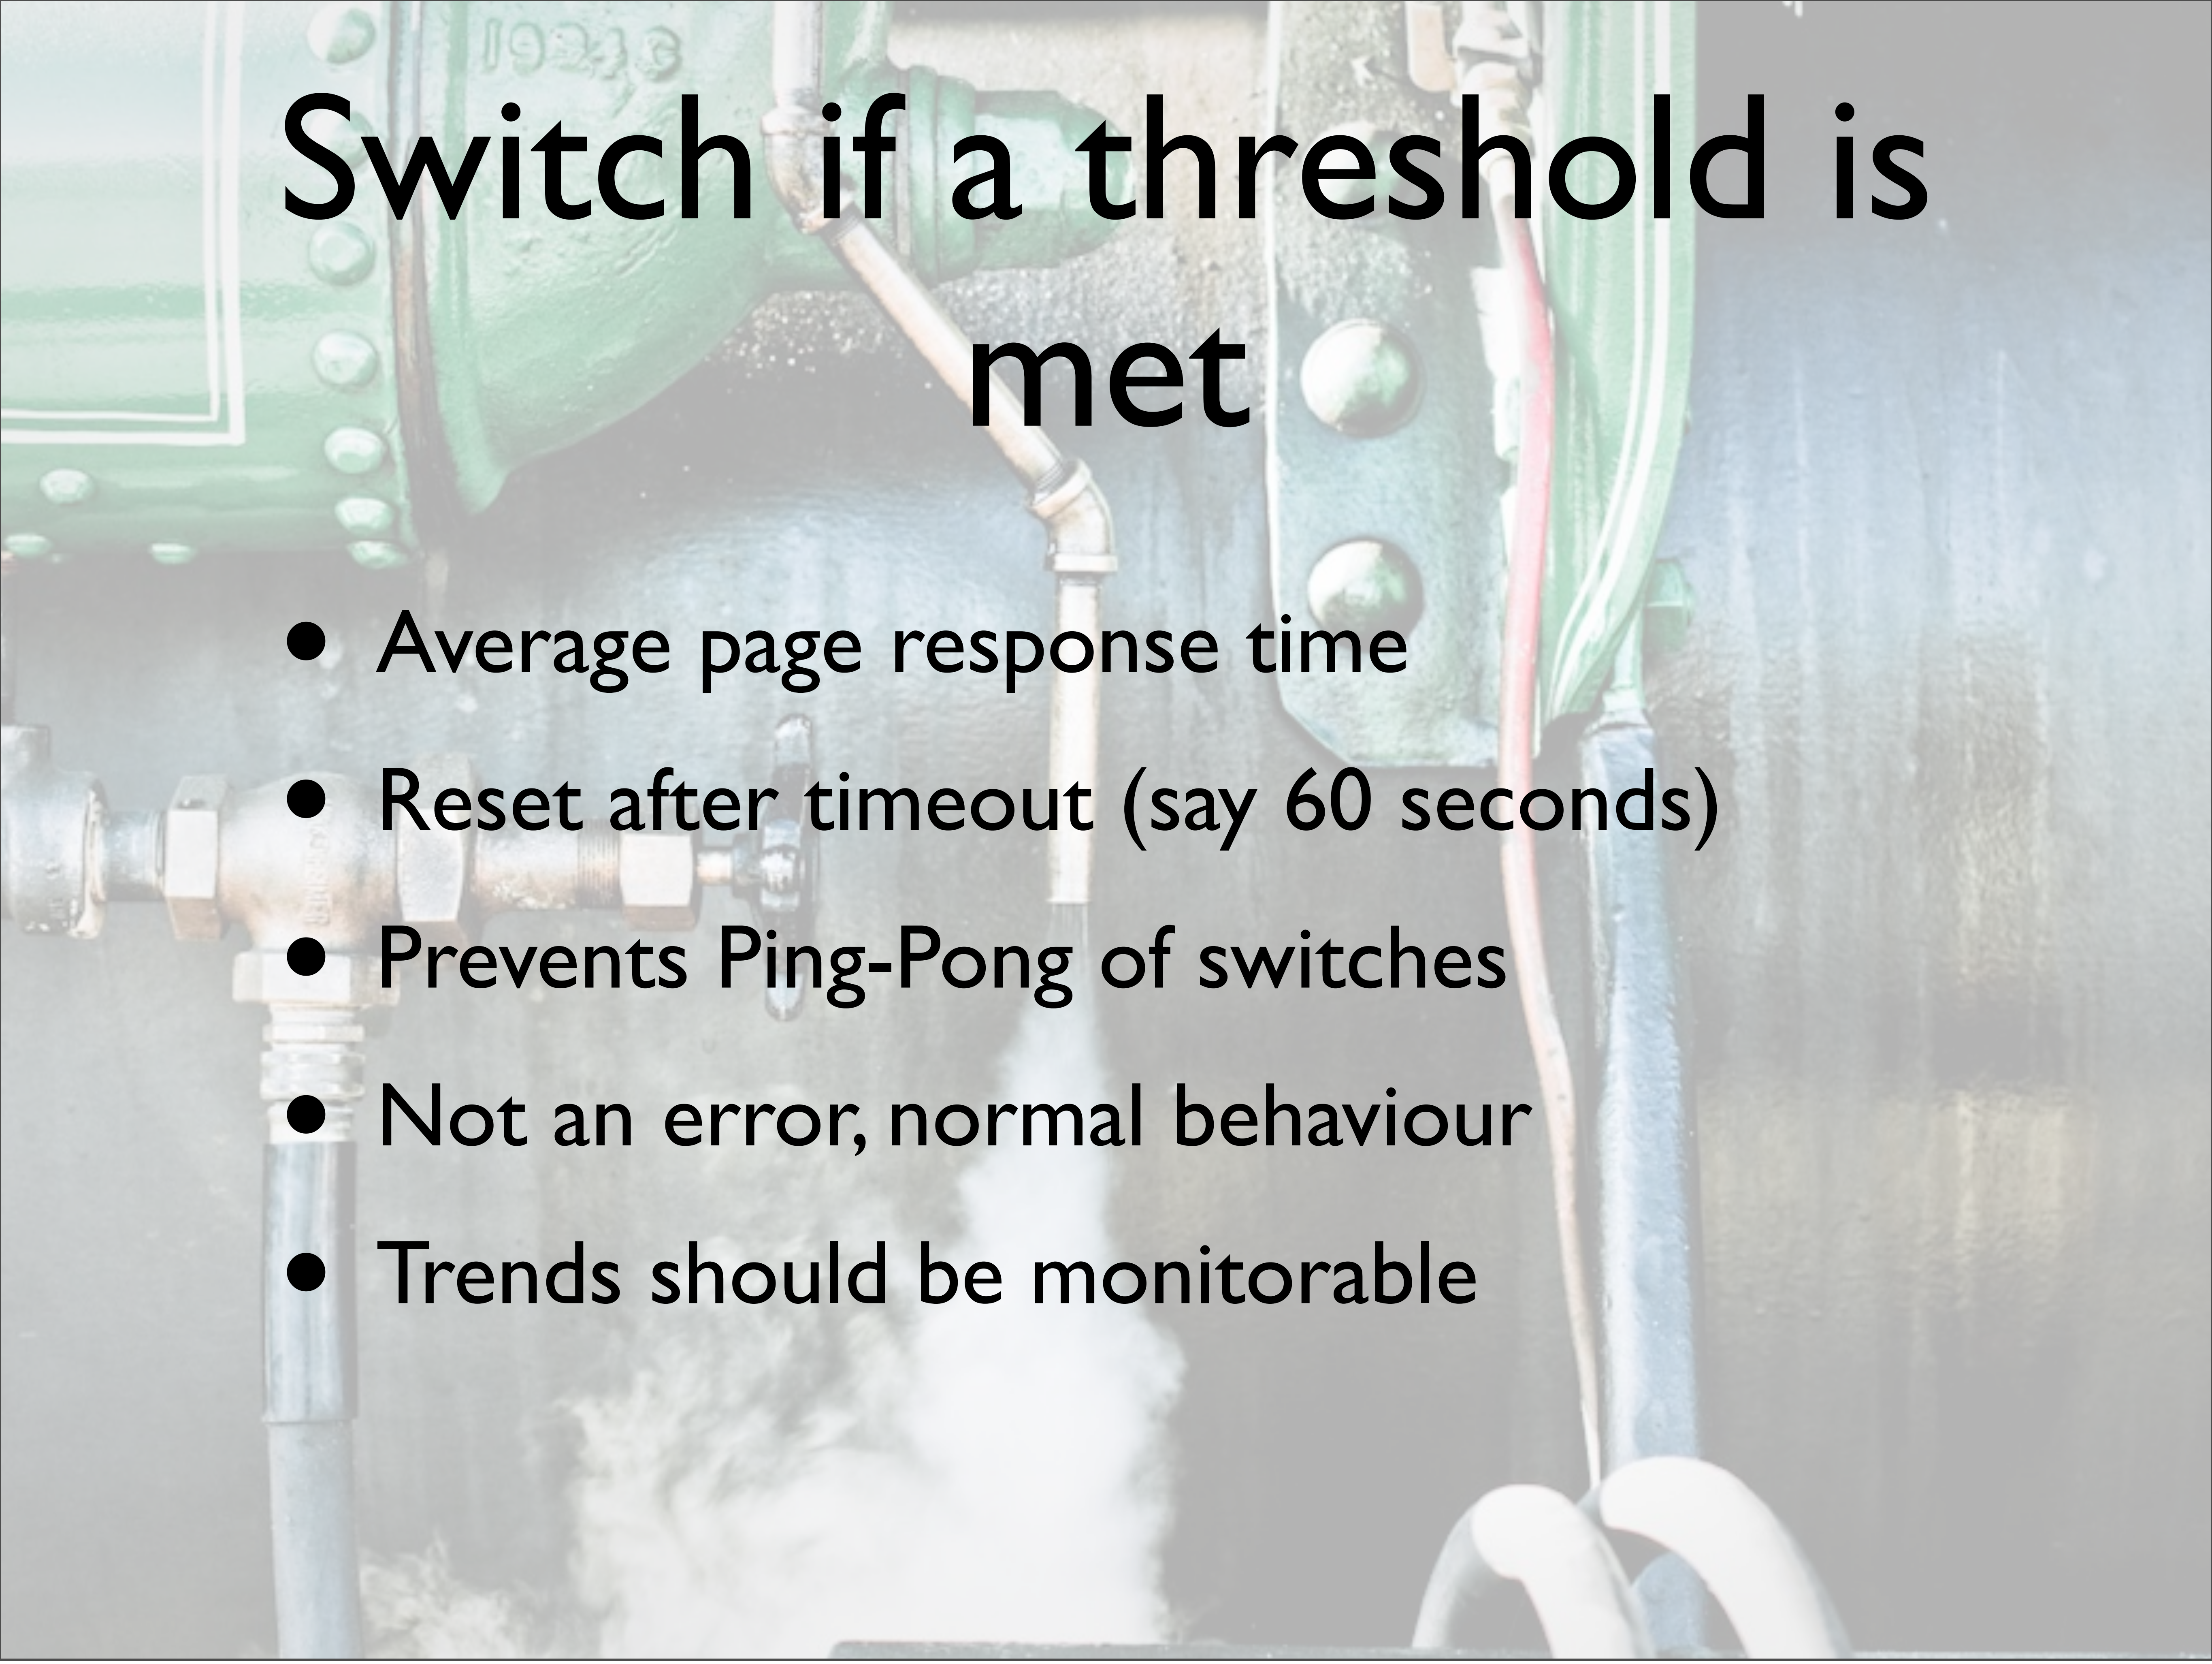

Switch if a threshold is met

- Average page response time
- Reset after timeout (say 60 seconds)
- Prevents Ping-Pong of switches
- Not an error, normal behaviour
- Trends should be monitorable

Diagnosing failure

Why do I care?

- Your architecture must be easy to diagnose
- These skills aren’t common enough
    - Basic unix skills (sed, grep, cut, sort)
    - Log analysis
- Take these into account when you design

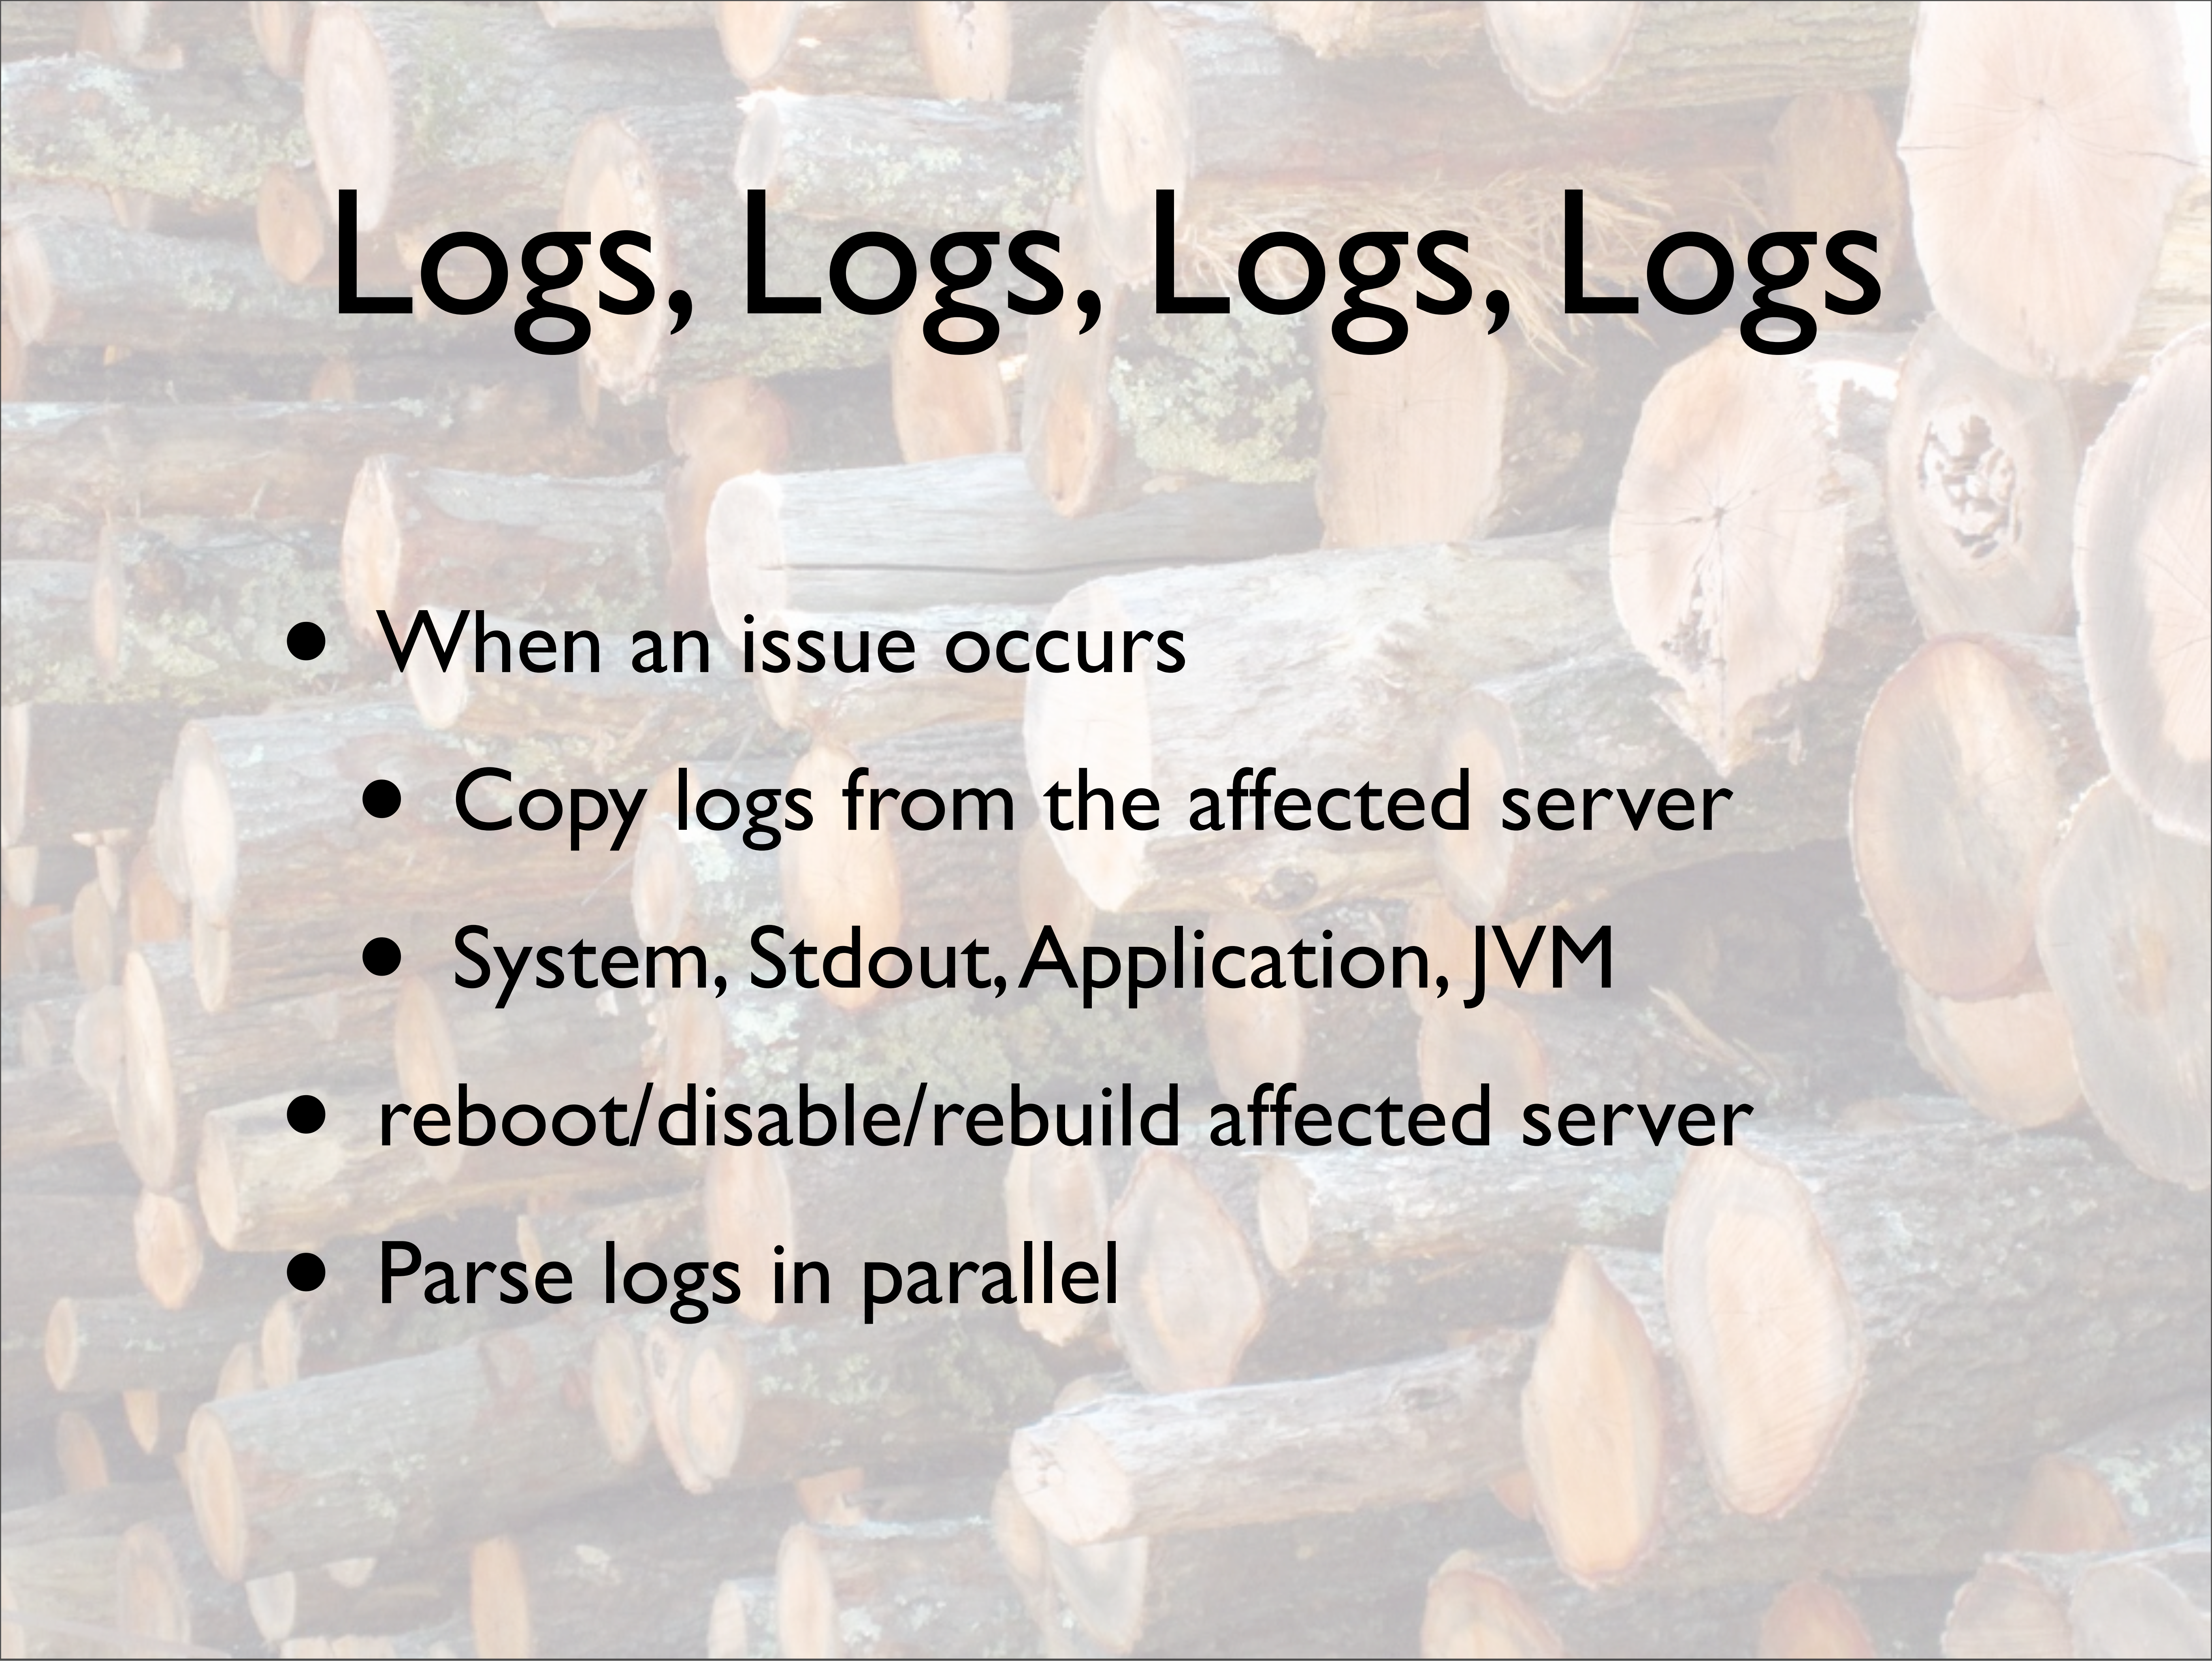

Logs, Logs, Logs, Logs

- When an issue occurs
  - Copy logs from the affected server
  - System, Stdout, Application, JVM
- reboot/disable/rebuild affected server
- Parse logs in parallel

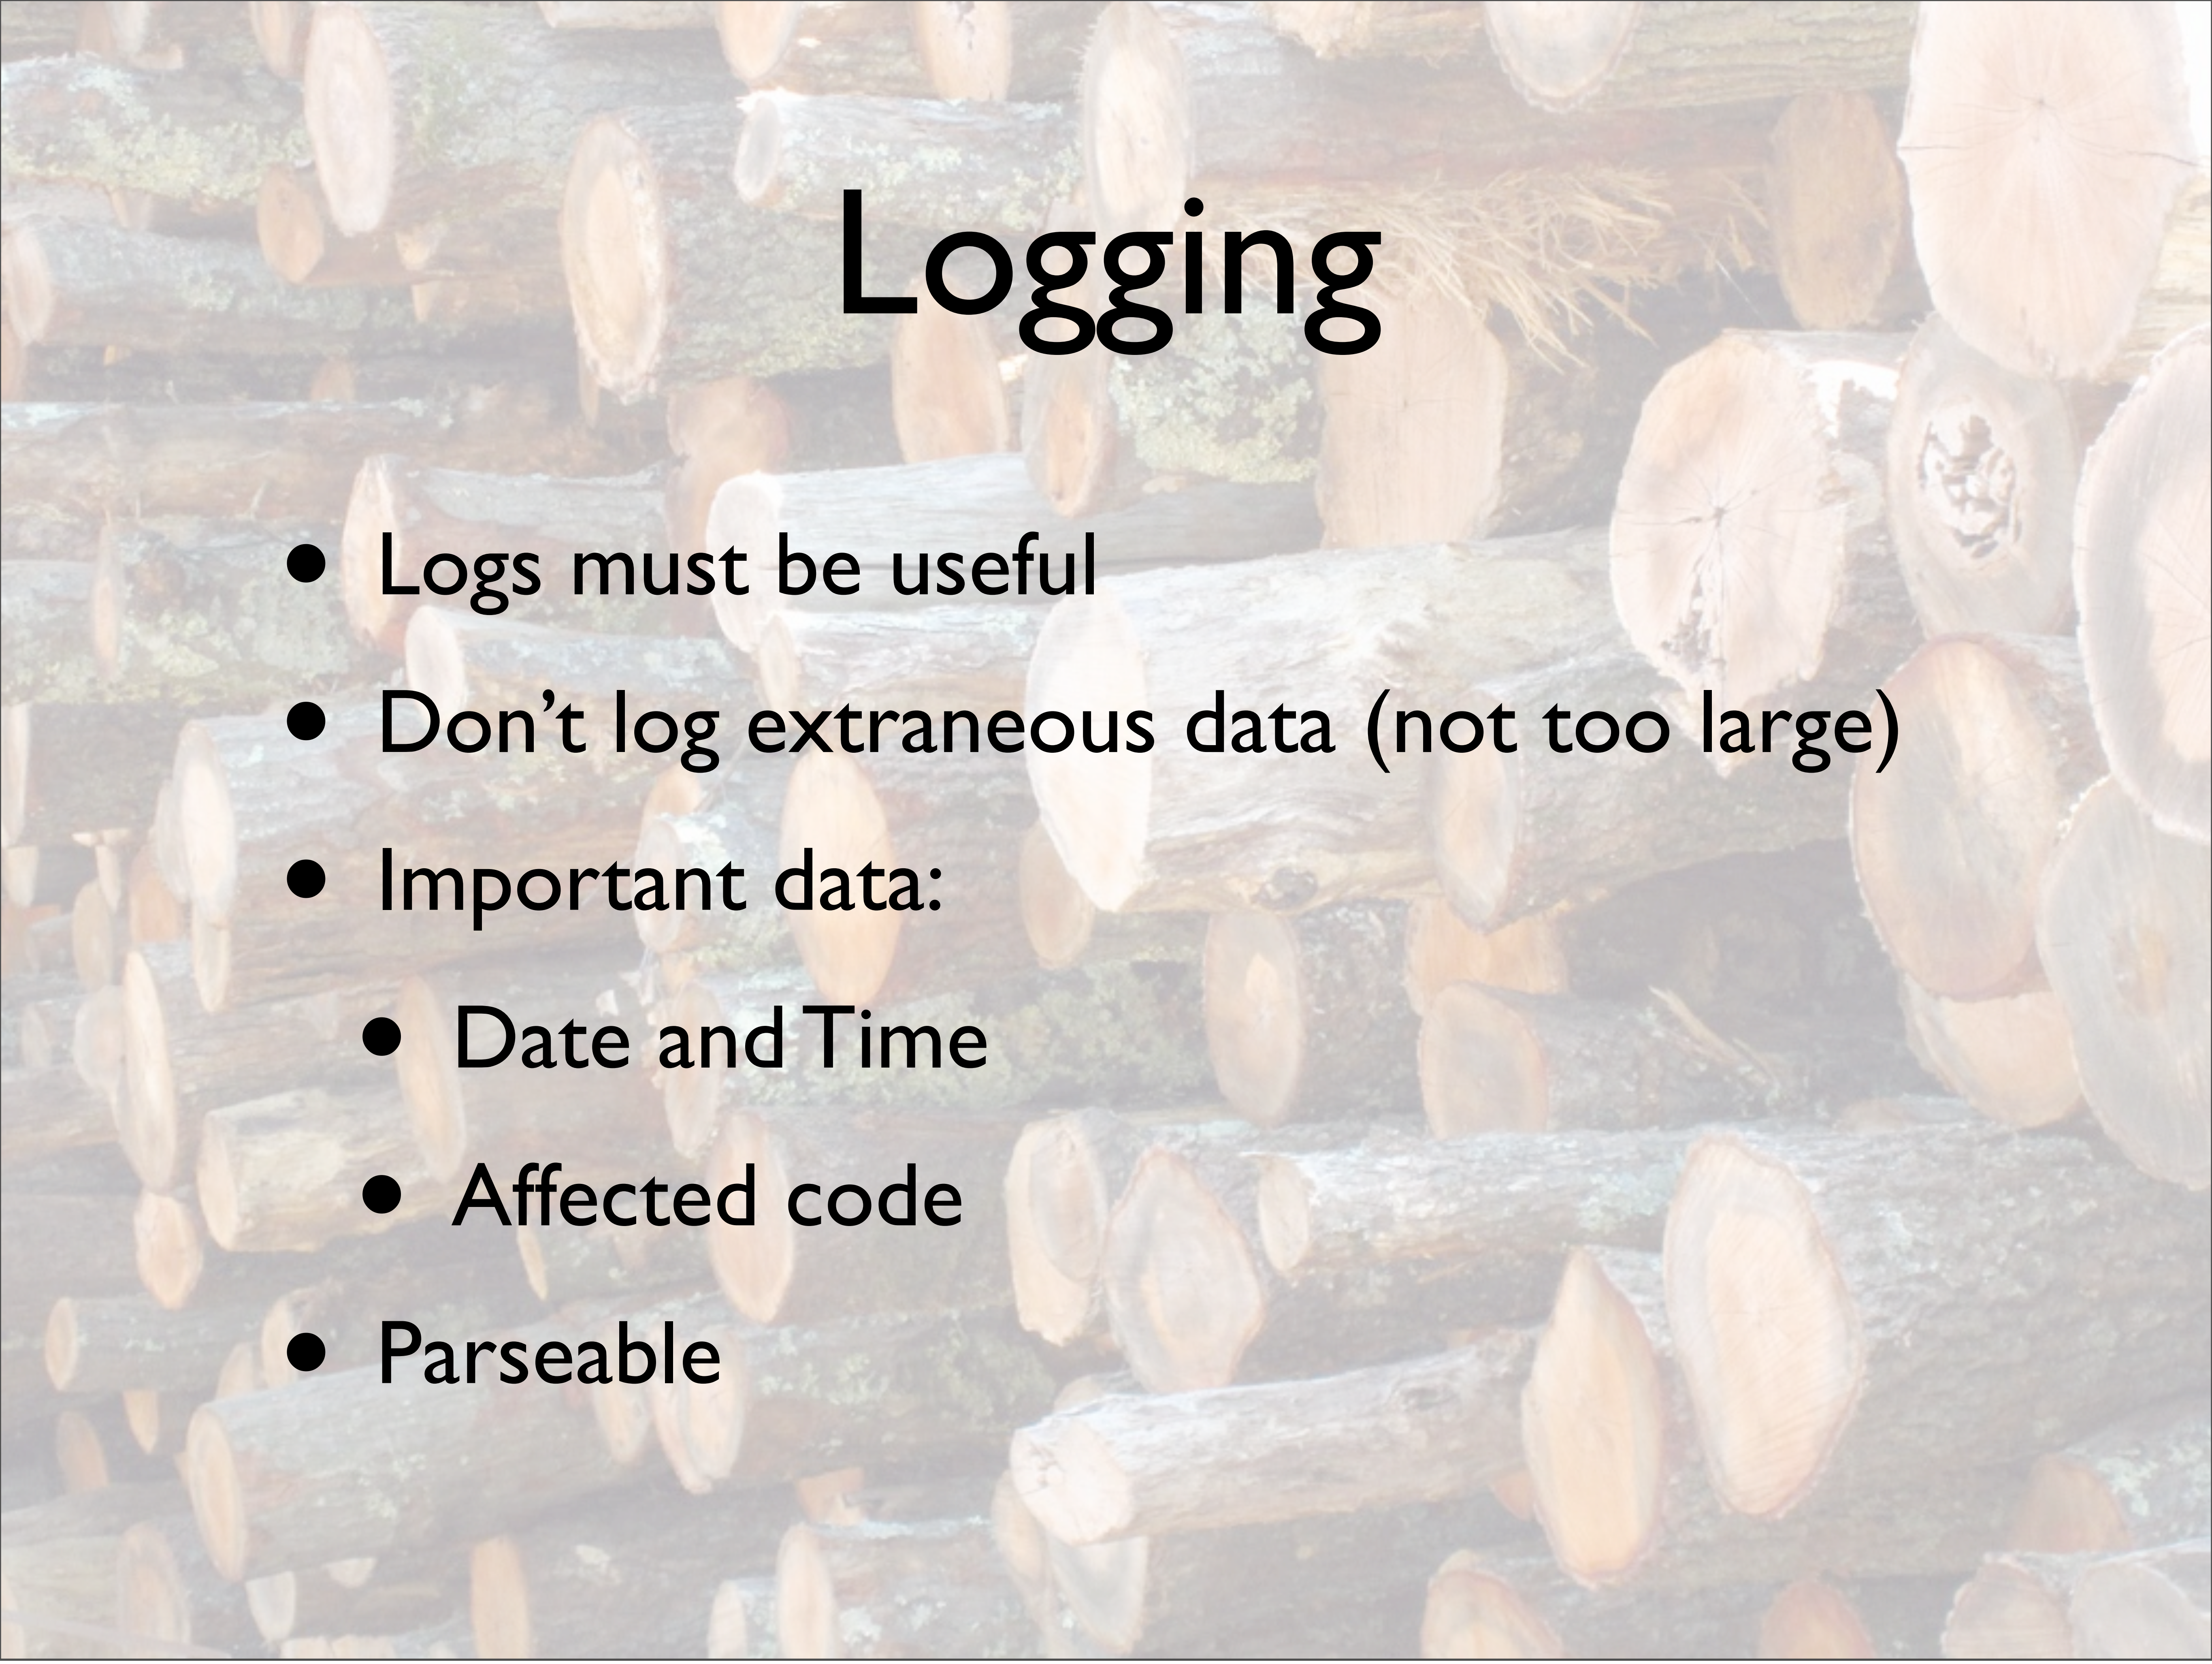

Logging

- Logs must be useful
- Don’t log extraneous data (not too large)
- Important data:
  - Date and Time
  - Affected code
- Parseable

Loggable Events

- Request Logging (after including time)
    - External service requests (after including time)
- Interesting events
- Exceptions
- Database calls?

Loggable Events

```
2012-03-06 14:58:19,351 [resin-tcp-connection-*:8080-19]
INFO  com.gu.management.logging.RequestLoggingFilter -
Request for /pages/Guardian/artanddesign/artblog/2008/jan/
31/catchofthedaysecondlifes completed in 231 ms
```

Loggable Events

```
2012-03-06 14:58:19,351 [resin-tcp-connection-*:8080-19] INFO  com.gu.management.logging.RequestLoggingFilter - Request for /pages/Guardian/artanddesign/artblog/2008/jan/31/catchofthedaysecondlifes completed in 231 ms
```

Date and Time

Loggable Events

```
2012-03-06 14:58:19,351 [resin-tcp-connection-*:8080-19]
INFO  com.gu.management.logging.RequestLoggingFilter -
Request for /pages/Guardian/artanddesign/artblog/2008/jan/
31/catchofthedaysecondlifes completed in 231 ms
```

Thread name

Date and Time

Loggable Events

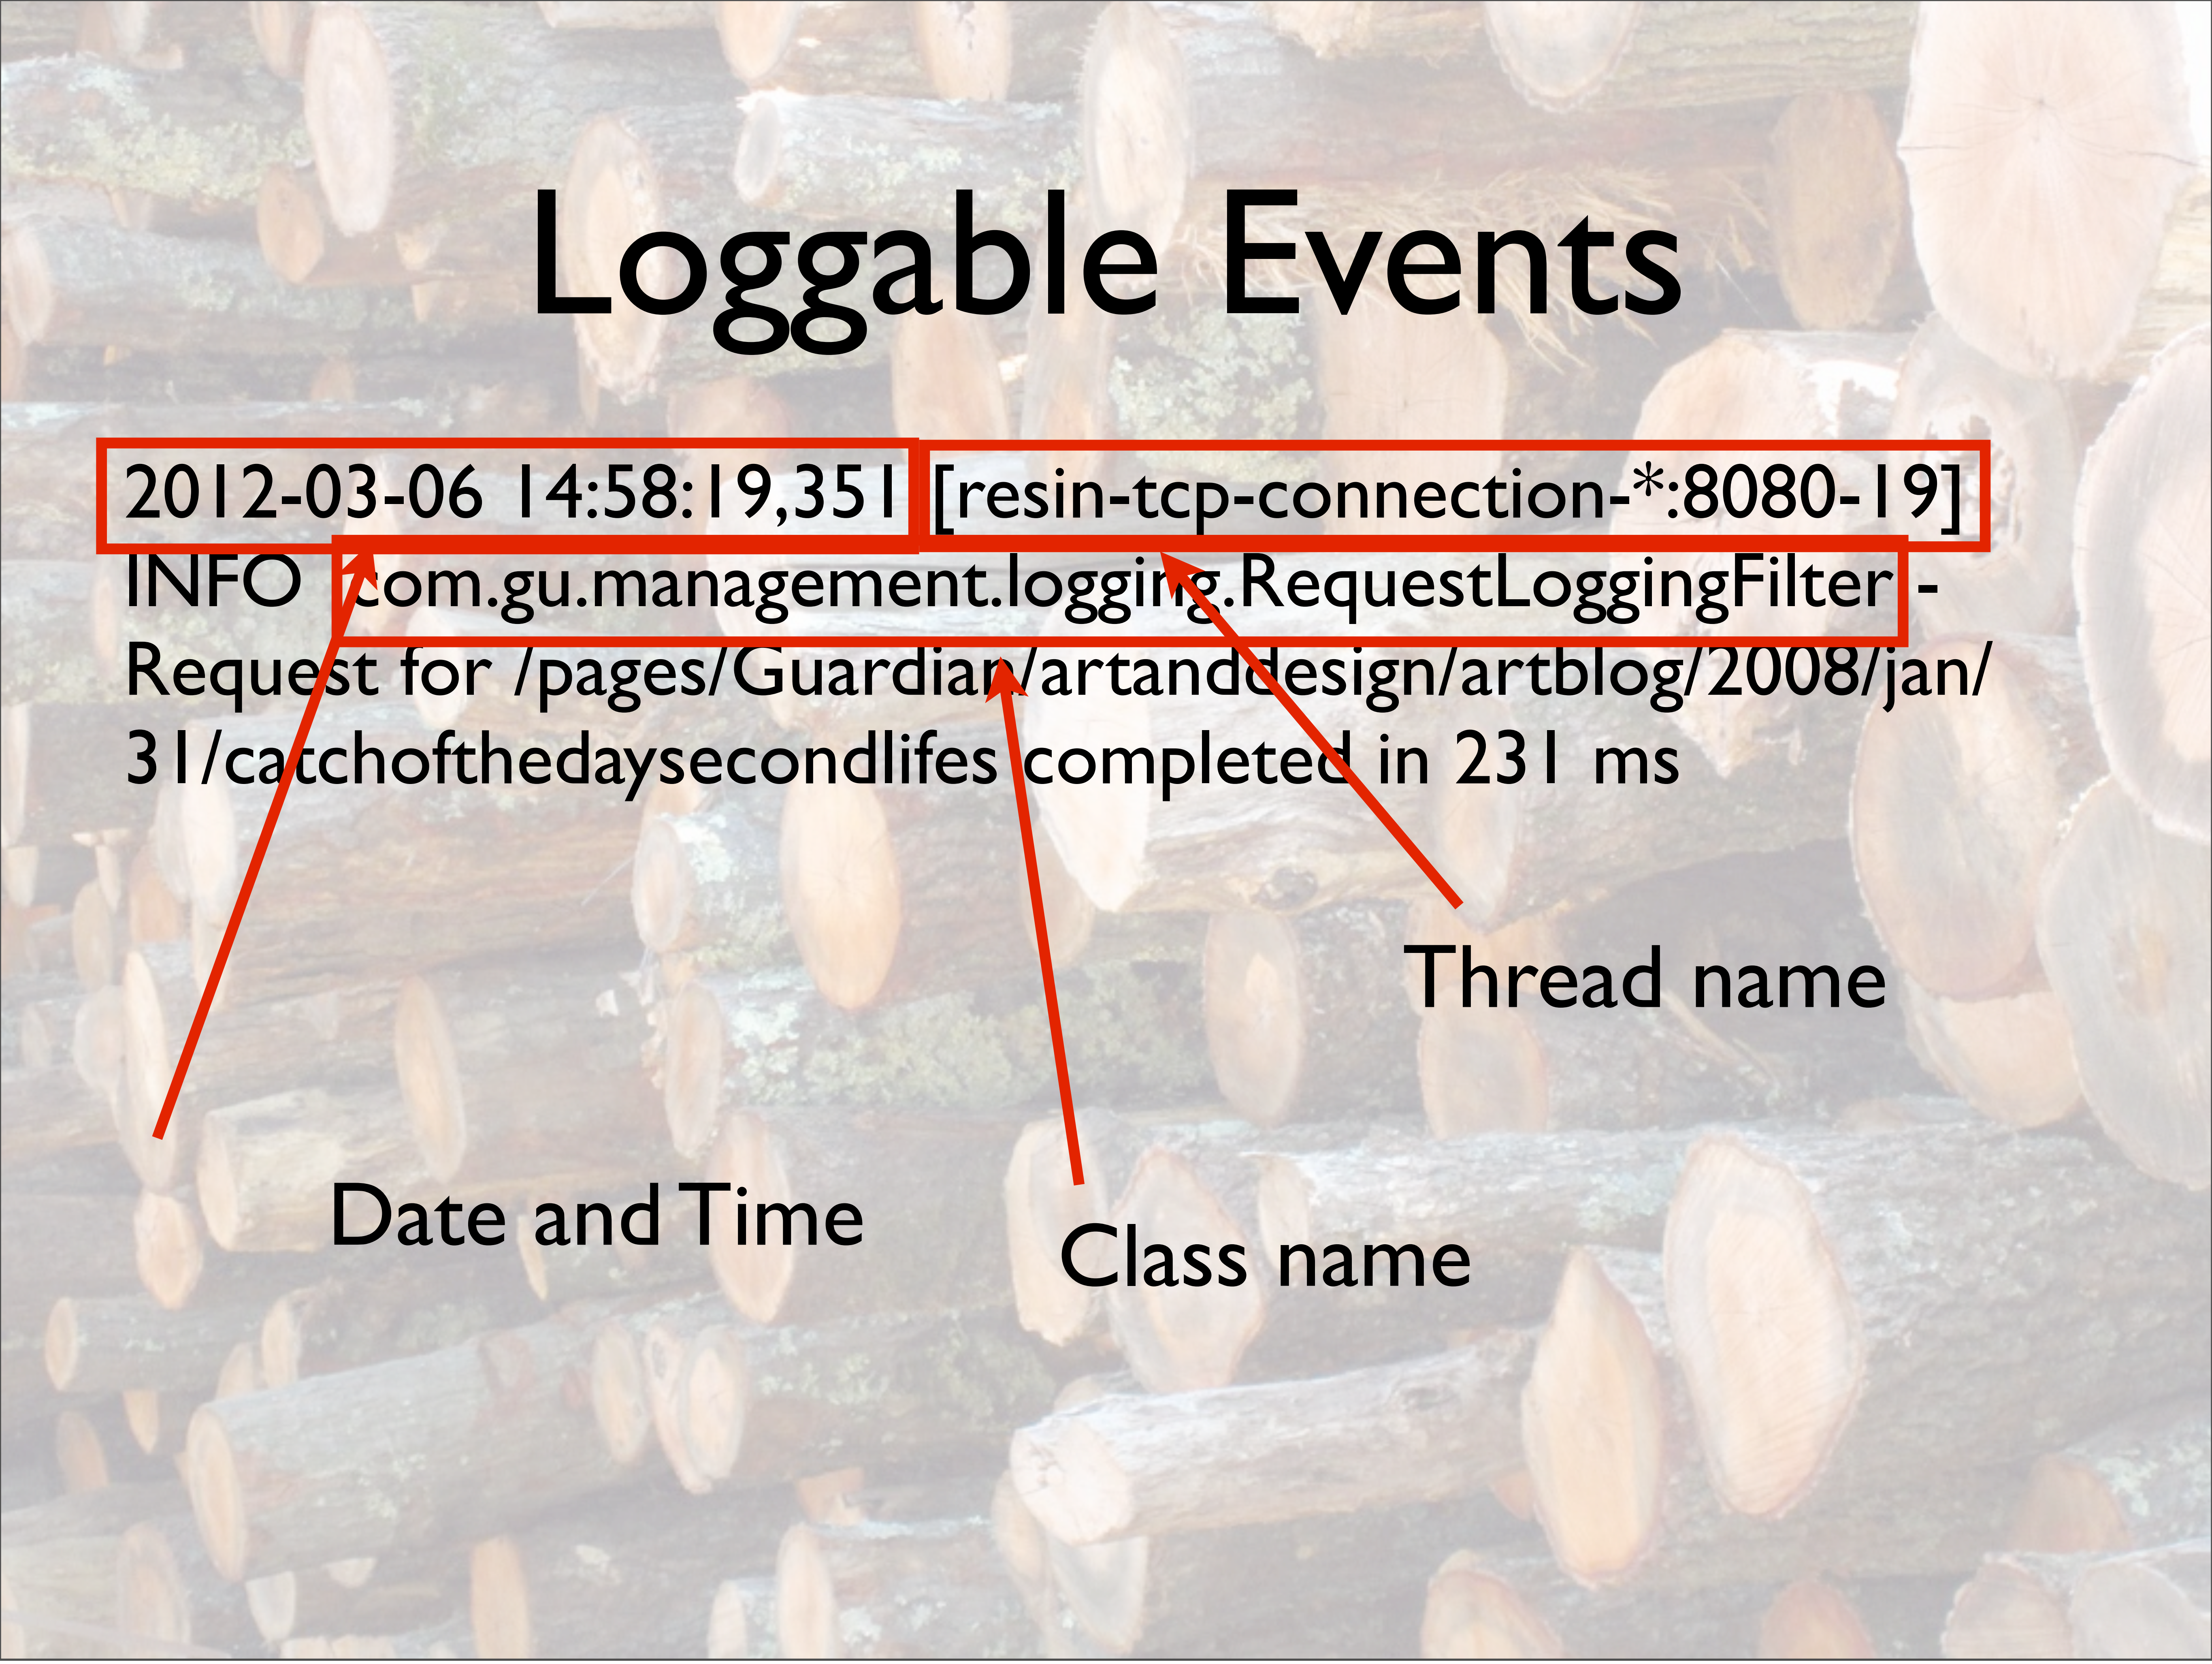

```
2012-03-06 14:58:19,351 [resin-tcp-connection-*:8080-19]
INFO  com.gu.management.logging.RequestLoggingFilter -
Request for /pages/Guardian/artanddesign/artblog/2008/jan/
31/catchofthedaysecondlifes completed in 231 ms
```

Thread name

Date and Time

Class name

Log analysis is your friend

- Simple tools for a simple life
  - Grep
  - Cut
  - Uniq
  - Sort
  - Sed and Awk

```
zgrep "RequestLoggingFilter - Request for.*completed in "
$LOGFILE
| grep -v " /management/"
| cut -d" " -f1,2,3,10,13
> $COMPLETED_REQUESTS_FILE
```

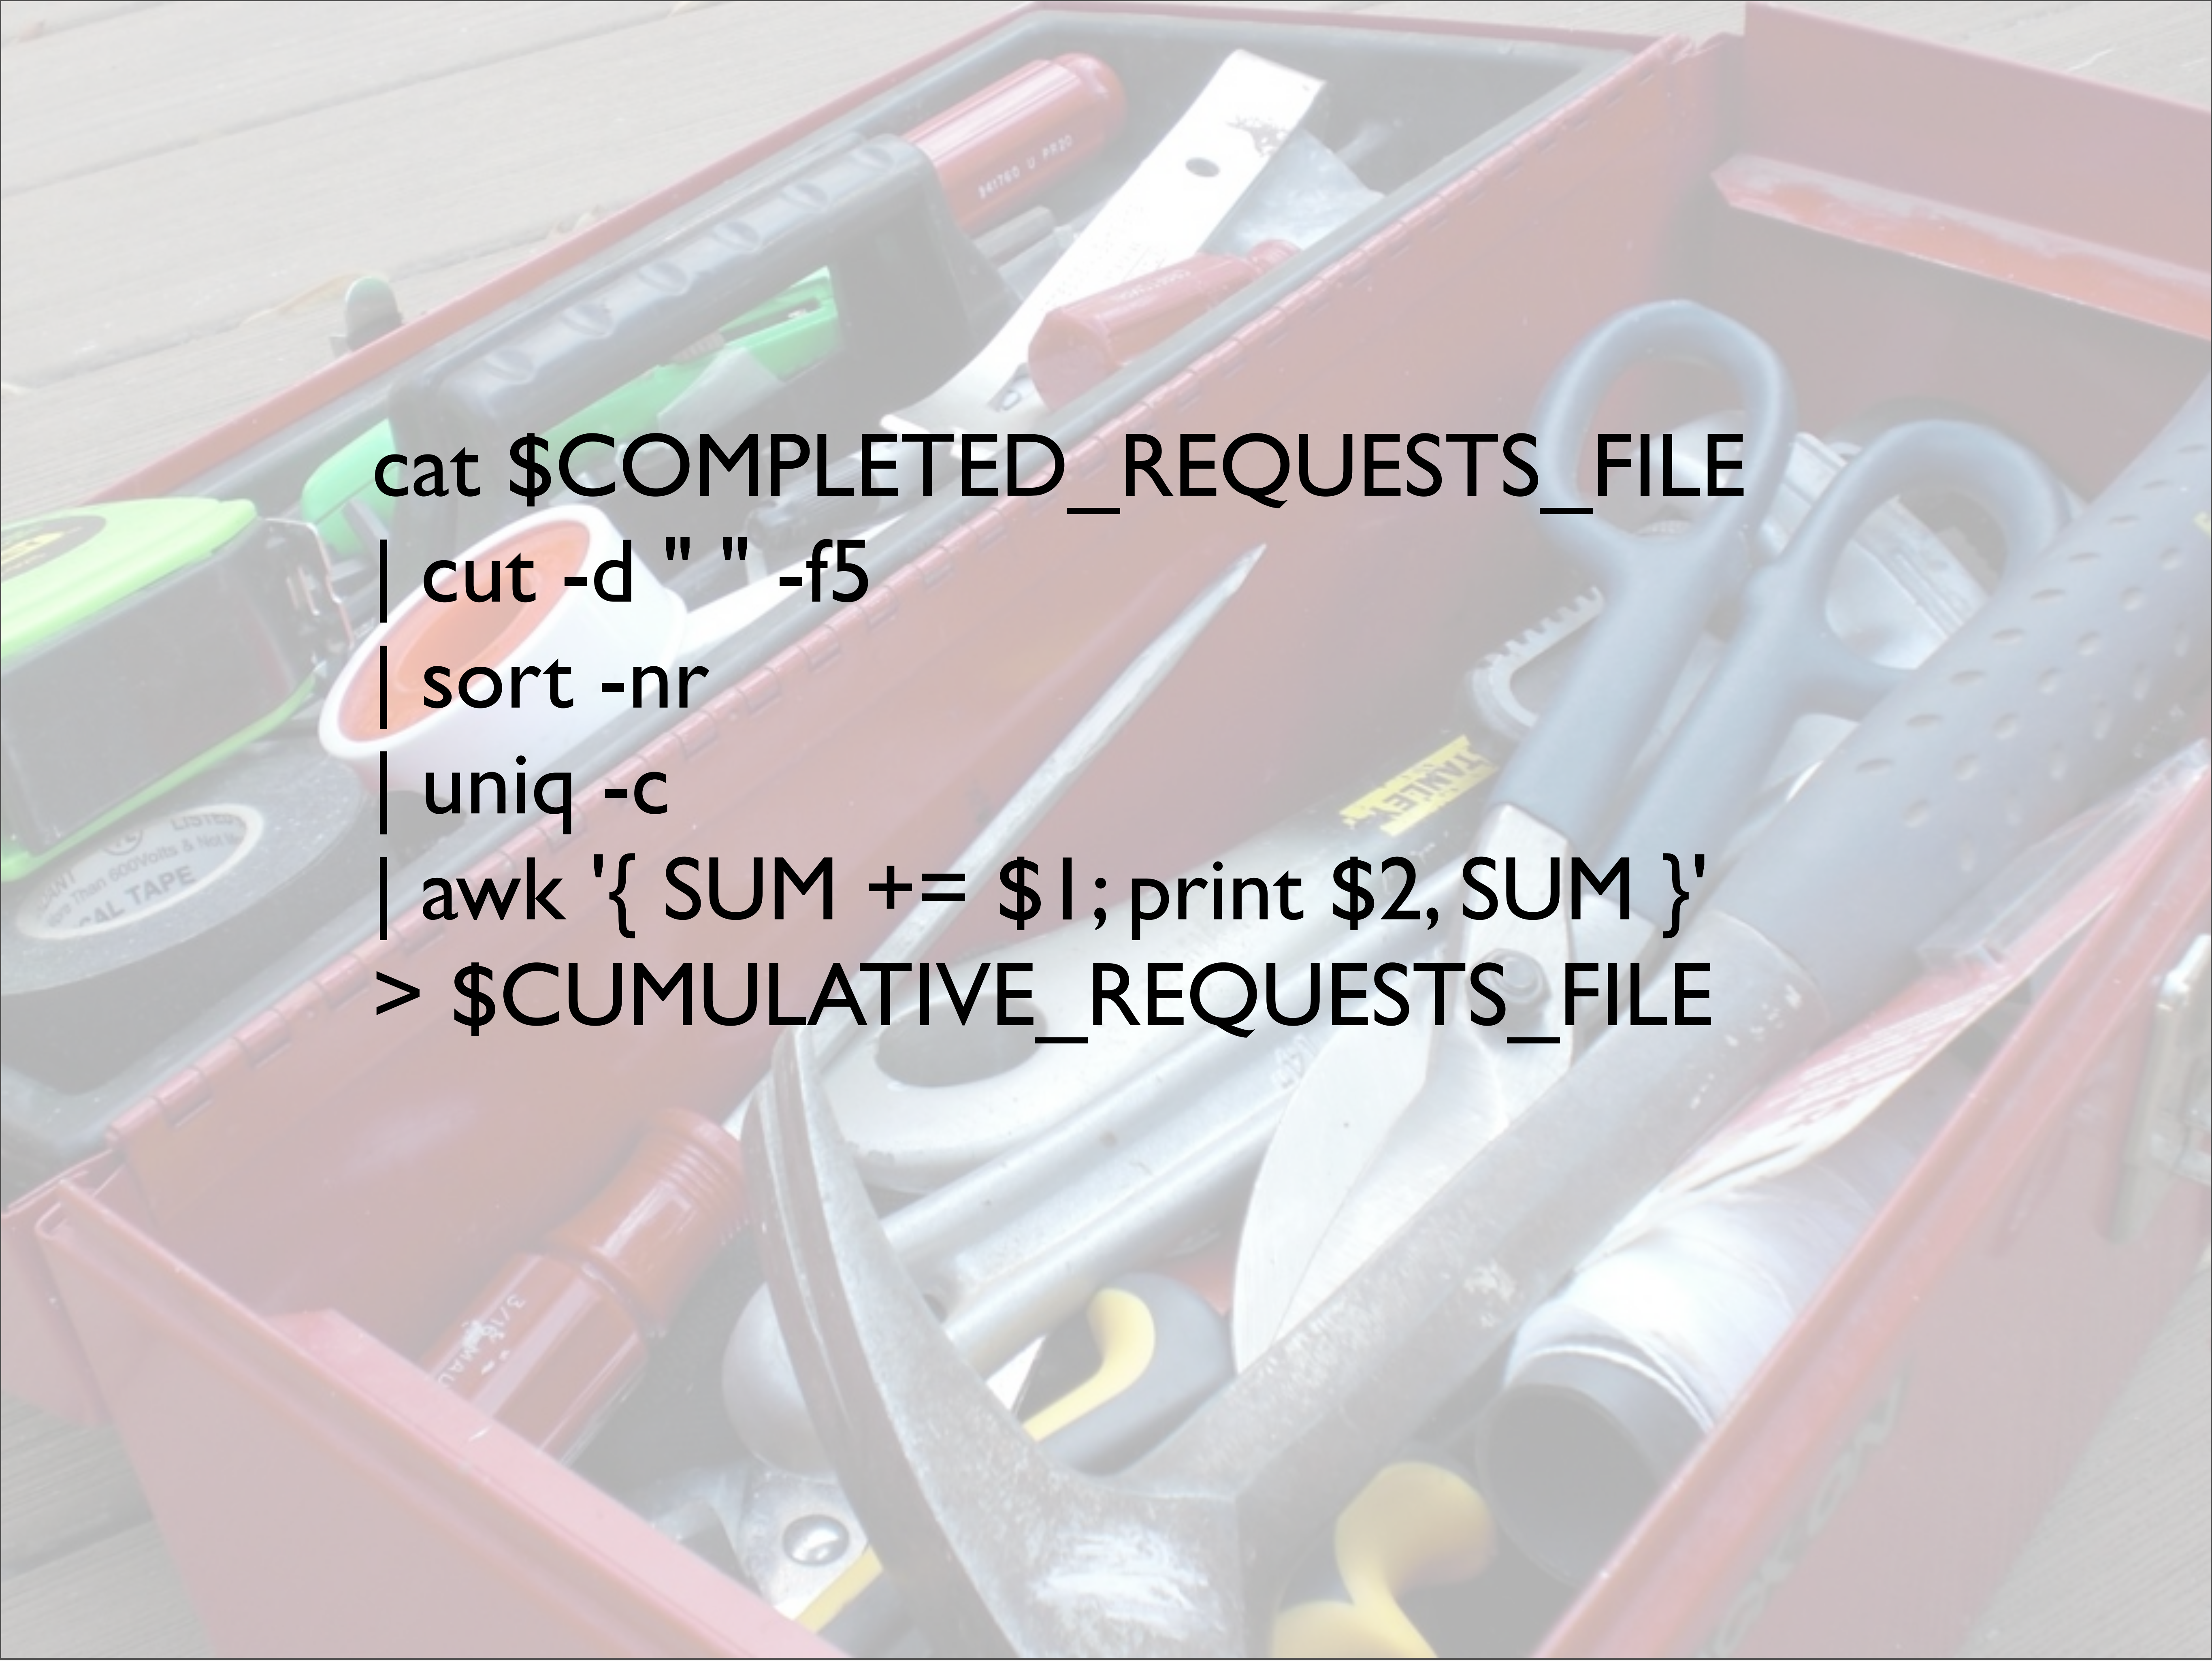

```
cat $COMPLETED_REQUESTS_FILE
| cut -d " " -f5
| sort -nr
| uniq -c
| awk '{ SUM += $1; print $2, SUM }'
> $CUMULATIVE_REQUESTS_FILE
```

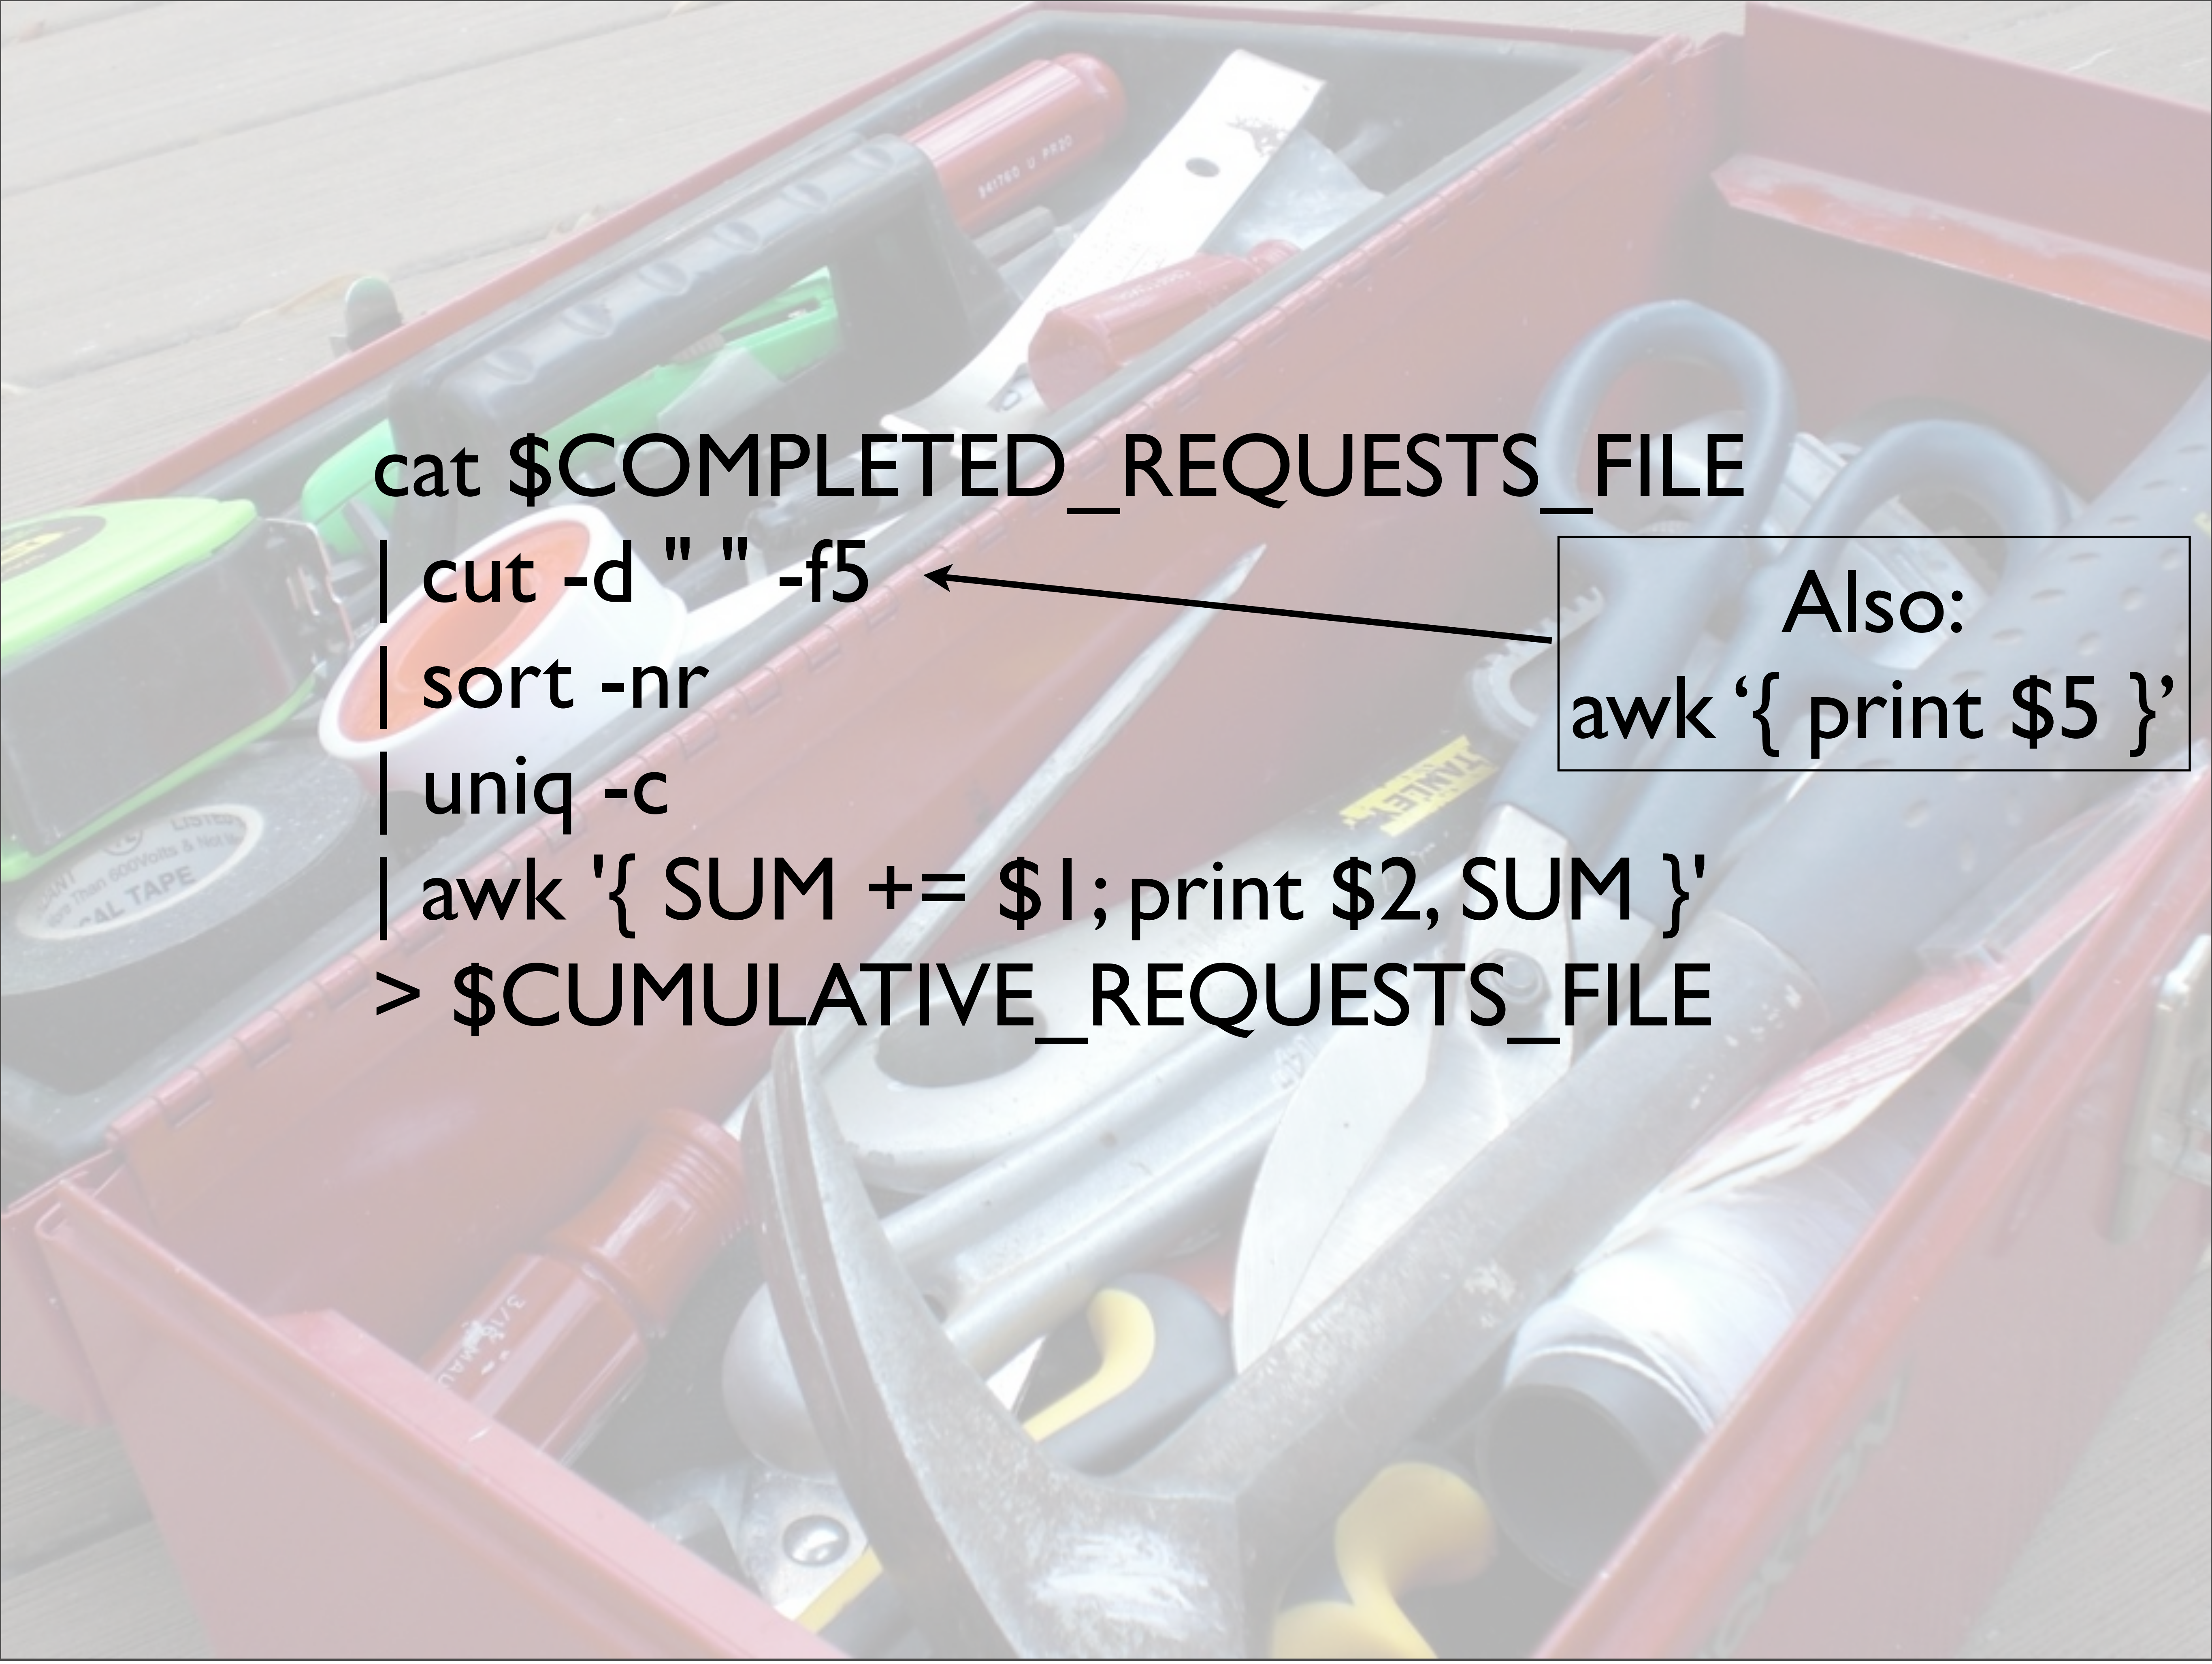

```
cat $COMPLETED_REQUESTS_FILE
| cut -d " " -f5
| sort -nr
| uniq -c
| awk '{ SUM += $1; print $2, SUM }'
> $CUMULATIVE_REQUESTS_FILE
```

Also:
awk ‘{ print $5 }’

You can get complicated

- When sed/awk et al aren’t enough
- Write your own
  - Log parsing into mysql
  - select count(*) from database_calls where request_id in (select id from requests where path like ‘/travel/france/%’)

Other kinds of failure

Not all about software

- Your application
- The system it runs on
- Infrastructure failures
- Network failures
- Bugs

Systems Failure

- Your system itself might get inconsistent
- Garbage collection loops
- Database connections
- Infinite loops
- File Handles

Infrastructure failure

- Power fails
- UPS fails
- Database machine fails
- Your own machine fails

Network failure

- Routers fail
- Uplinks fail
- Internet routing failures
- DNS failures
- Browser issues

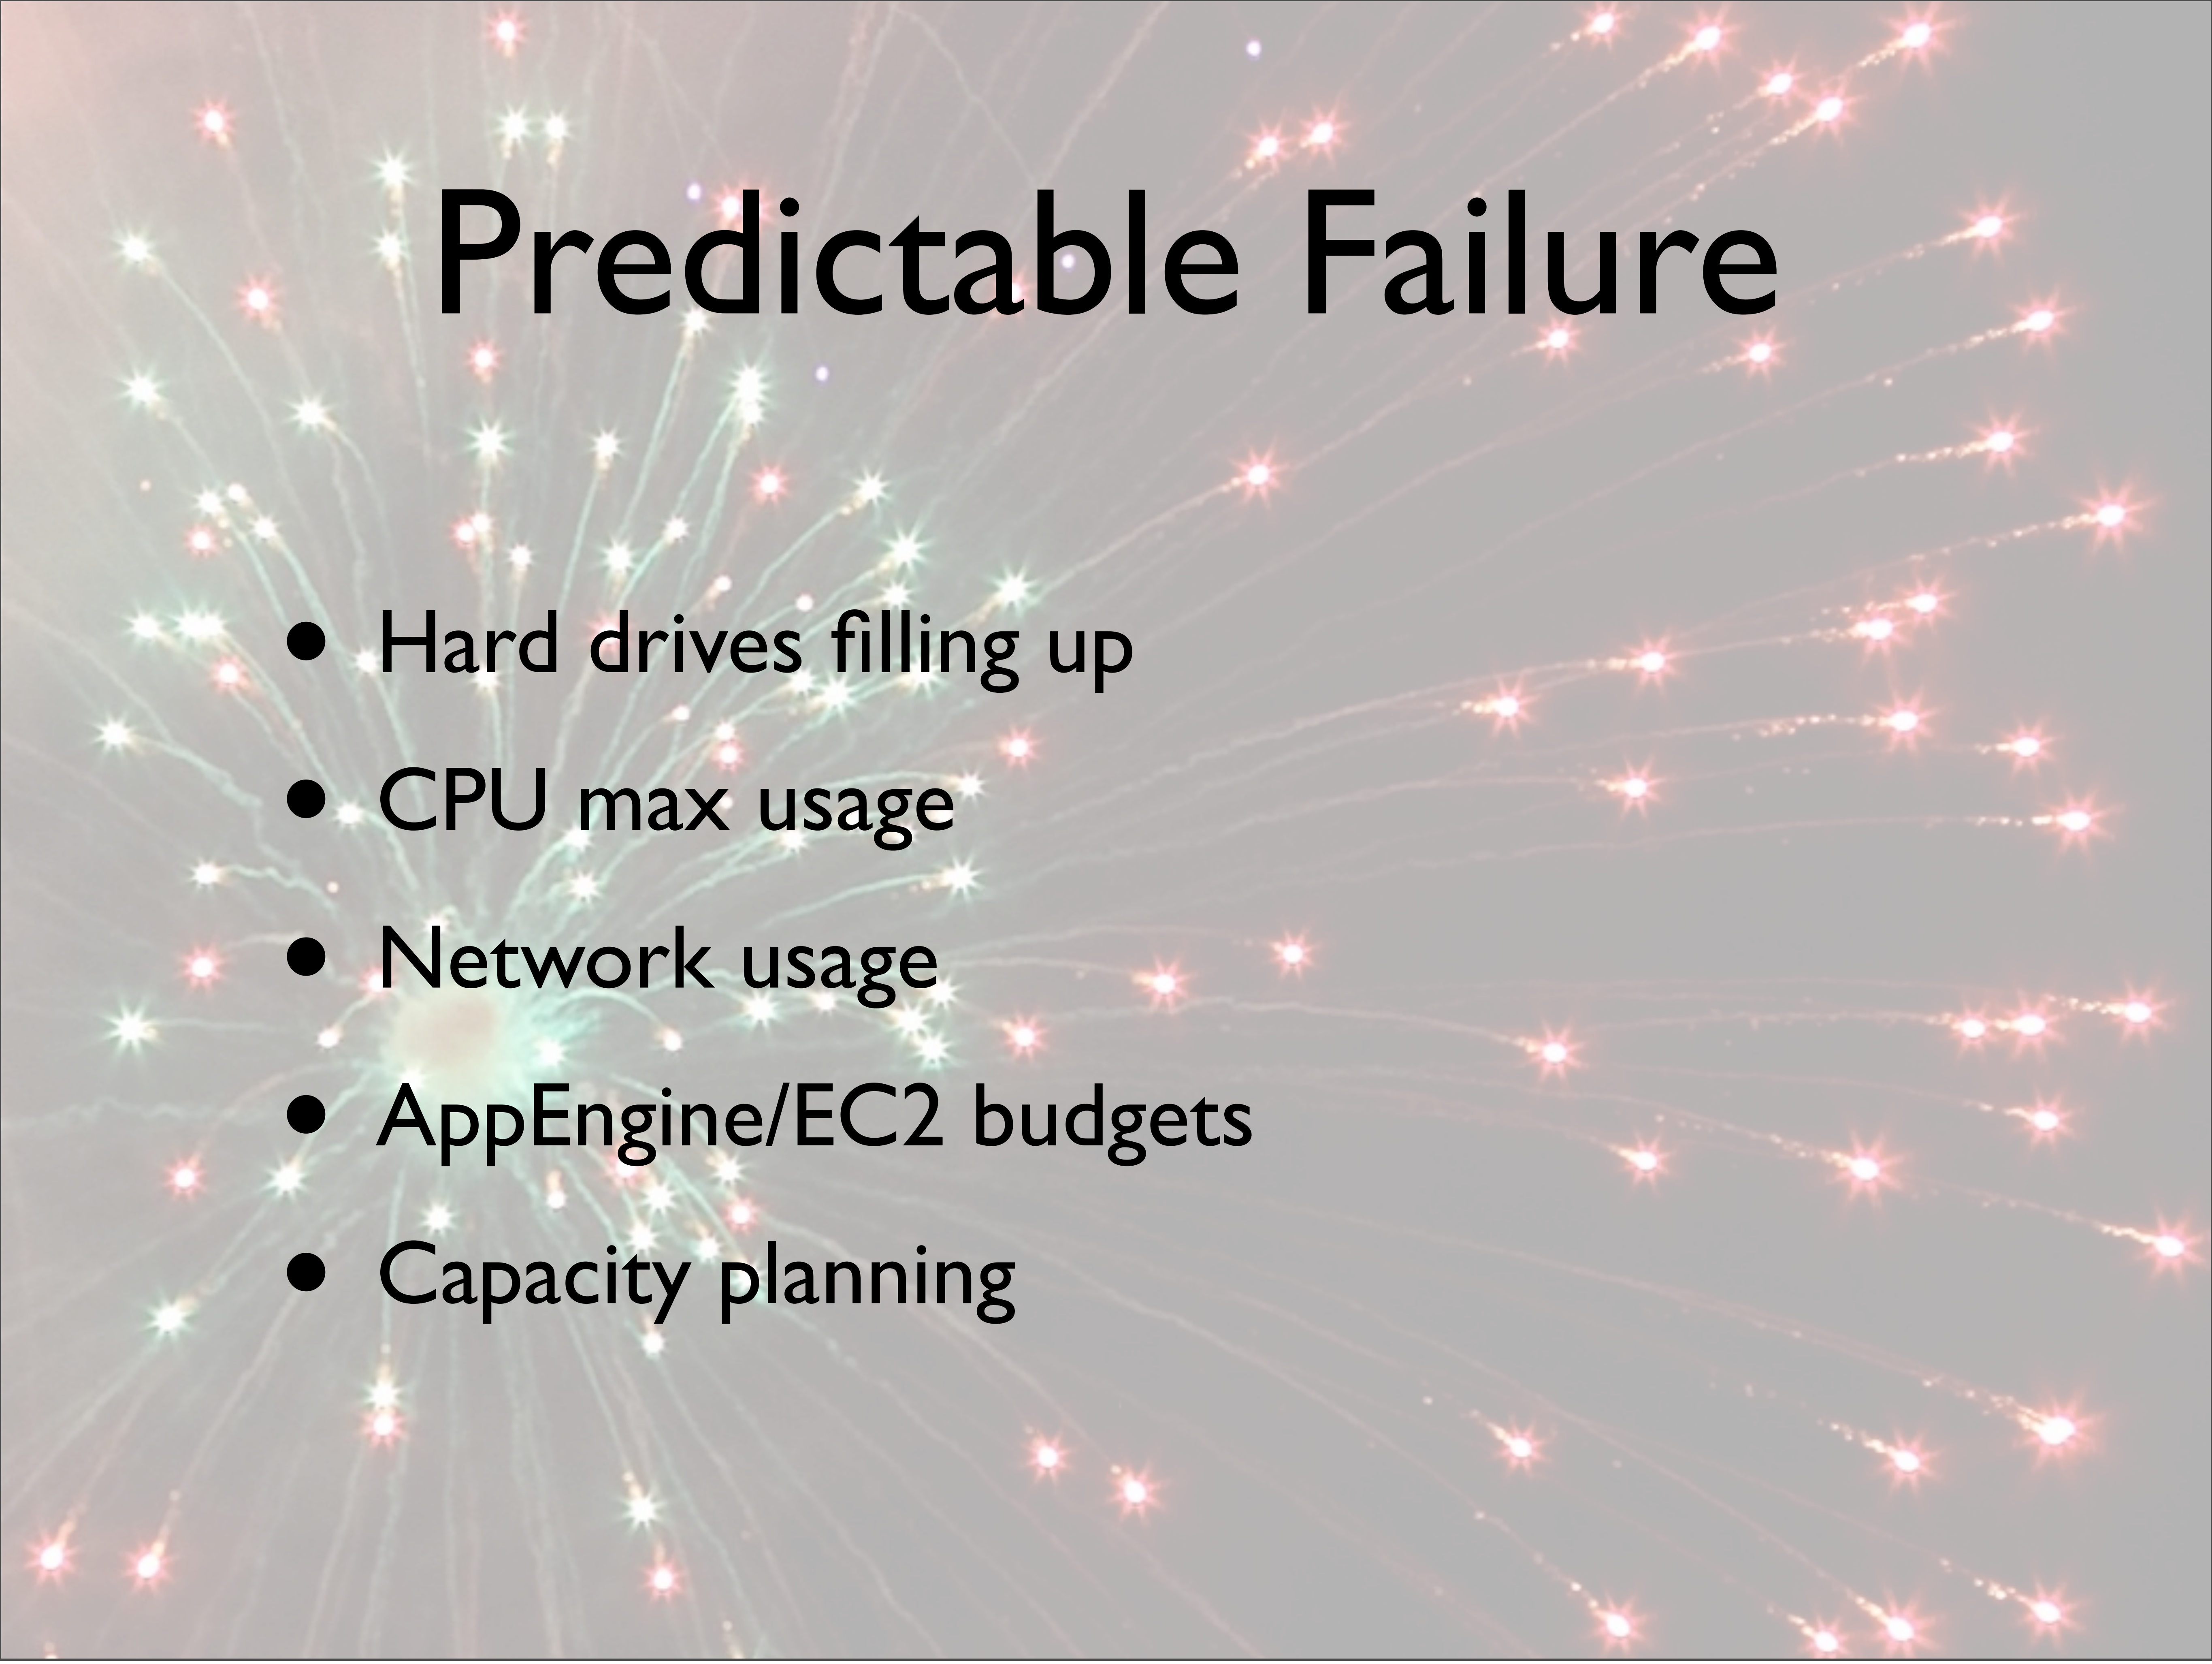

Predictable Failure

- Hard drives filling up
- CPU max usage
- Network usage
- AppEngine/EC2 budgets
- Capacity planning

Unpredictable failure

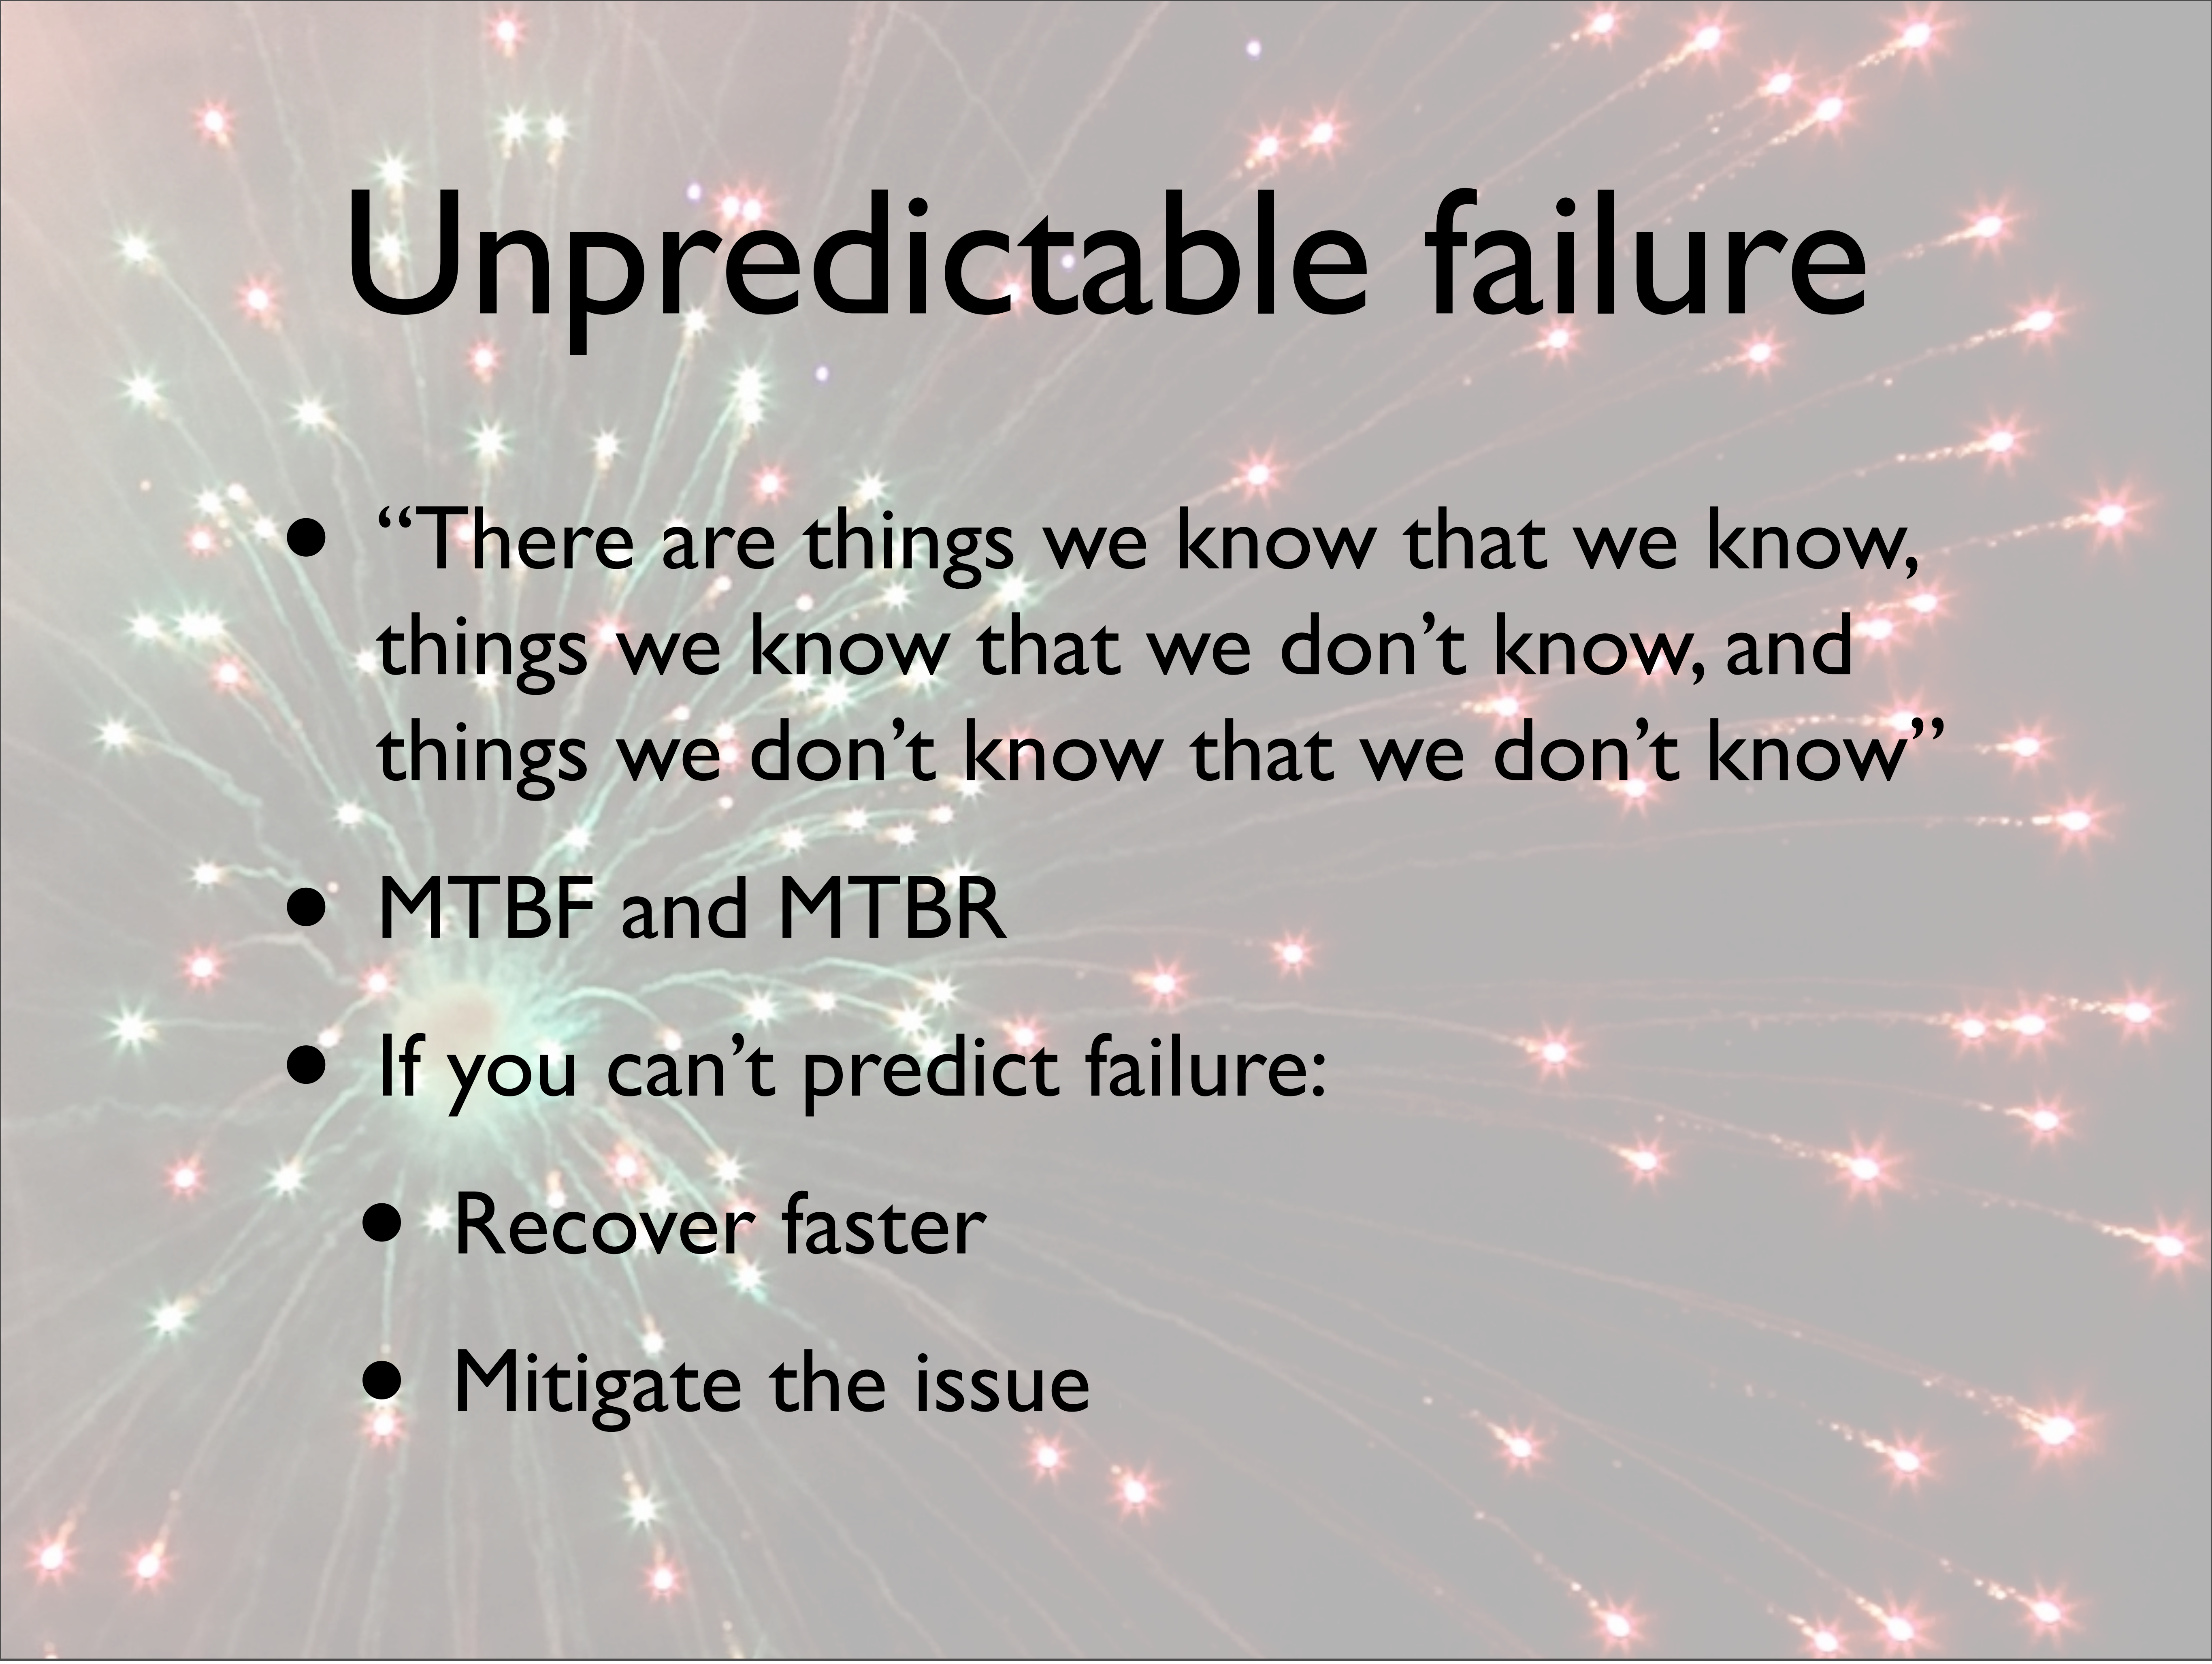

- “There are things we know that we know, things we know that we don’t know, and things we don’t know that we don’t know”
- MTBF and MTBR
- If you can’t predict failure:
  - Recover faster
  - Mitigate the issue

External dependencies

- Who is more likely to break, you or twitter?
- But can you predict when twitter will break?
- Never depend on a third party
- They will let you down
- At the worst possible time

So what have we learnt?

Open Platform

- Need to handle peaky load
- Fault isolation from main database
- feels like we’ve been here before...

Content API Architecture

Apache

AppServer

Solr

Content API Architecture

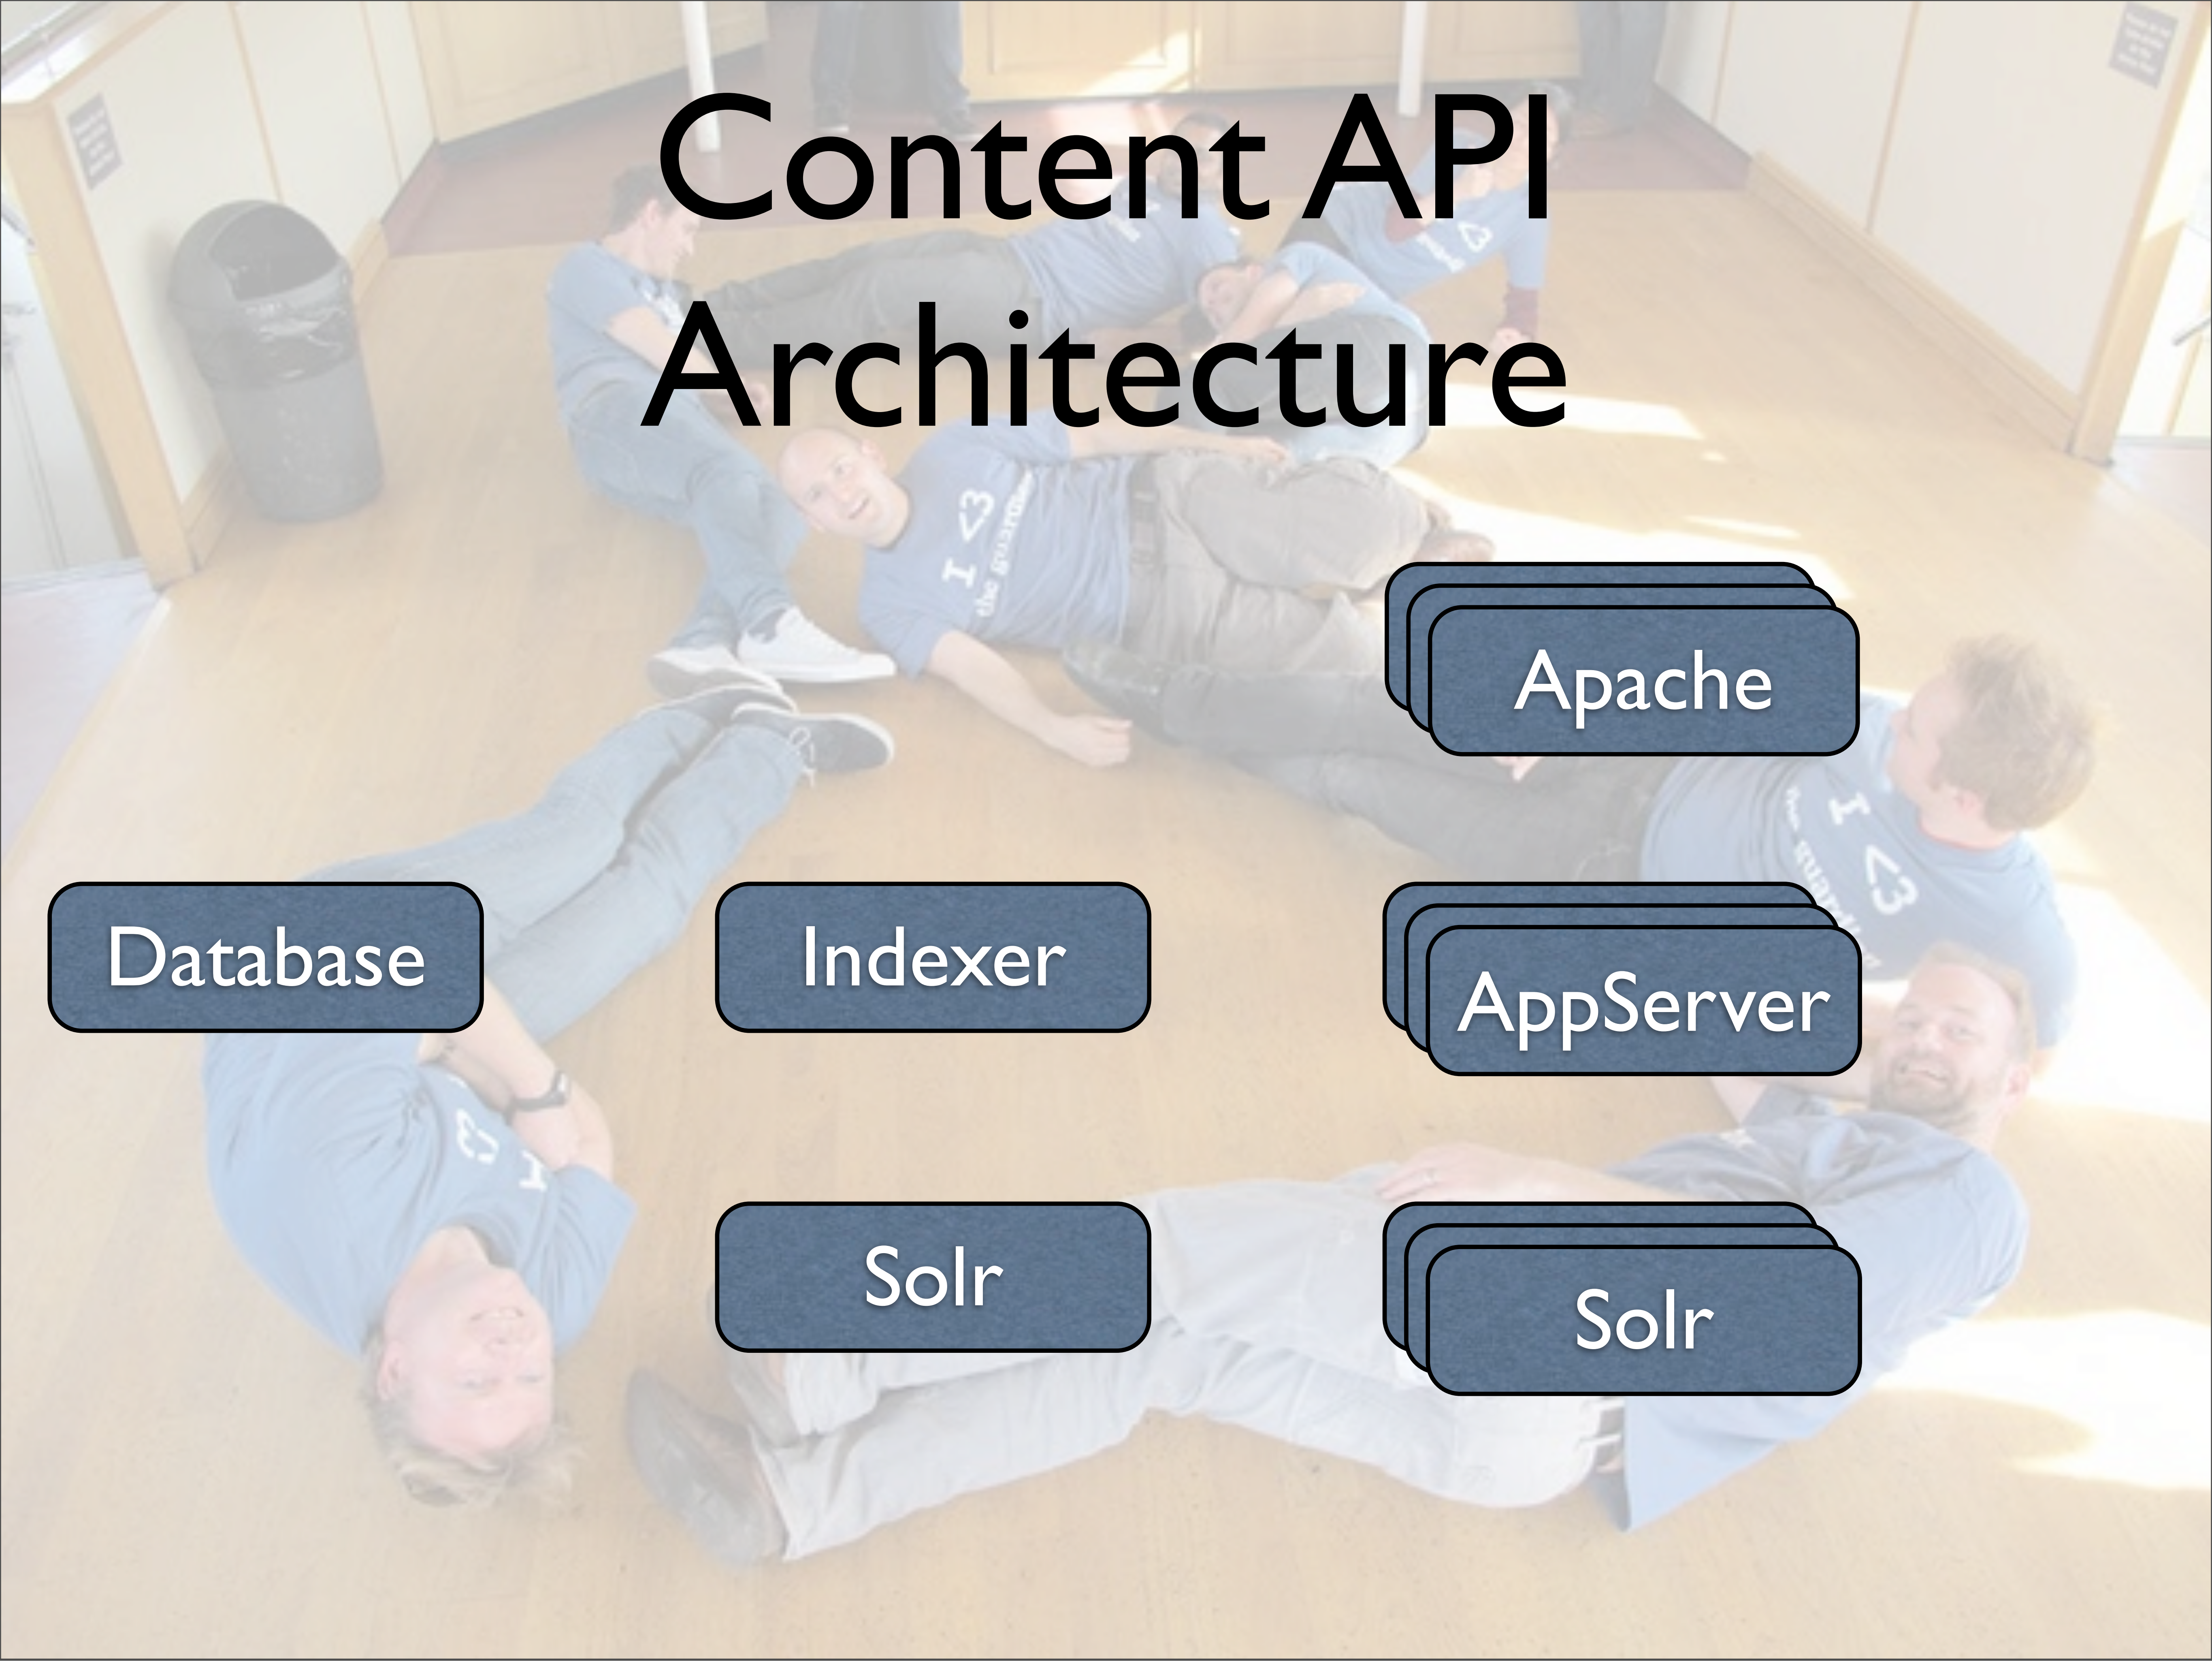

Apache

Database

Indexer

AppServer

Solr

Solr

Content API Architecture

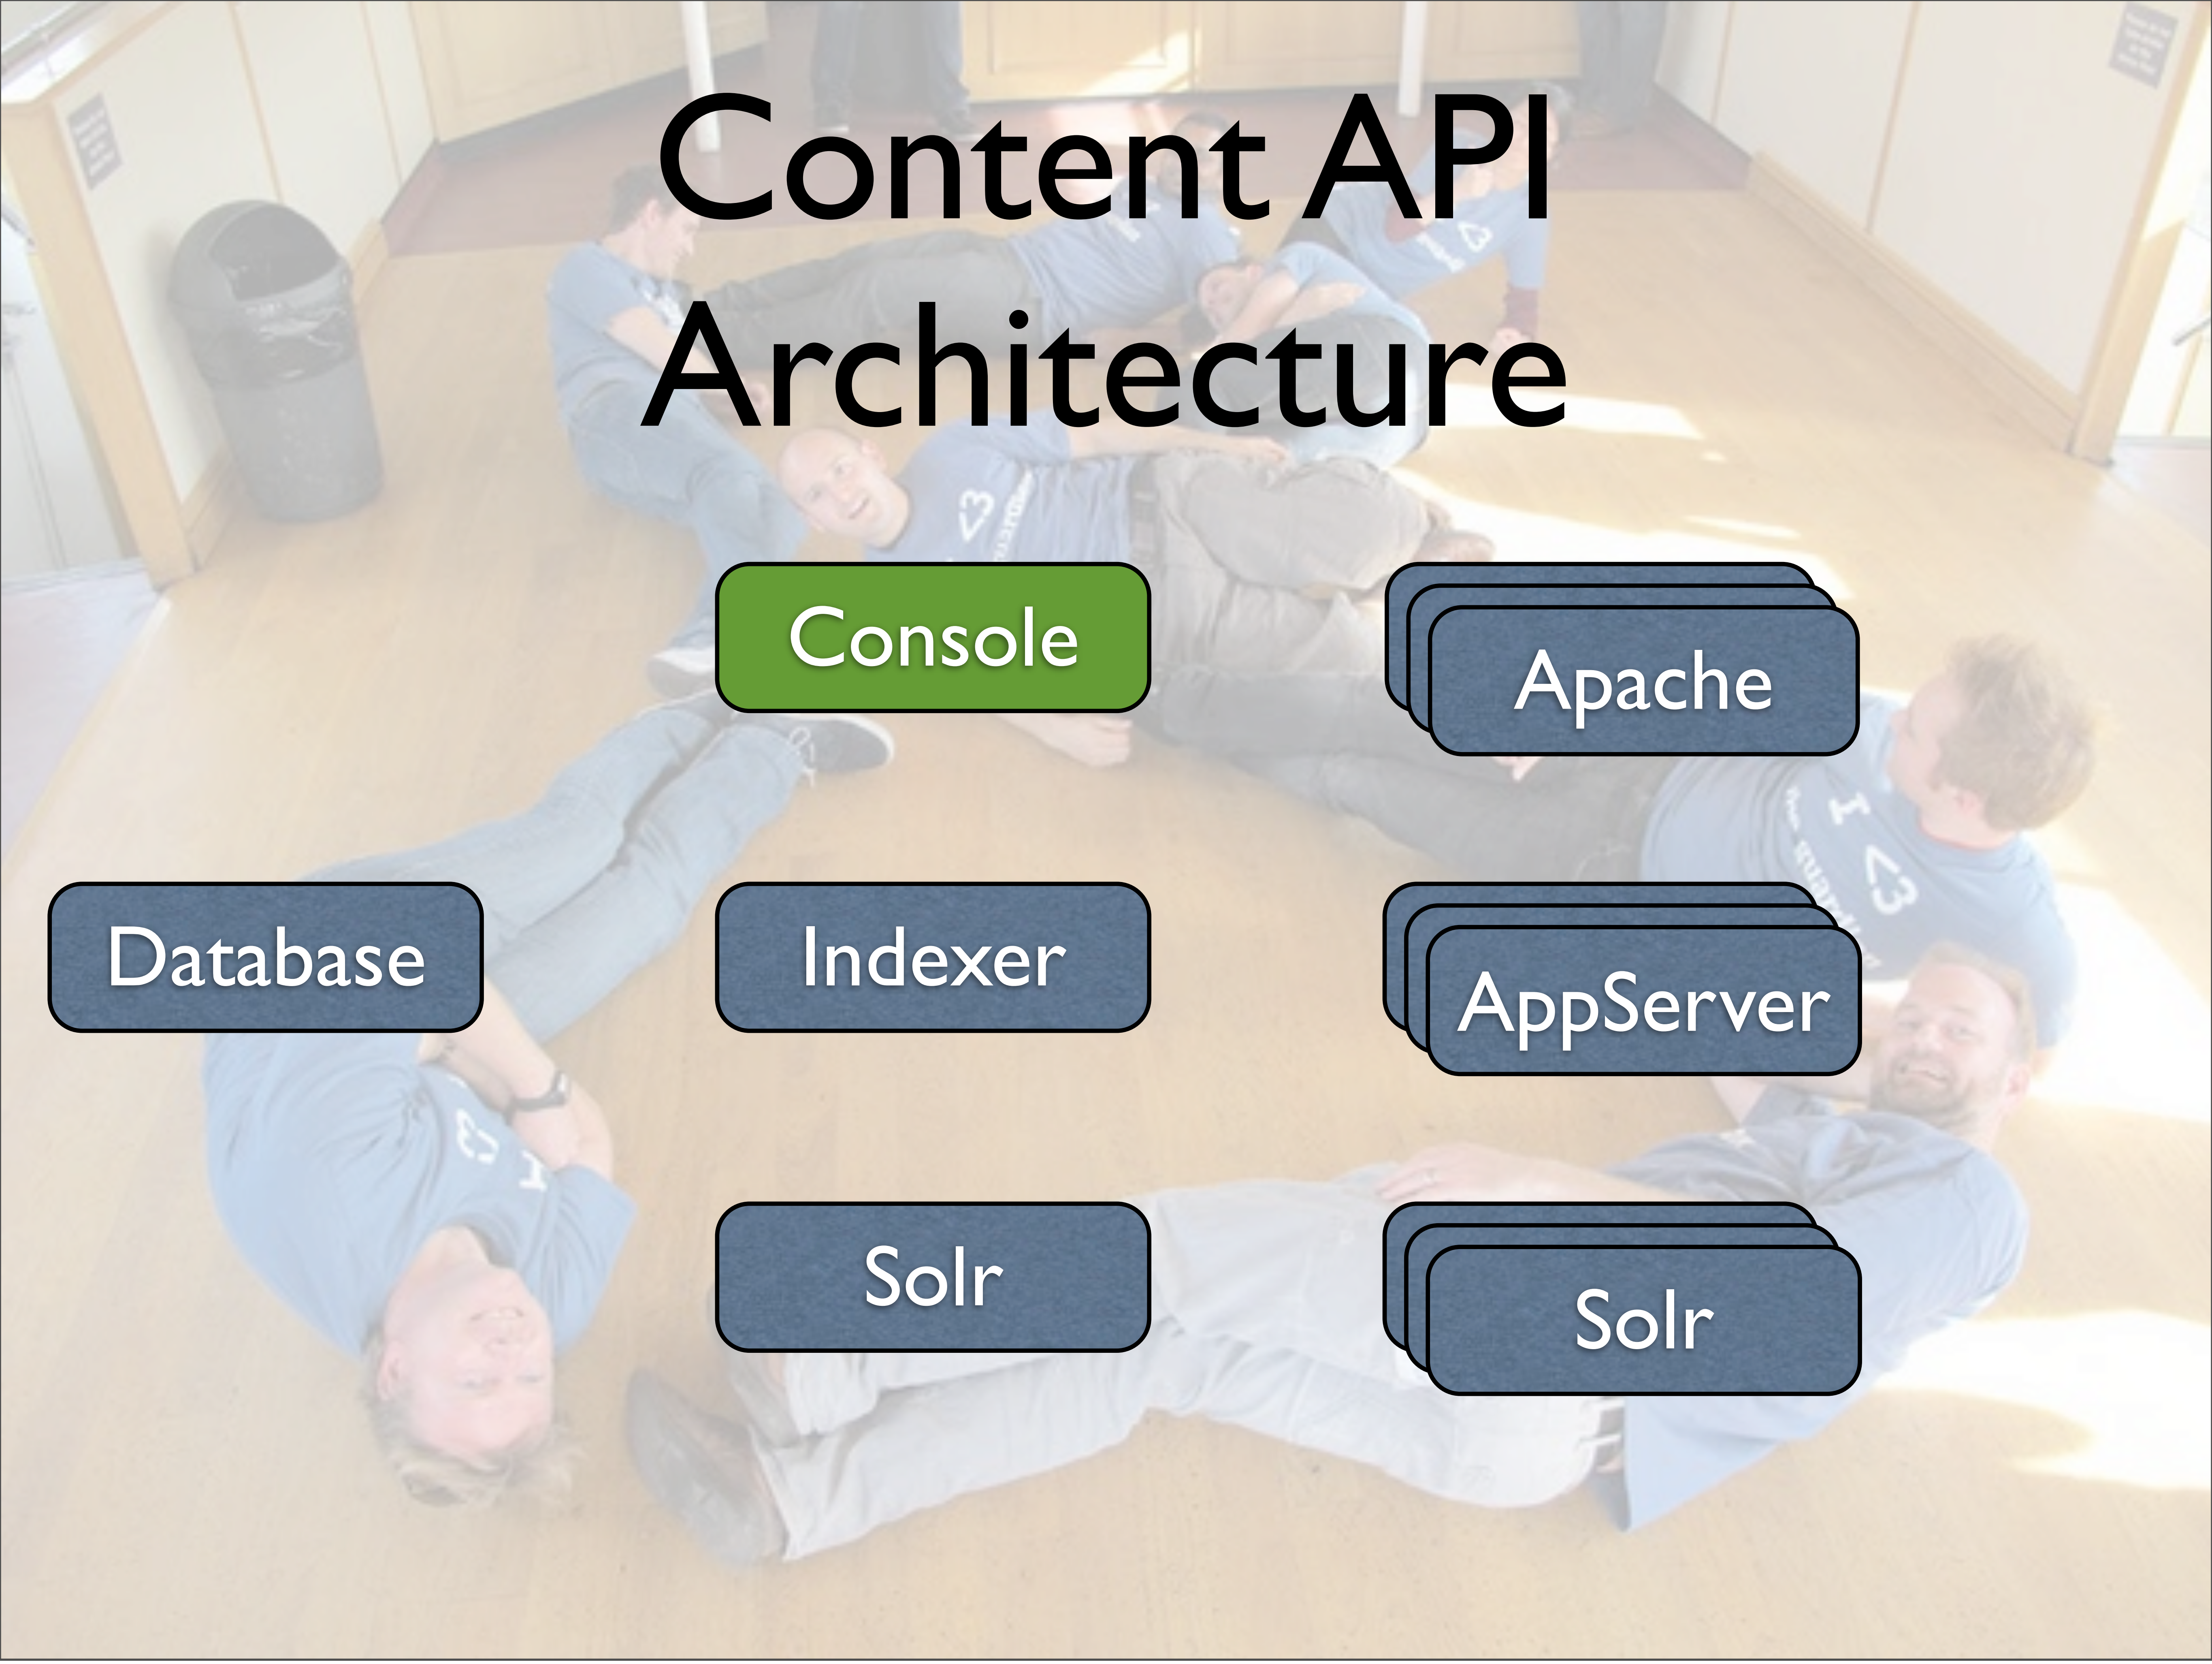

Console

Apache

Database

Indexer

AppServer

Solr

Solr

Benefits

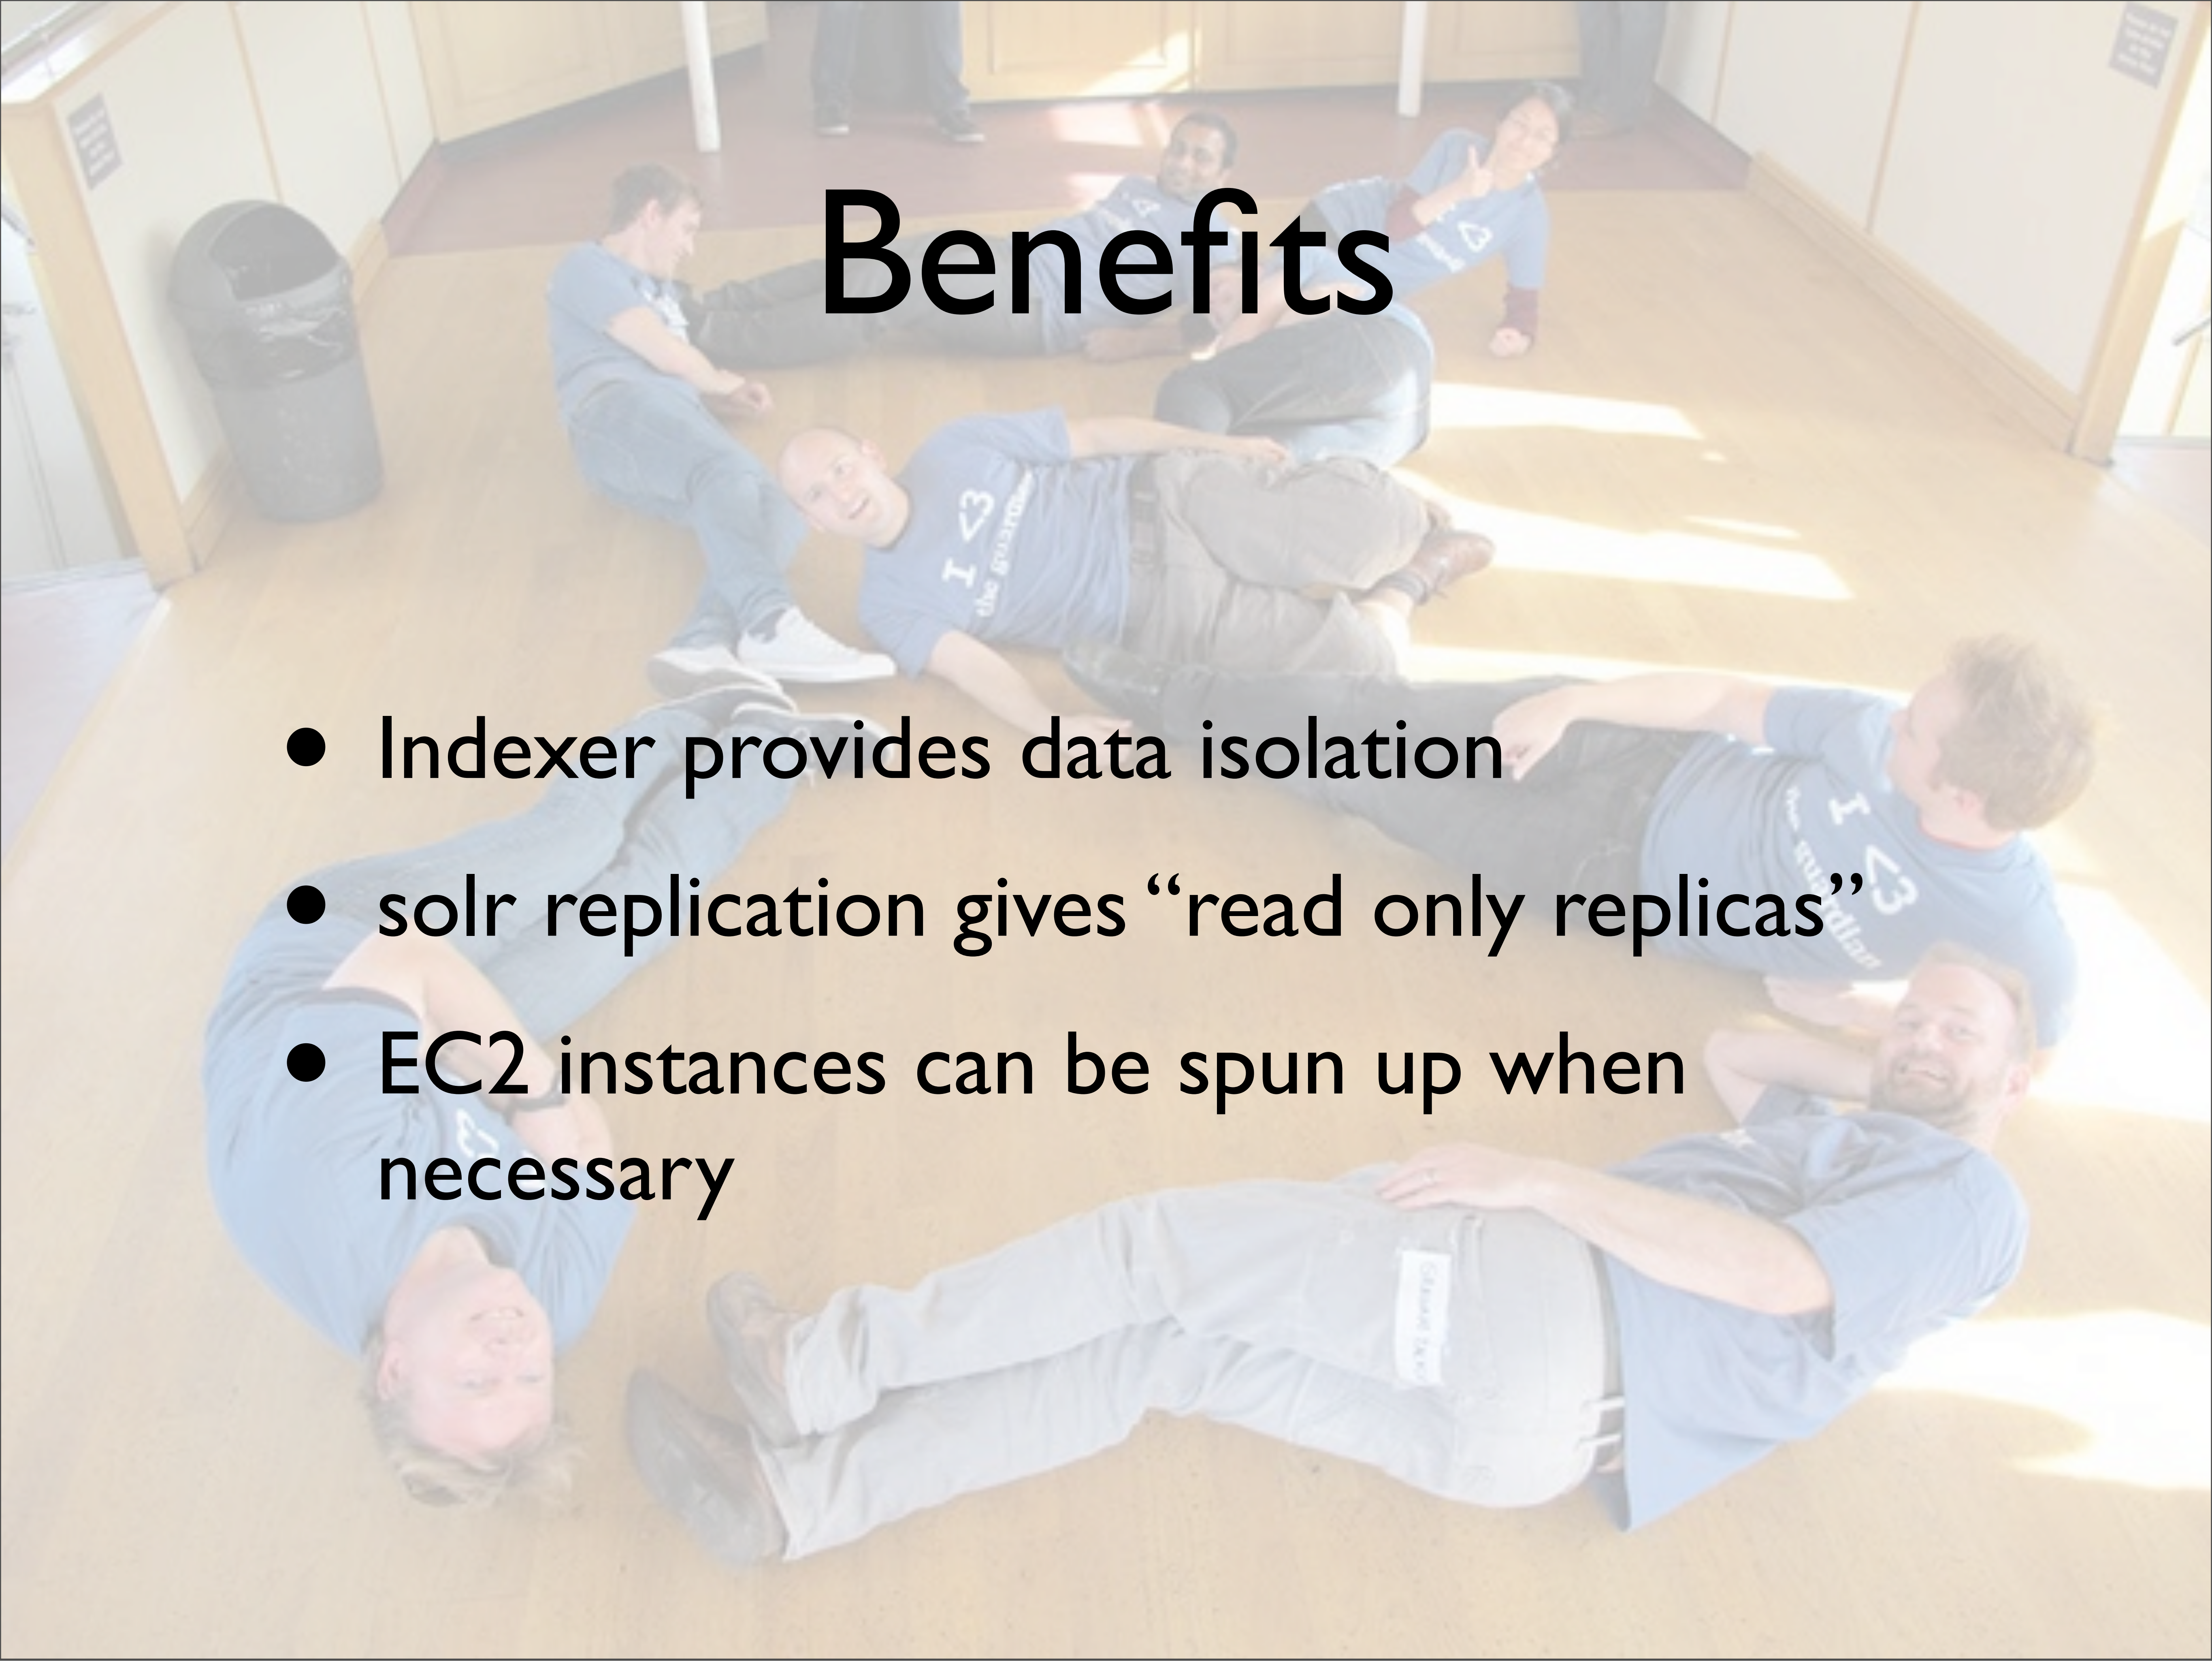

- Indexer provides data isolation
- solr replication gives “read only replicas”
- EC2 instances can be spun up when necessary

Benefits

- Switches on backend
  - Indexing
  - Features
  - Replication
- Switches in API
  - content.guardianapis.com/.json?show-

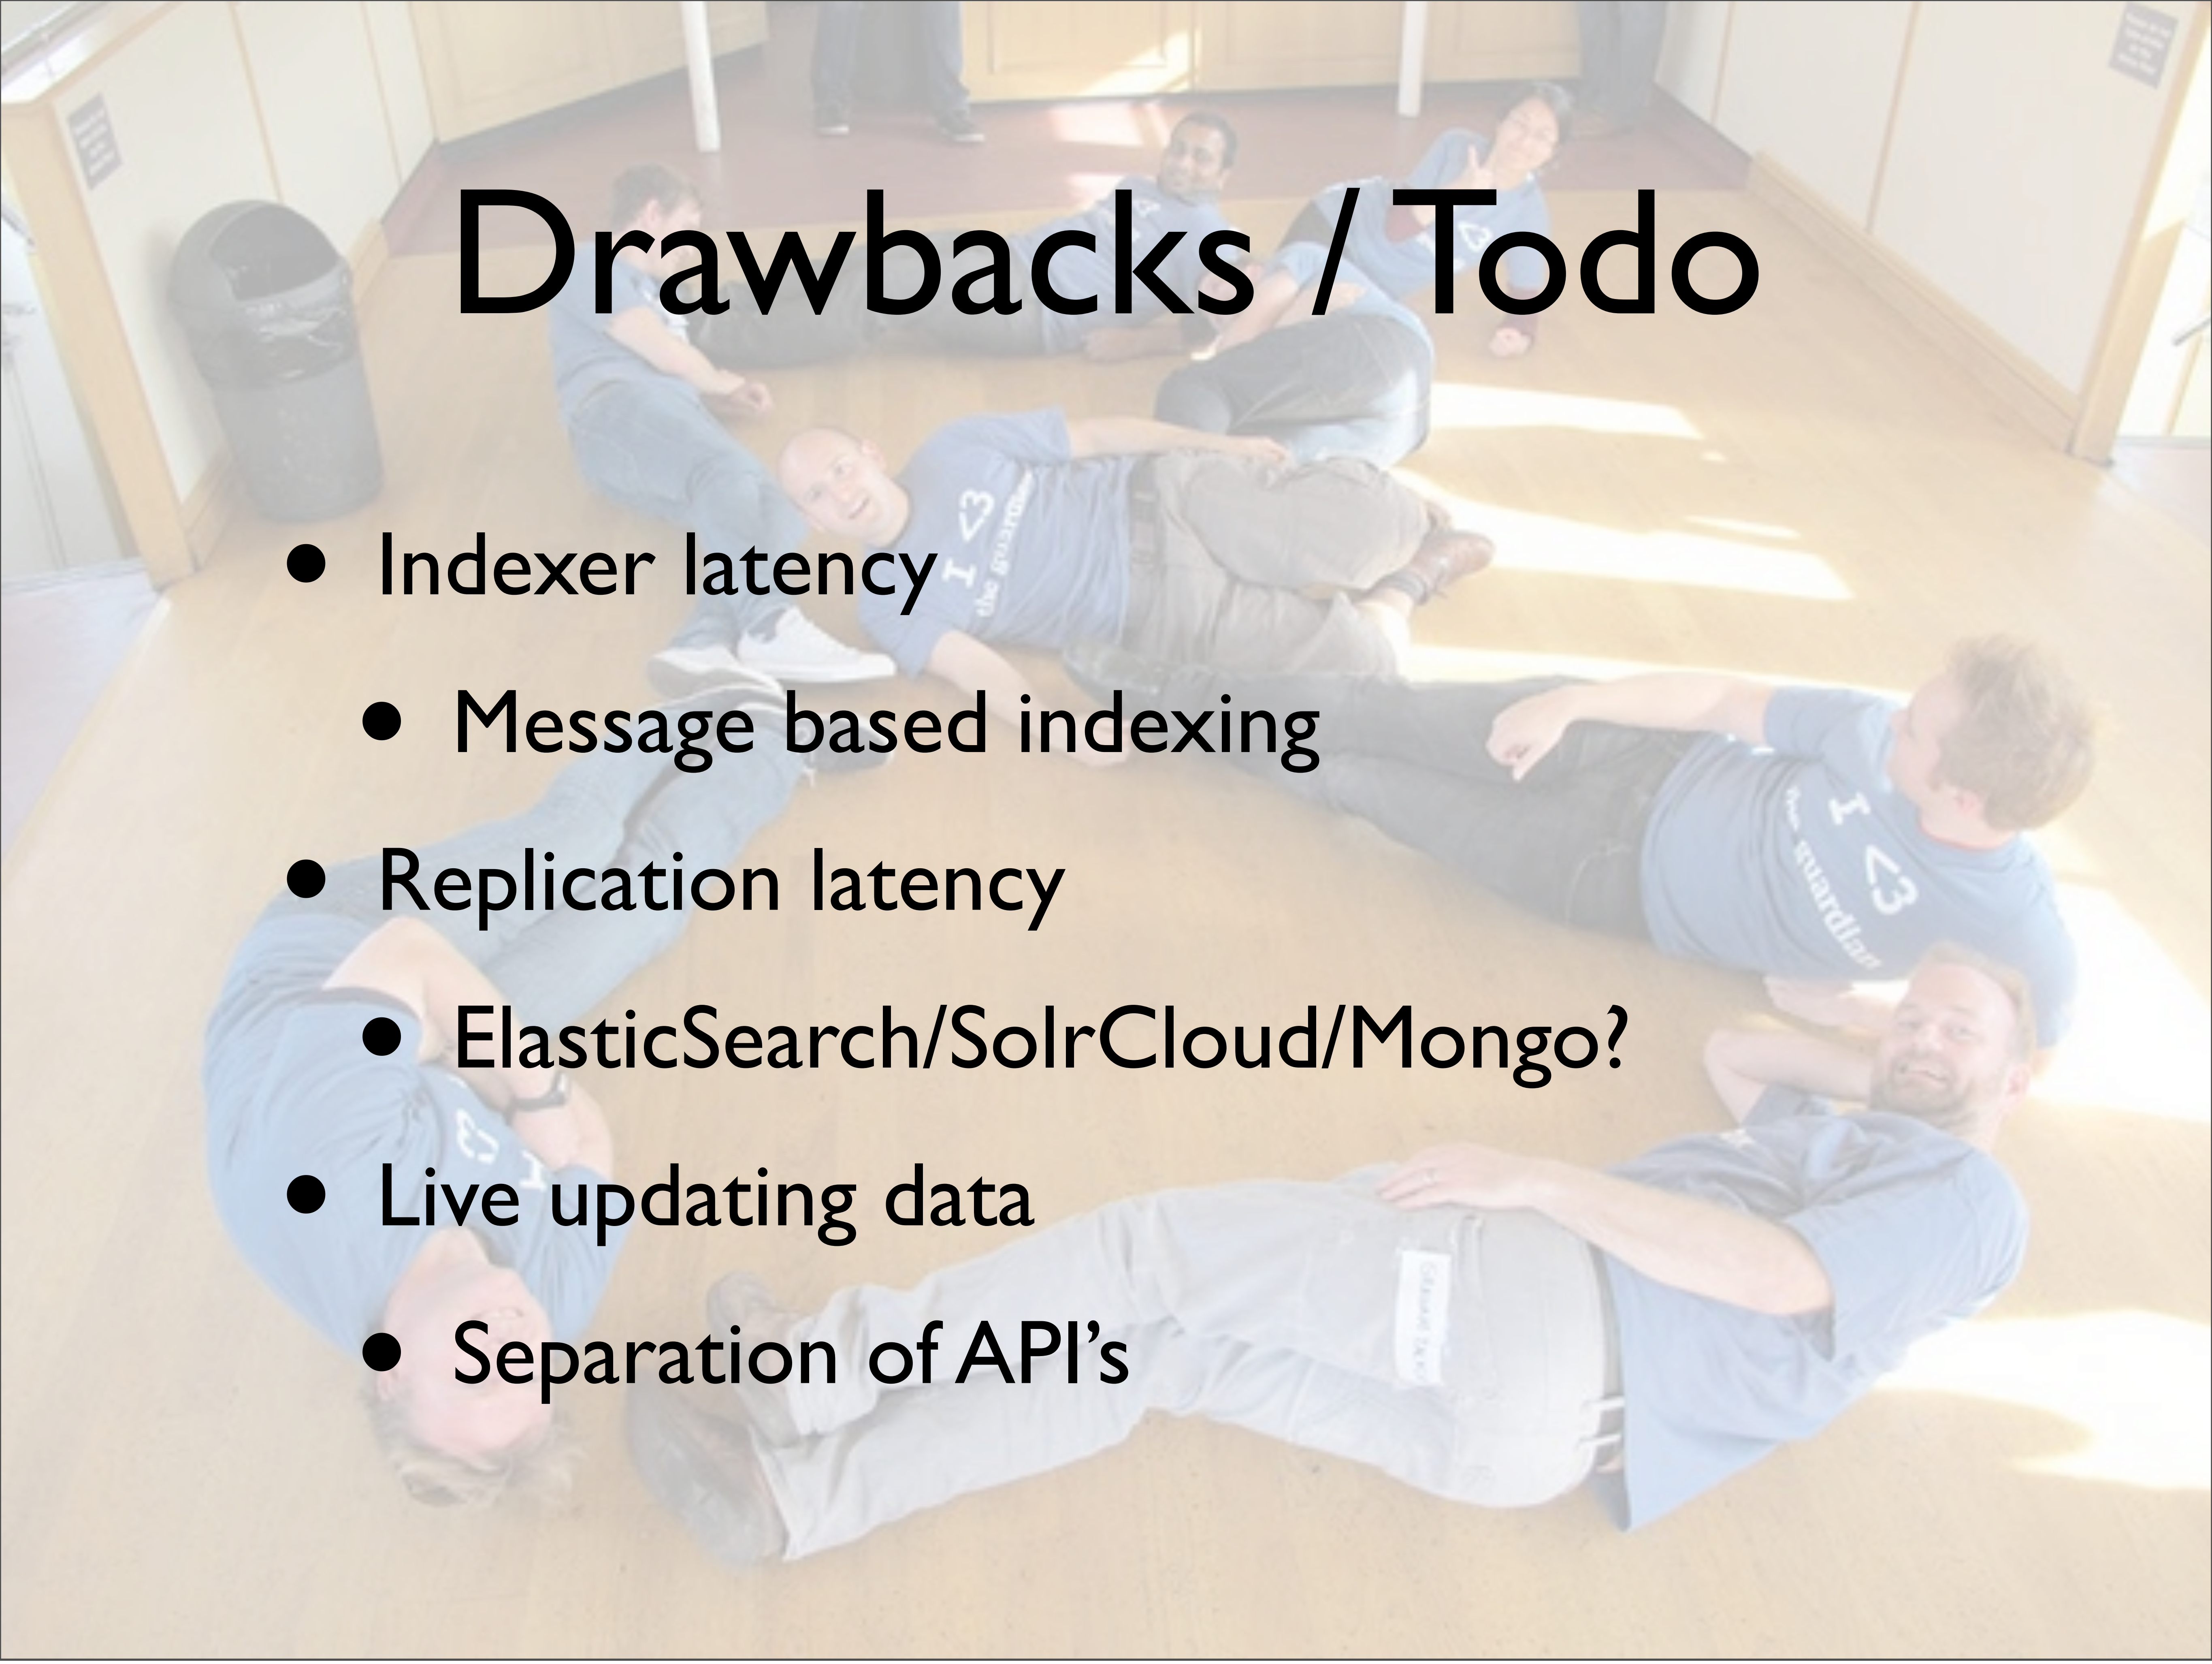

Drawbacks / Todo

- Indexer latency
  - Message based indexing
- Replication latency
  - ElasticSearch/SolrCloud/Mongo?
- Live updating data
  - Separation of API’s

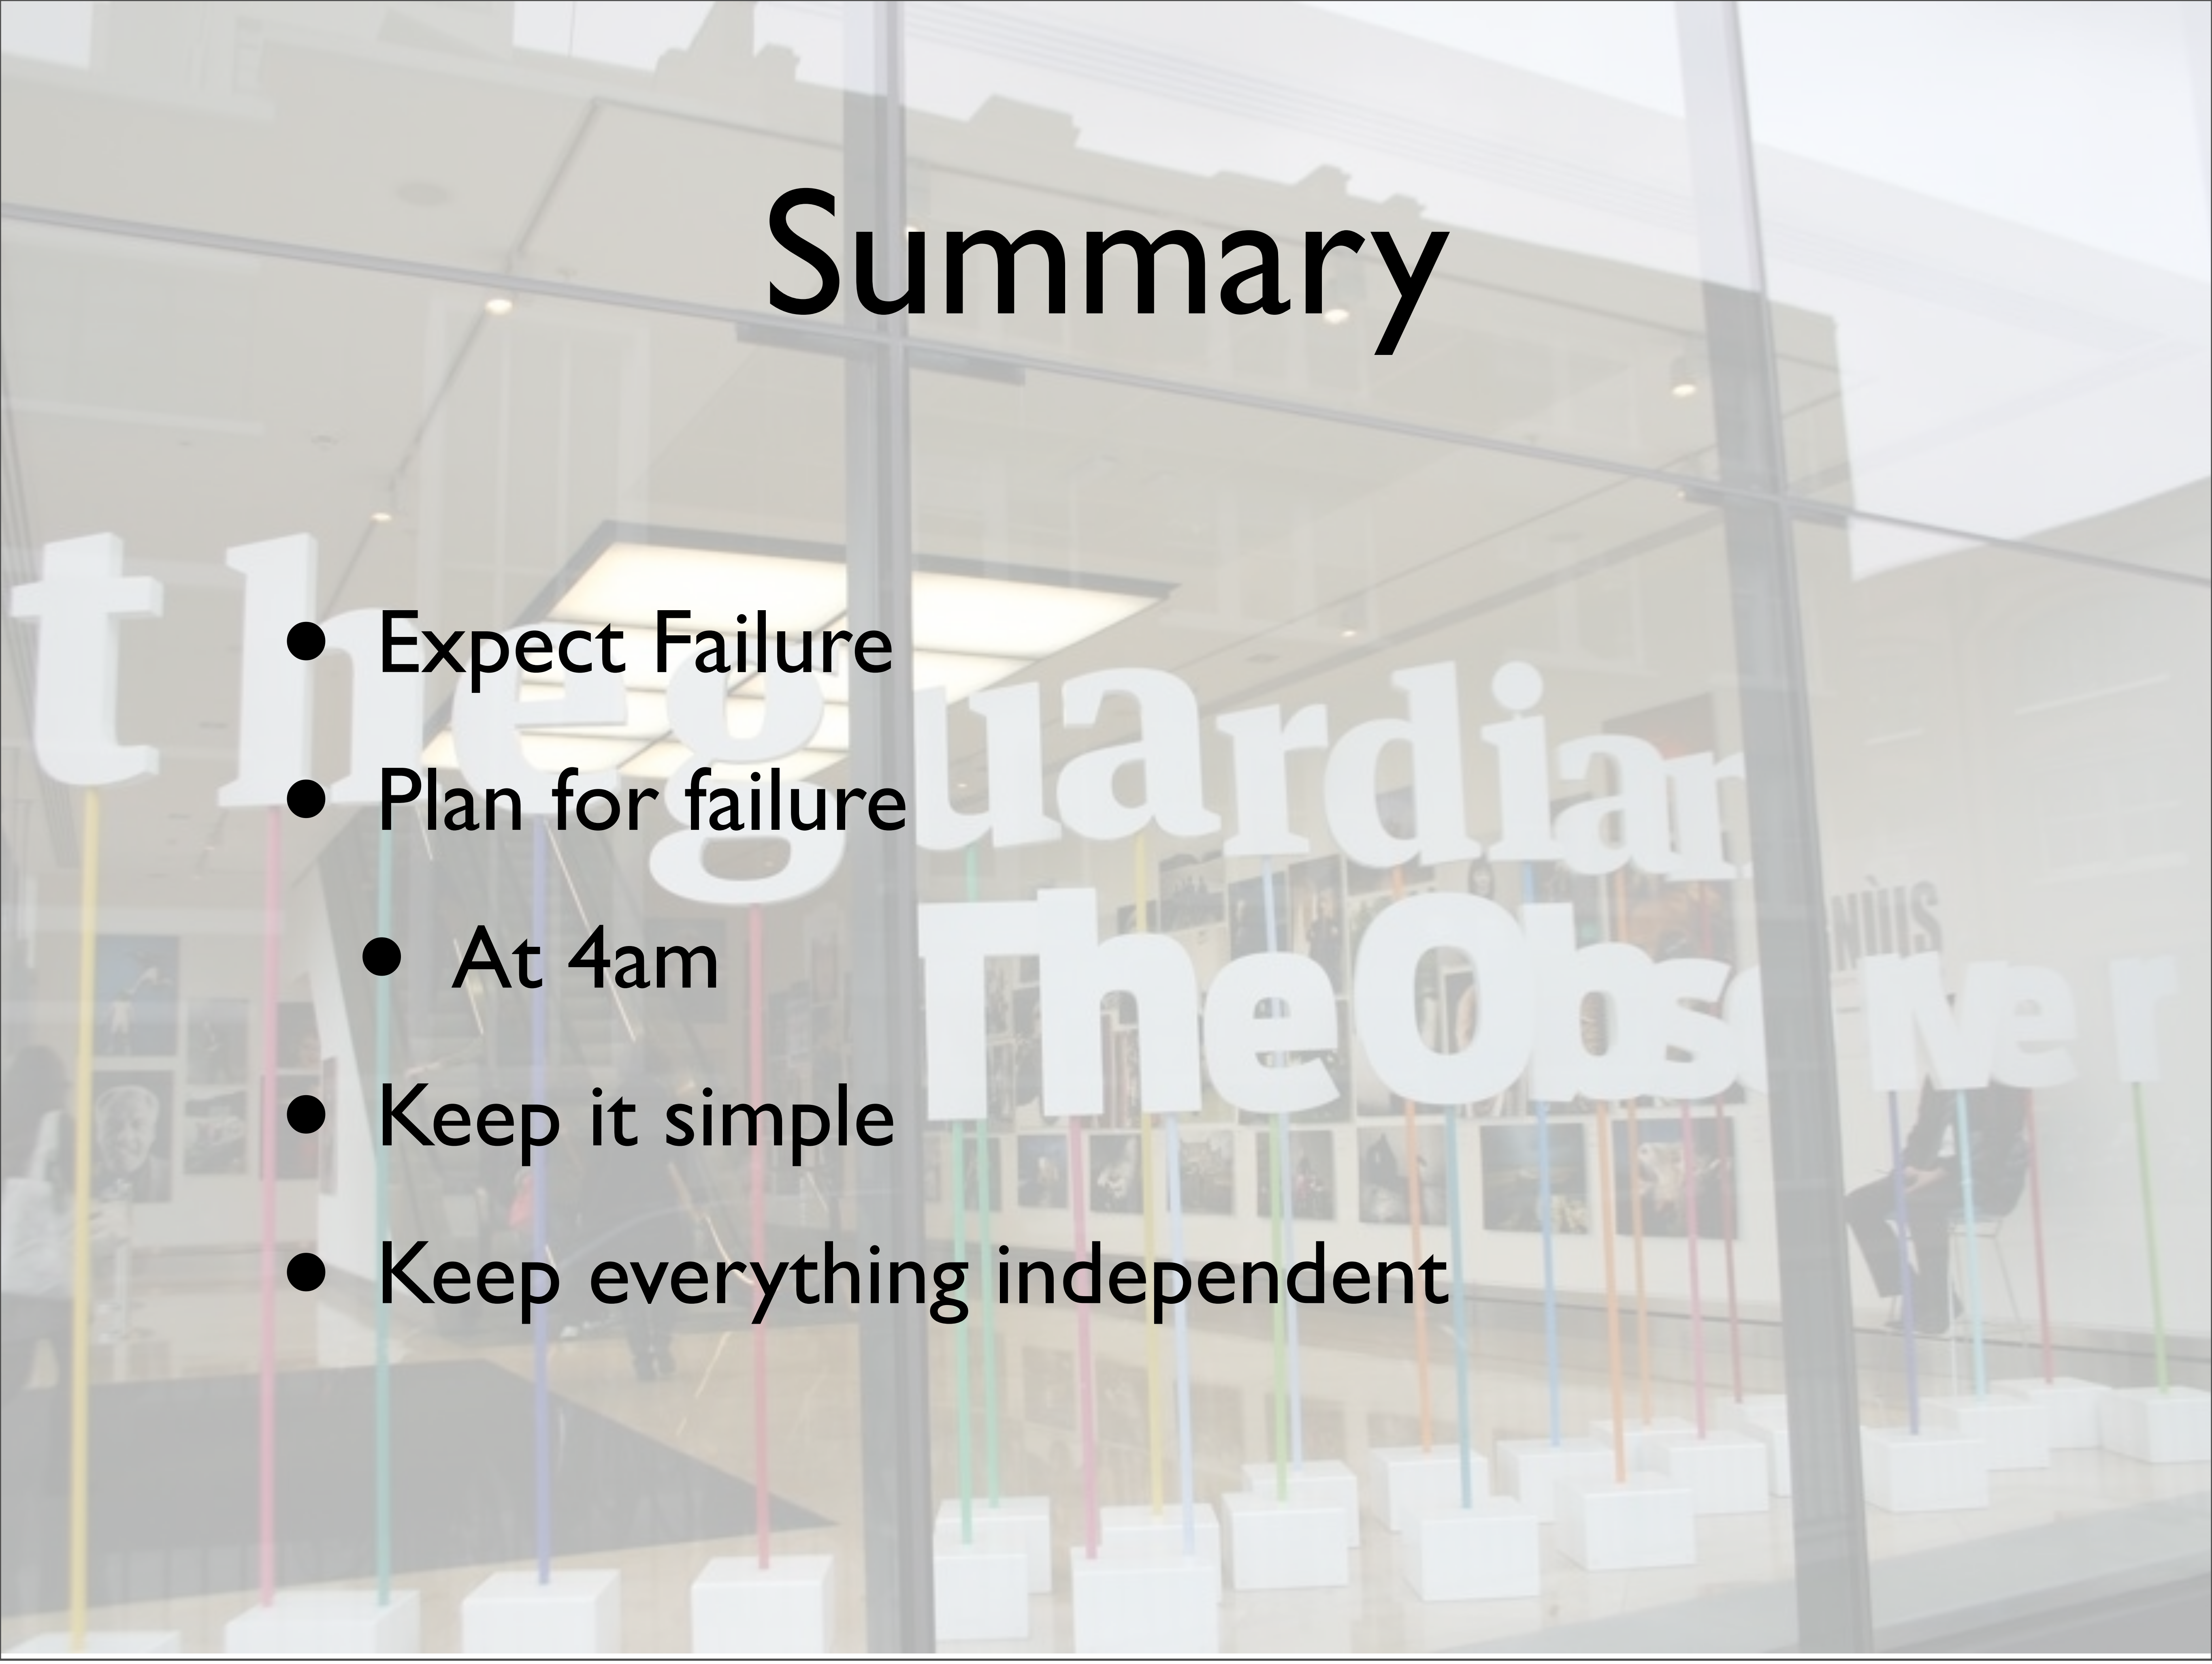

Summary

- Expect Failure
- Plan for failure
    - At 4am
- Keep it simple
- Keep everything independent

Thank You

- michael.brunton-spall@guardian.co.uk
- @bruntonspall
- Thanks to Lisa van Gelder (@techbint), Mat Wall (@matwall), Philip Wills (@philwills) and Graham Tackley (@tackers)

Giant Furry Rat - "Lost land of the Volcano"courtesy of BBC natural history unit
Panic Button - http://www.flickr.com/photos/trancemist/361935363/
Long Meg Sidings - http://www.flickr.com/photos/ingythewingy/5243875486/
Server Rack - http://www.flickr.com/photos/jamisonjudd/2433102356/
Release Valve - http://www.flickr.com/photos/kayveeinc/4107697872
Ancient Planet - http://www.flickr.com/photos/gsfc/4479185727/
Solar system - http://www.flickr.com/photos/gsfc/4479185727
Gauges - http://www.flickr.com/photos/dgoodphoto/5264024028
Logs - http://www.flickr.com/photos/catzrule/5693655199
Higgs boson - http://www.flickr.com/photos/jurvetson/4233962874
Toolbox - http://www.flickr.com/photos/jrhode/4632887921