Breaking the Monolith

Microservice Extraction at SoundCloud

Soundcloud

- 11 hours uploaded every minute
- 150 million tracks
- 300 million users

History

- Rails 2.3
- MySQL
- S3

What happened then?

Success!

History

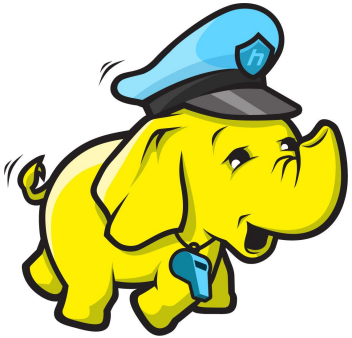

- Rails 2.3
- MySQL
- AWS
- Cassandra
- Hadoop
- SolR
- RabbitMQ

https://developers.soundcloud.com/blog/evolution-of-soundclouds-architecture

Still not enough

- More servers
- Add caching layer
- Defer long running tasks to workers

Still not enough

- Optimize database schema
- Introduce read slaves
- Dedicated databases for some models


Hello
my name is
Monolith

Major pain points

- Testing, building and deploying
- Dependency hell
- “*I’d rather not touch this*”

Rails problems I - No service layer

```
<% Category.all.each do |cat| %>
  <li><%= cat.name %></li>
<% end %>
```

Rails problems I - No service layer

```
<% Category.all.each do |cat| %>
  <li><%= cat.name %></li>
<% end %>
```

⇒ *Tight coupling with storage layer!*

Rails problems II - Active Record Magic

```ruby
class User < ActiveRecord::Base

  validates_length_of :username, :within => 2..64

  before_save :encrypt_password, :accept_terms_of_use

  has_many :comments, :dependent => :destroy

  # ...
end
```

Rails problems II - Active Record Magic

```ruby
class User < ActiveRecord::Base

  validates_length_of :username, :within => 2..64

  before_save :encrypt_password, :accept_terms_of_use

  has_many :comments, :dependent => :destroy

  # ...
end
```

⇒ *Easy to write, hard to maintain*

Ruby Problems

- GIL
- Native extensions
- Dependency management can be painful


Fire
exit

Extract features as Services

- Less painful to maintain
- Easy to replace
- Fun to build

An Example: Messages

- 200 million messages
- MySQL database on shared host
- Features:
  - embedded sounds
  - email notifications
  - spam detection

Migration Requirements

- New schema to support upcoming features
- Dedicated database
- Zero downtime

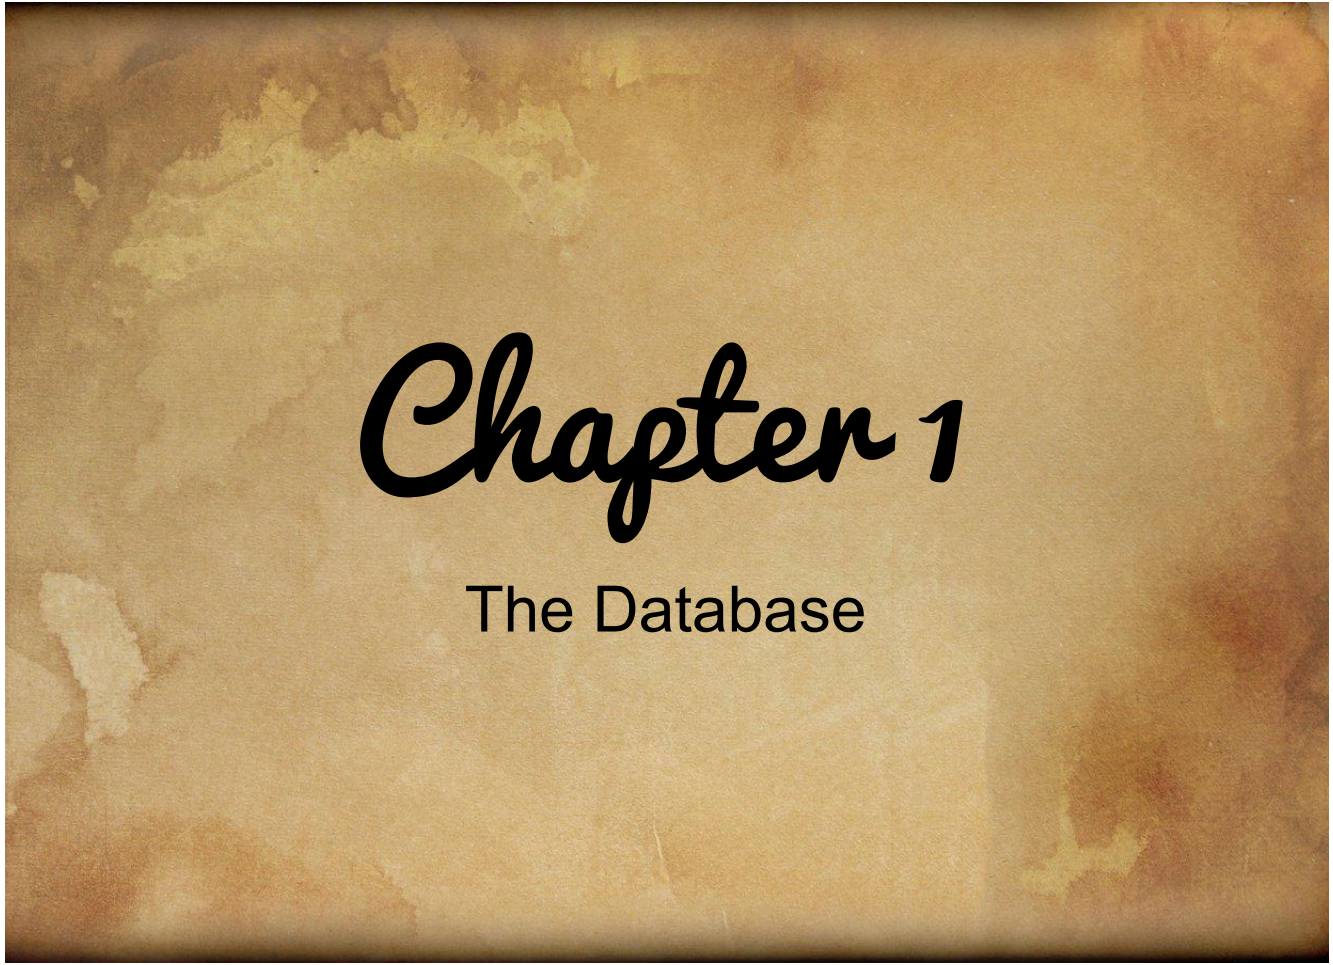
Chapter 1
The Database

Migrating the Schema

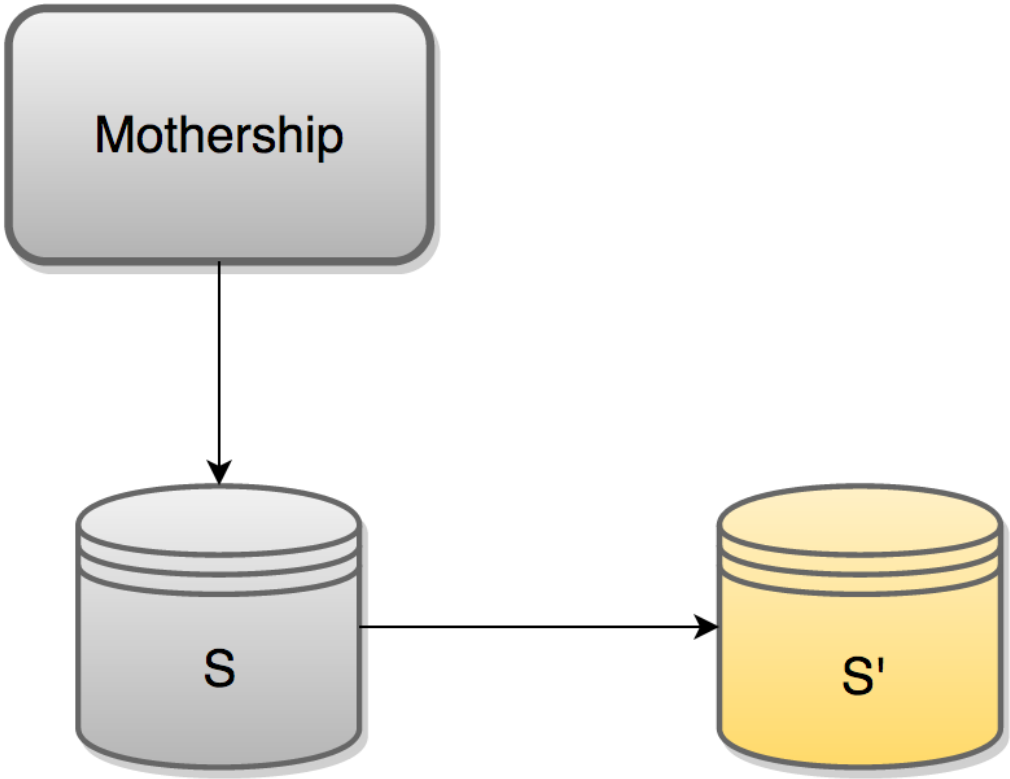

Migrating the Schema

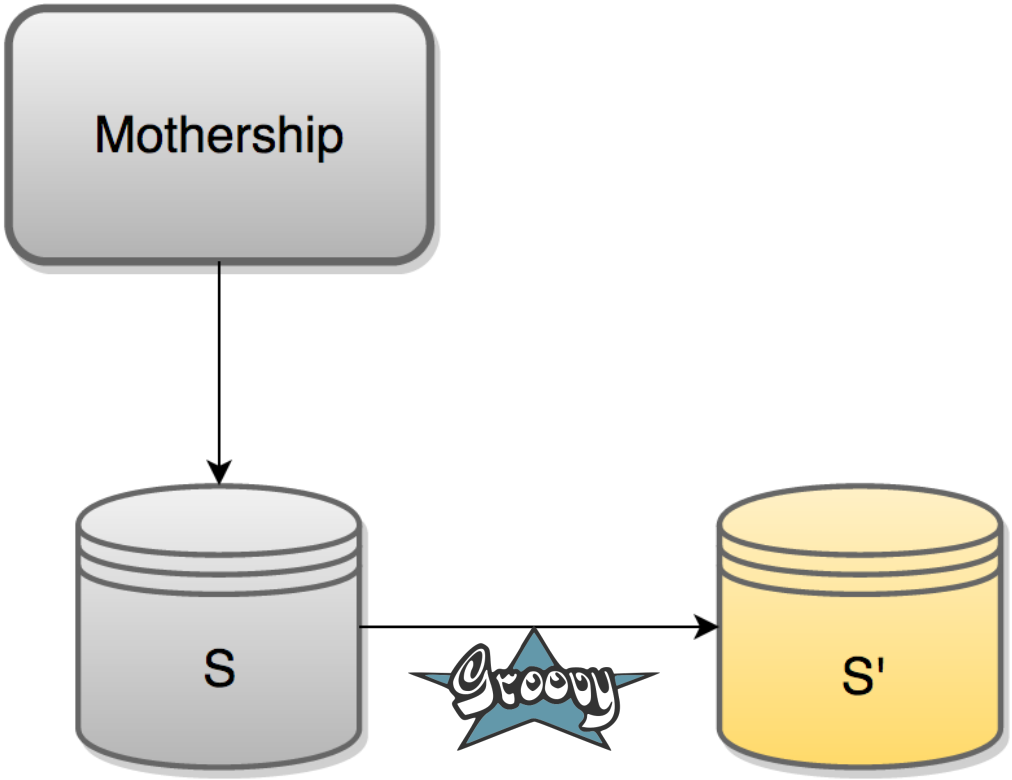

Migrating the Schema

```
#!/usr/bin/env groovy

@Grapes([
  @Grab(group='org.yaml', module='snakeyaml', version='1.12'),
  @Grab(group='mysql', module='mysql-connector-java', version='5.1.24')
])
import groovy.sql.Sql
import org.yaml.snakeyaml.Yaml
```

Convenient Dependency Management with `@Grapes`

Strategy

- Import all messages
- Setup cron job to get new messages
- Listen to events for updates

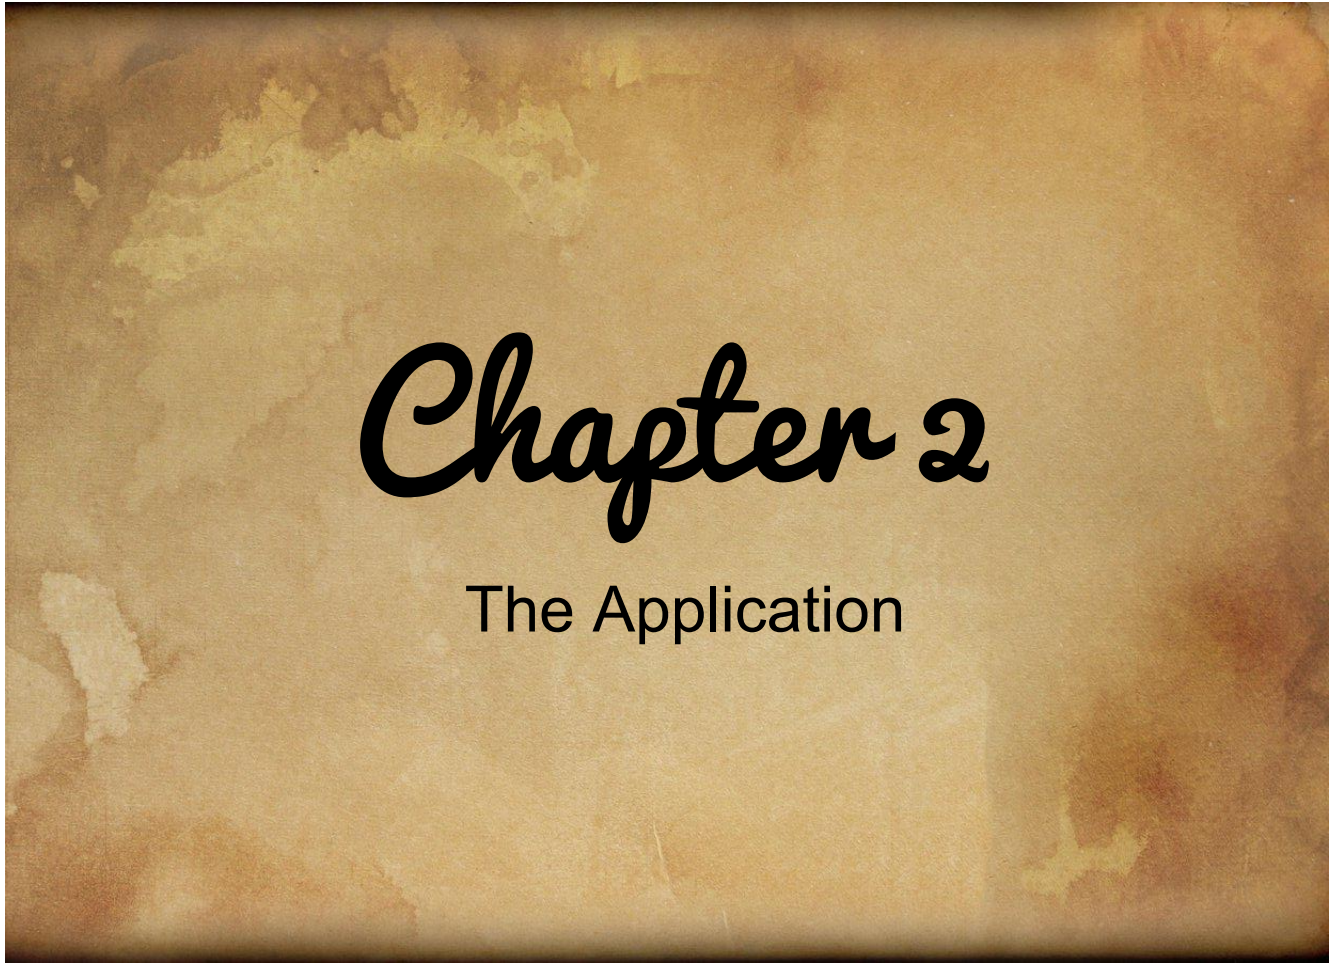
Chapter 2
The Application

Creating a new service

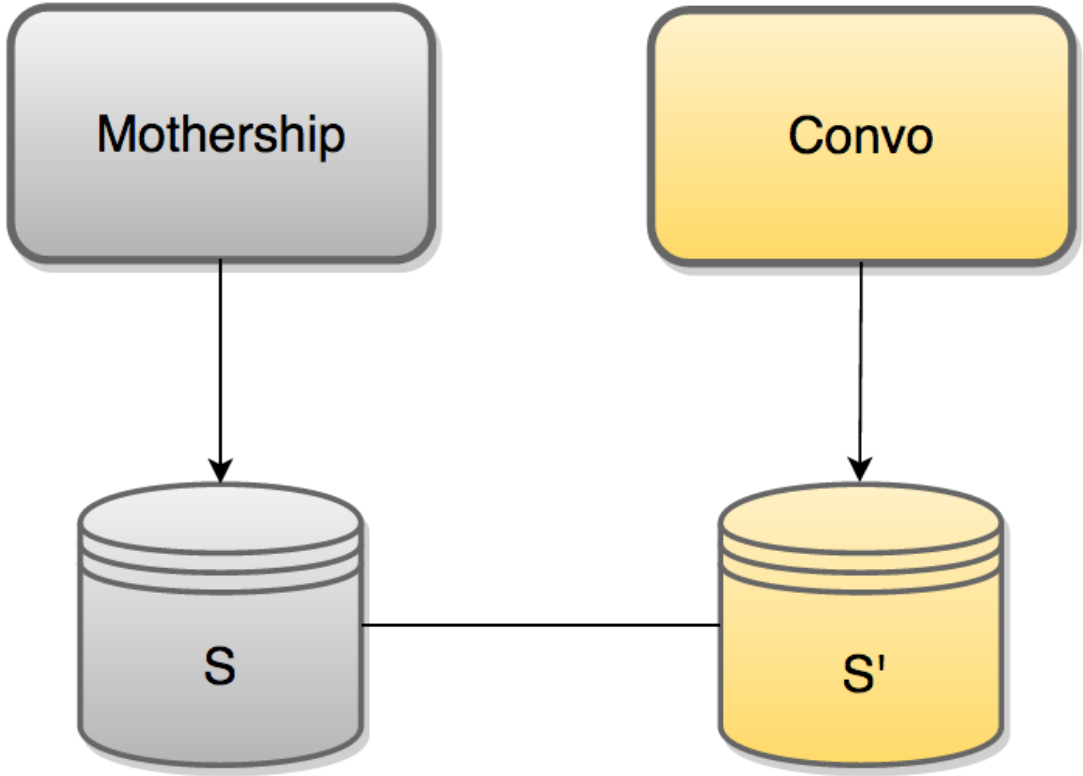

Convo

- Scala
- Twitter Finagle
- Scalatra Framework

Convo architecture

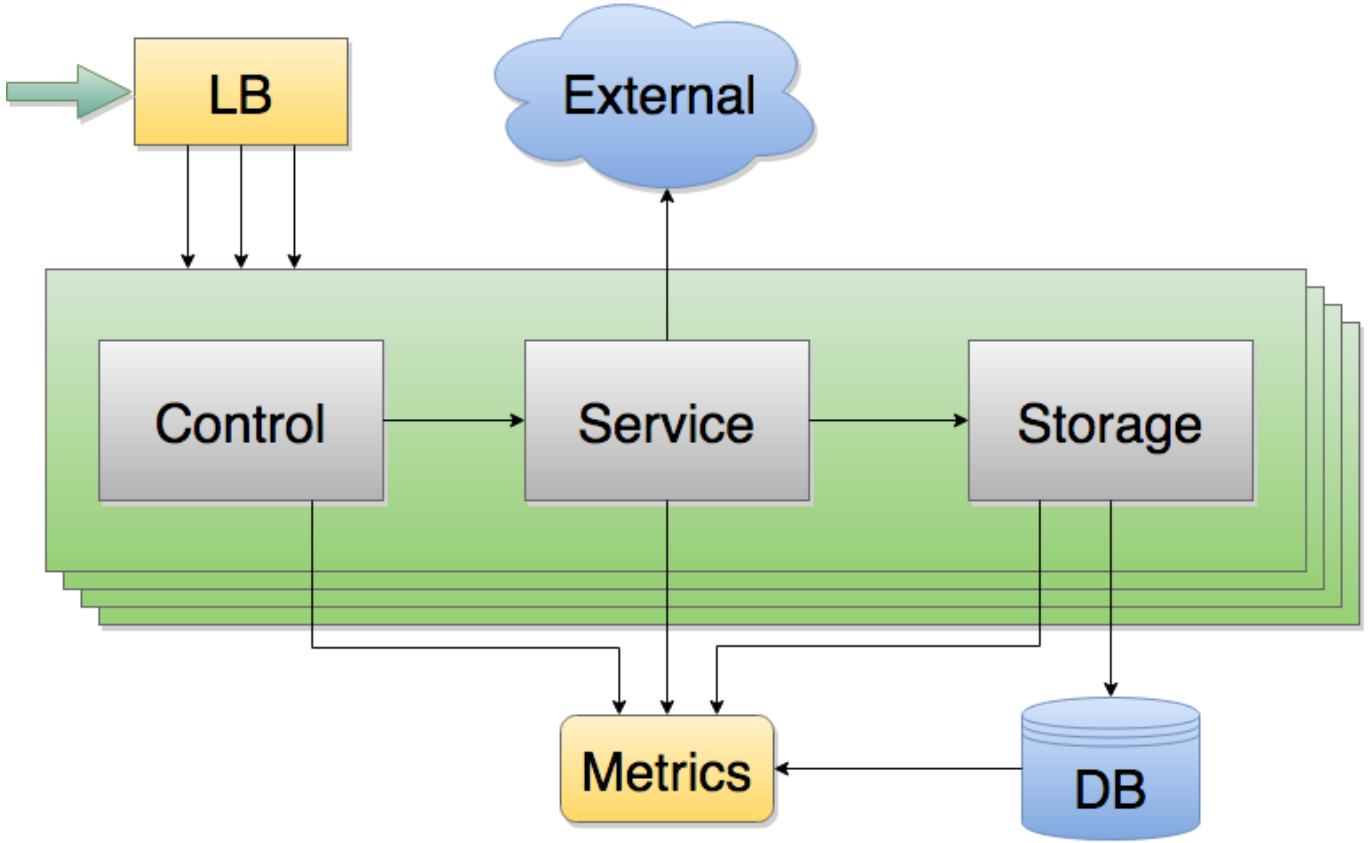

Scala

- Functional
- OOP
- Static but inferred typing

Scala Joy – Options I

```
val opt: Option[String] = params.get("id")

val id: Int = opt.map(id => id.toInt).getOrElse(10)
```

Good bye `NullPointerException`

Scala Joy – Options II

```
for {
  id    <- params.get("id")
  user  <- users.lookup(id)
  count <- counts.forUser(user)
} yield count
```

Safe chaining with `for` comprehensions

Scala Joy – Pattern Matching

```
Urn("soundcloud:users:20") match {
  case Urn(_, "tracks", _)       => None,
  case Urn(_, "messages", "20") => None,
  case Urn(_, "users", id)       => Some(id)
}
```

Expressive code with decomposition

Scala Joy – Functional Goodness

```
delete("/playlist/:urn/likes")(destroy)

def destroy(request: Request) =
  write(request, 200)(repo.deleteLike)

def write
  (request: Request, statusCode: Int)
  (f: (UserSession, Urn) => Future[Like]) = {
  // ...
}
```

Function arguments and references

Futures!

Finagle

- Twitter rpc library on top of Netty
- Support for multiple protocols
- Future composition

Futures

```
class Future[A] {

  def get(): A

  def map[B](f: A => B): Future[B]

  def flatMap[B](f: A => Future[B]]): Future[B]

  def onSuccess(f : A => Unit): Future[A]
}
```

Instance API (excerpt)

Futures

```
object Future {

  def value[A](a: A): Future[A]

  def exception[A](e: Throwable): Future[A]

  def collect[A](fs : Seq[Future[A]]): Future[Seq[A]]
}
```

Object API (excerpt)

Futures - Examples

```
service.getUsers().flatMap { users =>
  service.tracksFor(users).flatMap { tracks =>
    asJson(tracks)
  }
}.onSuccess(json => log(s"found $json"))
```

Multiple transformations - The ugly way

Futures - Example

```
val response = for {
  users  <- service.getUsers()
  tracks <- service.tracksFor(users)
  json   <- asJson(tracks)
} yield json

response.onSuccess(json => log(s"found $json"))
```

Multiple transformations - The nice way

Scala Problems

- Implicit conversions
- Binary compatibility of libraries
- Tooling still not perfect

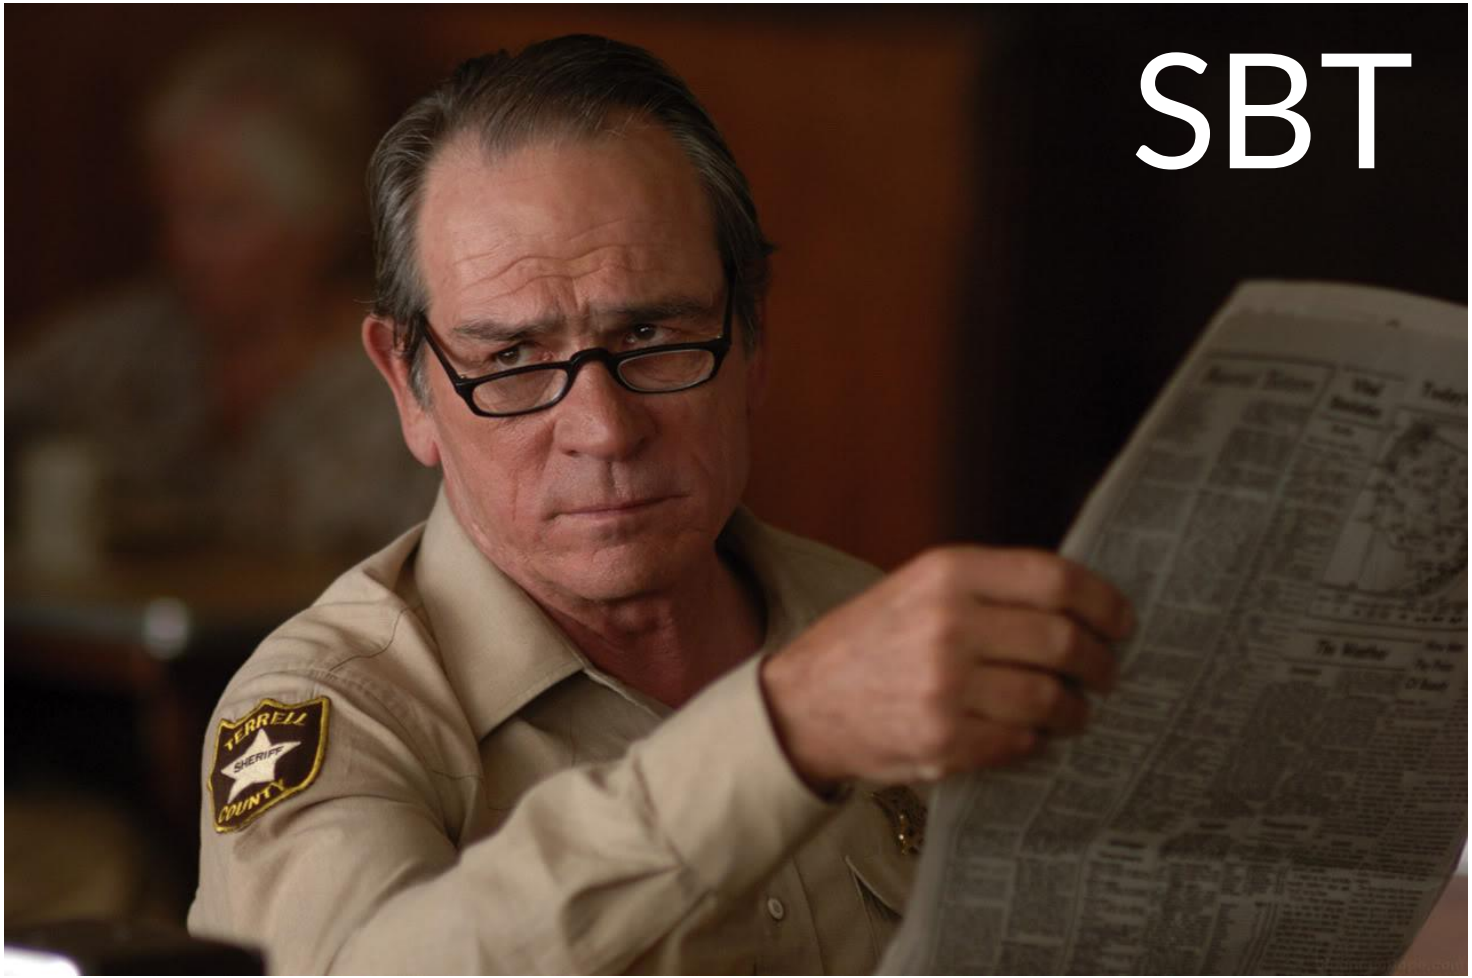
SBT
TERRELL
SHERIFF
COUNTY

IntelliJ

- Code inspection
- Debugging
- SBT support

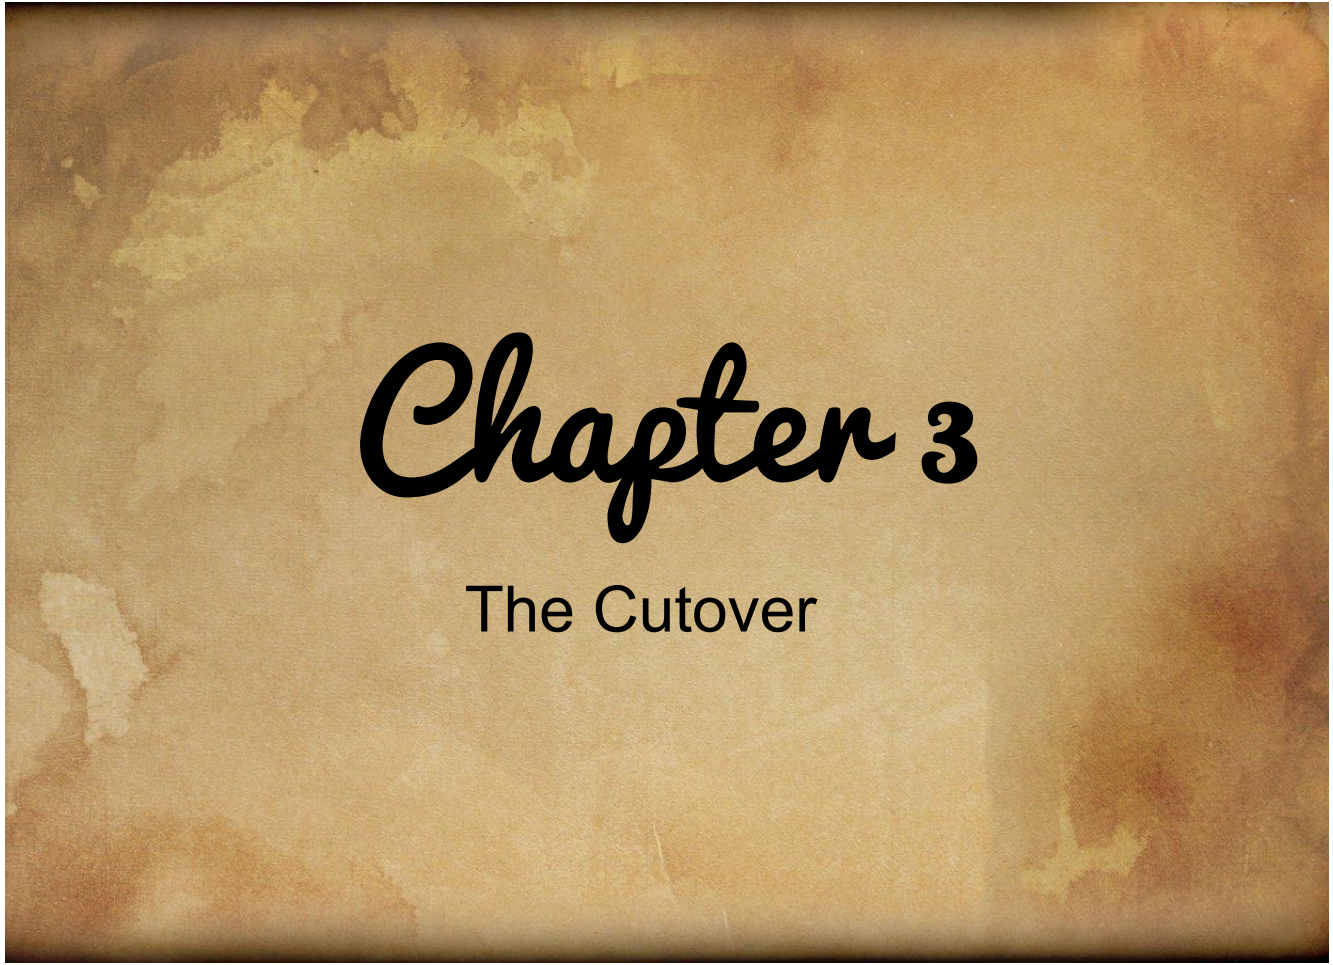
Chapter 3
The Cutover

Integrate Service

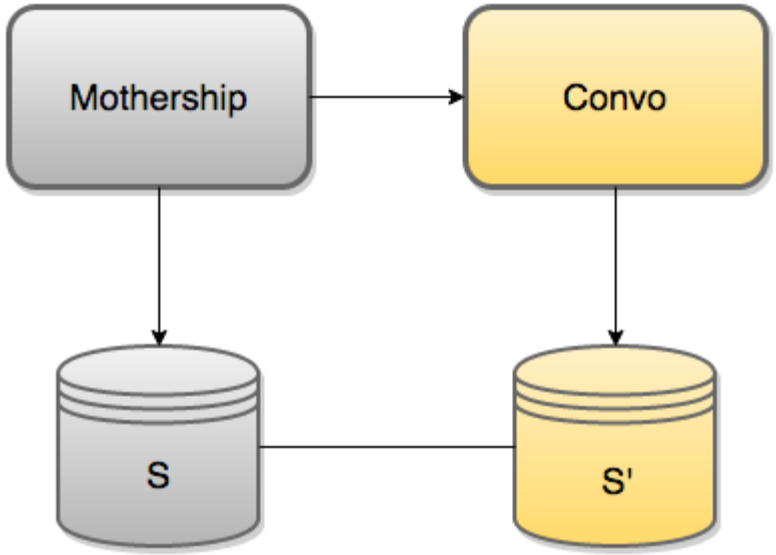

Integration Risks

- Service failure
- Data loss after rolling back
- Data loss caused by stale clients

Integration Risks

- Service failure → *load testing, A/B testing*
- Data loss after rolling back
- Data loss caused by stale clients

Integration Risks

- Service failure → *load testing, A/B testing*
- Data loss after rolling back → *prepare scripts, practice*
- Data loss caused by stale clients

Integration Risks

- Service failure → *load testing, A/B testing*
- Data loss after rolling back → *prepare scripts, practice*
- Data loss caused by stale clients → *keep migration running*

Enable Feature

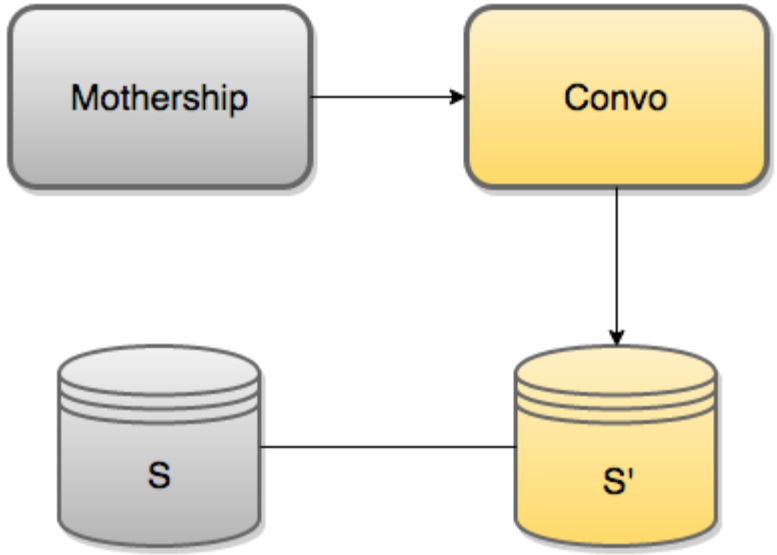

Retire Old Database

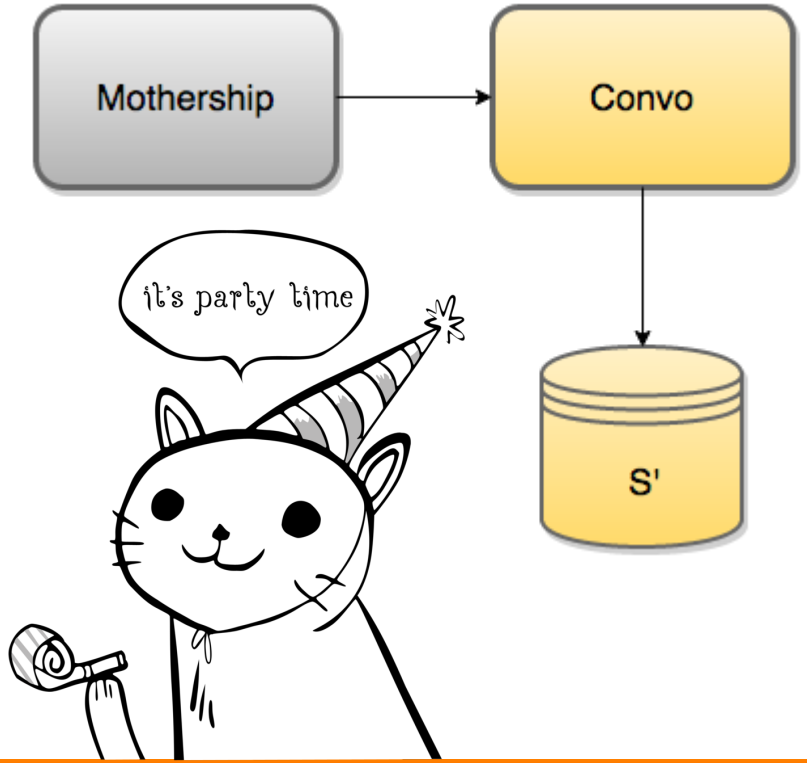

Convo

- 500 million requests per day
- 1000 qps during peak time
- 5 instances

Microservice Problems

- Event bus dependency
- Maintenance overhead
- Distributed tracing

Microservices

→ *Not a silver bullet*


That's all Folks!

Questions?

Images

- Slide 4,7 - Rails Logo http://en.wikipedia.org/wiki/File:Ruby_on_Rails.svg
- Slide 6,51 - Party Cat http://ghostexist.deviantart.com/art/Party-Cat-logo-287986071
- Silde 7 - MySQL Logo http://blogwifi.fr/?p=9990
- Slide 7 - Hadoop Cop https://svn.apache.org/repos/asf/hadoop/logos/out_rgb/hadoop-security-logo.jpg
- Slide 10 - Hello, My Name Is: http://commons.wikimedia.org/wiki/File:Hello_my_name_is_sticker.svg
- Slide 14 - Sad Panda: http://www.whatsupyasieve.com/2012/09/17/lockout-blues/sad-panda-2/
- Slide 16 - Exit Sign: http://logo-kid.com/emergency-exit-sign-left.htm
- Slide 22 - Groovy Logo: http://groovy.codehaus.org/images/groovy-logo-medium.png
- Slide 20, 25, 44 - Book Page: http://daviddiazolivares.deviantart.com/art/Old-Book-Page-345869530
- Slide 34 - Back to the future: http://i.huffpost.com/gen/1369403/thumbs/o-BACK-TO-THE-FUTURE-facebook.jpg
- Slide 42 - Tommy Lee Jones: http://persephonemagazine.com/2014/04/friday-news-bites-airline-pranks-gabriel-garcia-marquez-pulitzers-more/film-title-no-country-for-old-men/
- Slide 55 - That's all folks: http://www.hd2wallpapers.com/view/thats_all_folks-1280x800.php