

Database Scaling at Mozilla




Sheeri Cabral

Senior DB Admin/Architect, Mozilla
@sheeri www.sheeri.com




High Availability vs Scaling


High Availability




High Availability




Scaling



Scaling



Scaling



...is not necessarily highly available

The Difference



Availability

- Can I reach it?

Scaling

- Can 50,000 people reach it?

Ways to Scale



Scale up

vs

Scale out

Open Source Scaling



Multiple slaves

Open Source Scaling



Multiple slaves

Partitioning

Open Source Scaling



Multiple slaves

Partitioning

Sharding

Open Source Scaling



Multiple slaves

Partitioning

Sharding

MySQL Cluster

Open Source Scaling



Multiple slaves

Partitioning

Sharding

MySQL Cluster

Galera Cluster

Open Source Scaling



Multiple slaves

Partitioning

Sharding

MySQL Cluster

Galera Cluster

...some of these are also high availability solutions

Our Options



Multiple slaves

Partitioning

Sharding

MySQL Cluster

Galera Cluster

Our Options



Multiple slaves

Partitioning – needed HA

Sharding

MySQL Cluster

Galera Cluster

Our Options



Multiple slaves

Partitioning – needed HA

Sharding – could not re-architect app

MySQL Cluster

Galera Cluster

Our Options



Multiple slaves

~~Partitioning~~ – needed HA

~~Sharding~~ – could not re-architect app

MySQL Cluster

~~Galera Cluster~~ – GA Apr 2012

Our Options



~~Multiple slaves~~ – needed write scaling


~~Partitioning~~ – needed HA

~~Sharding~~ – could not re-architect app

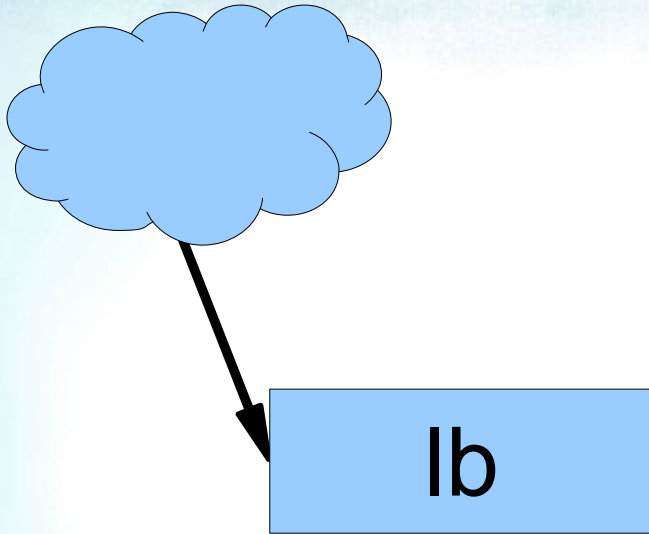
MySQL Cluster

~~Galera Cluster~~ – GA Apr 2012


Current Setup




Current Setup



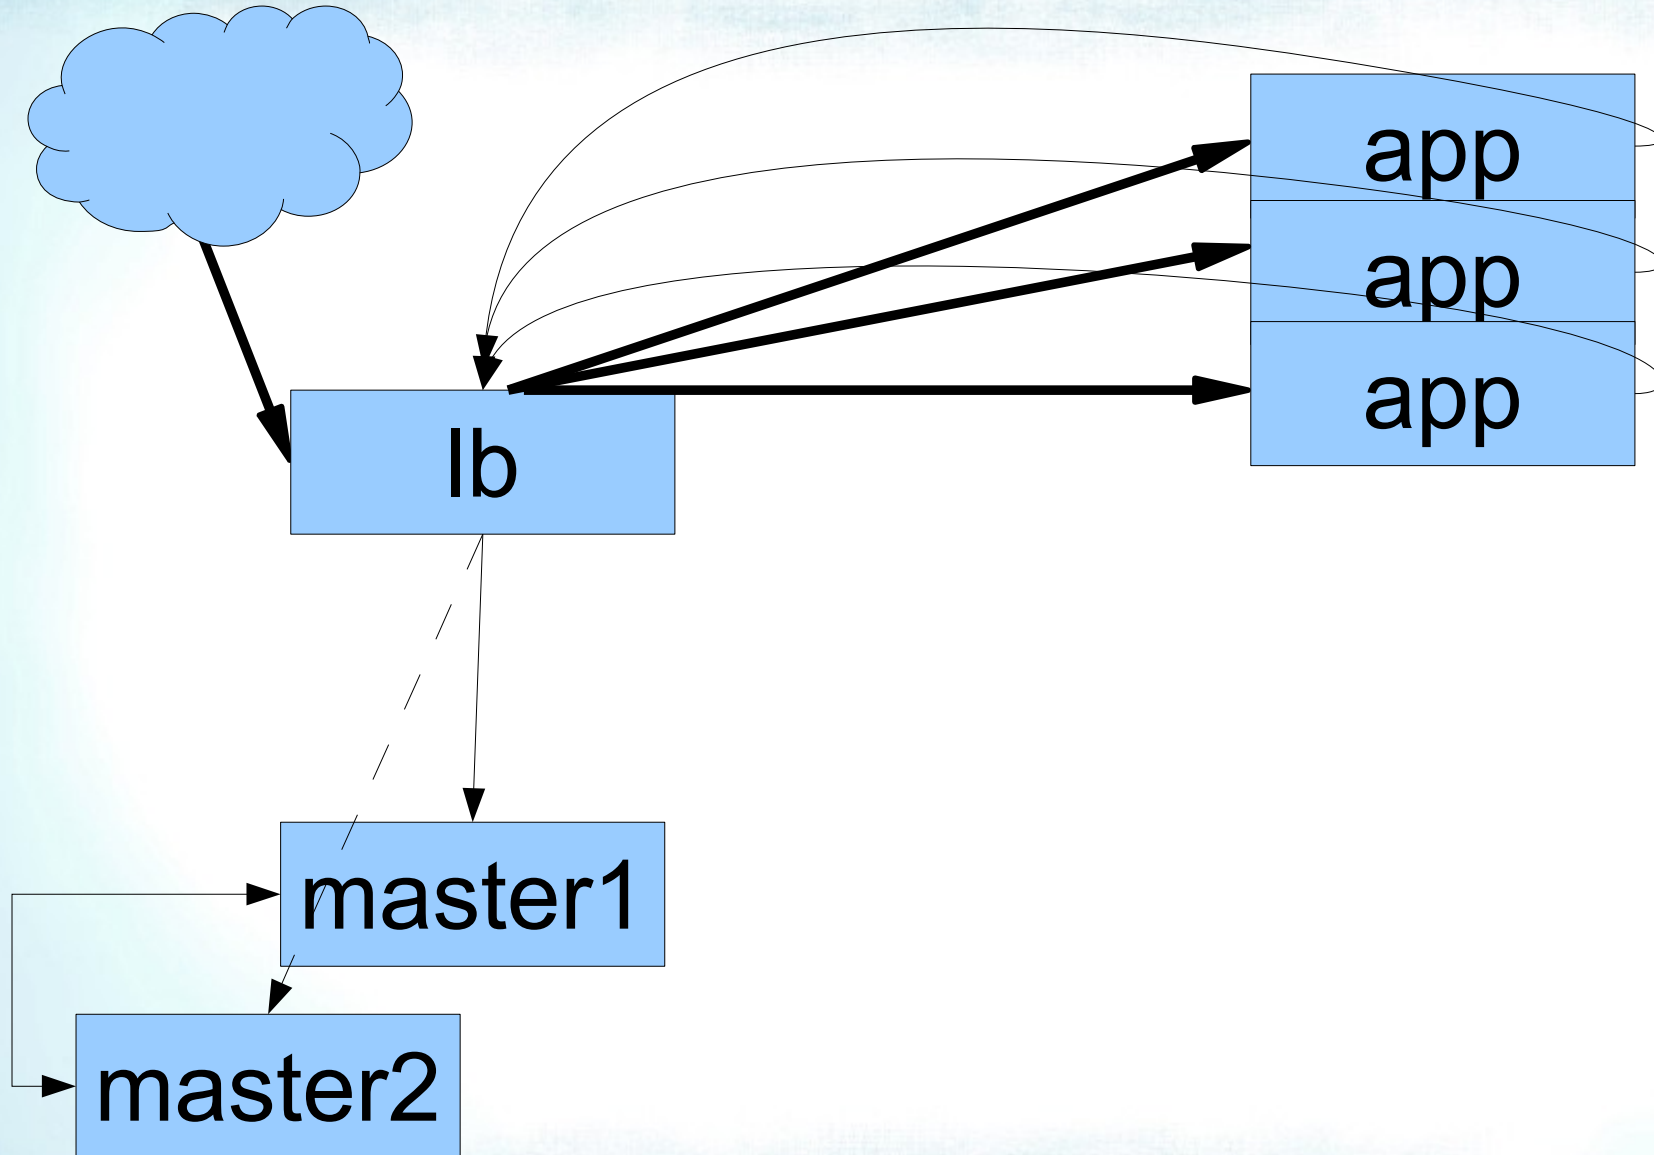
Current Setup



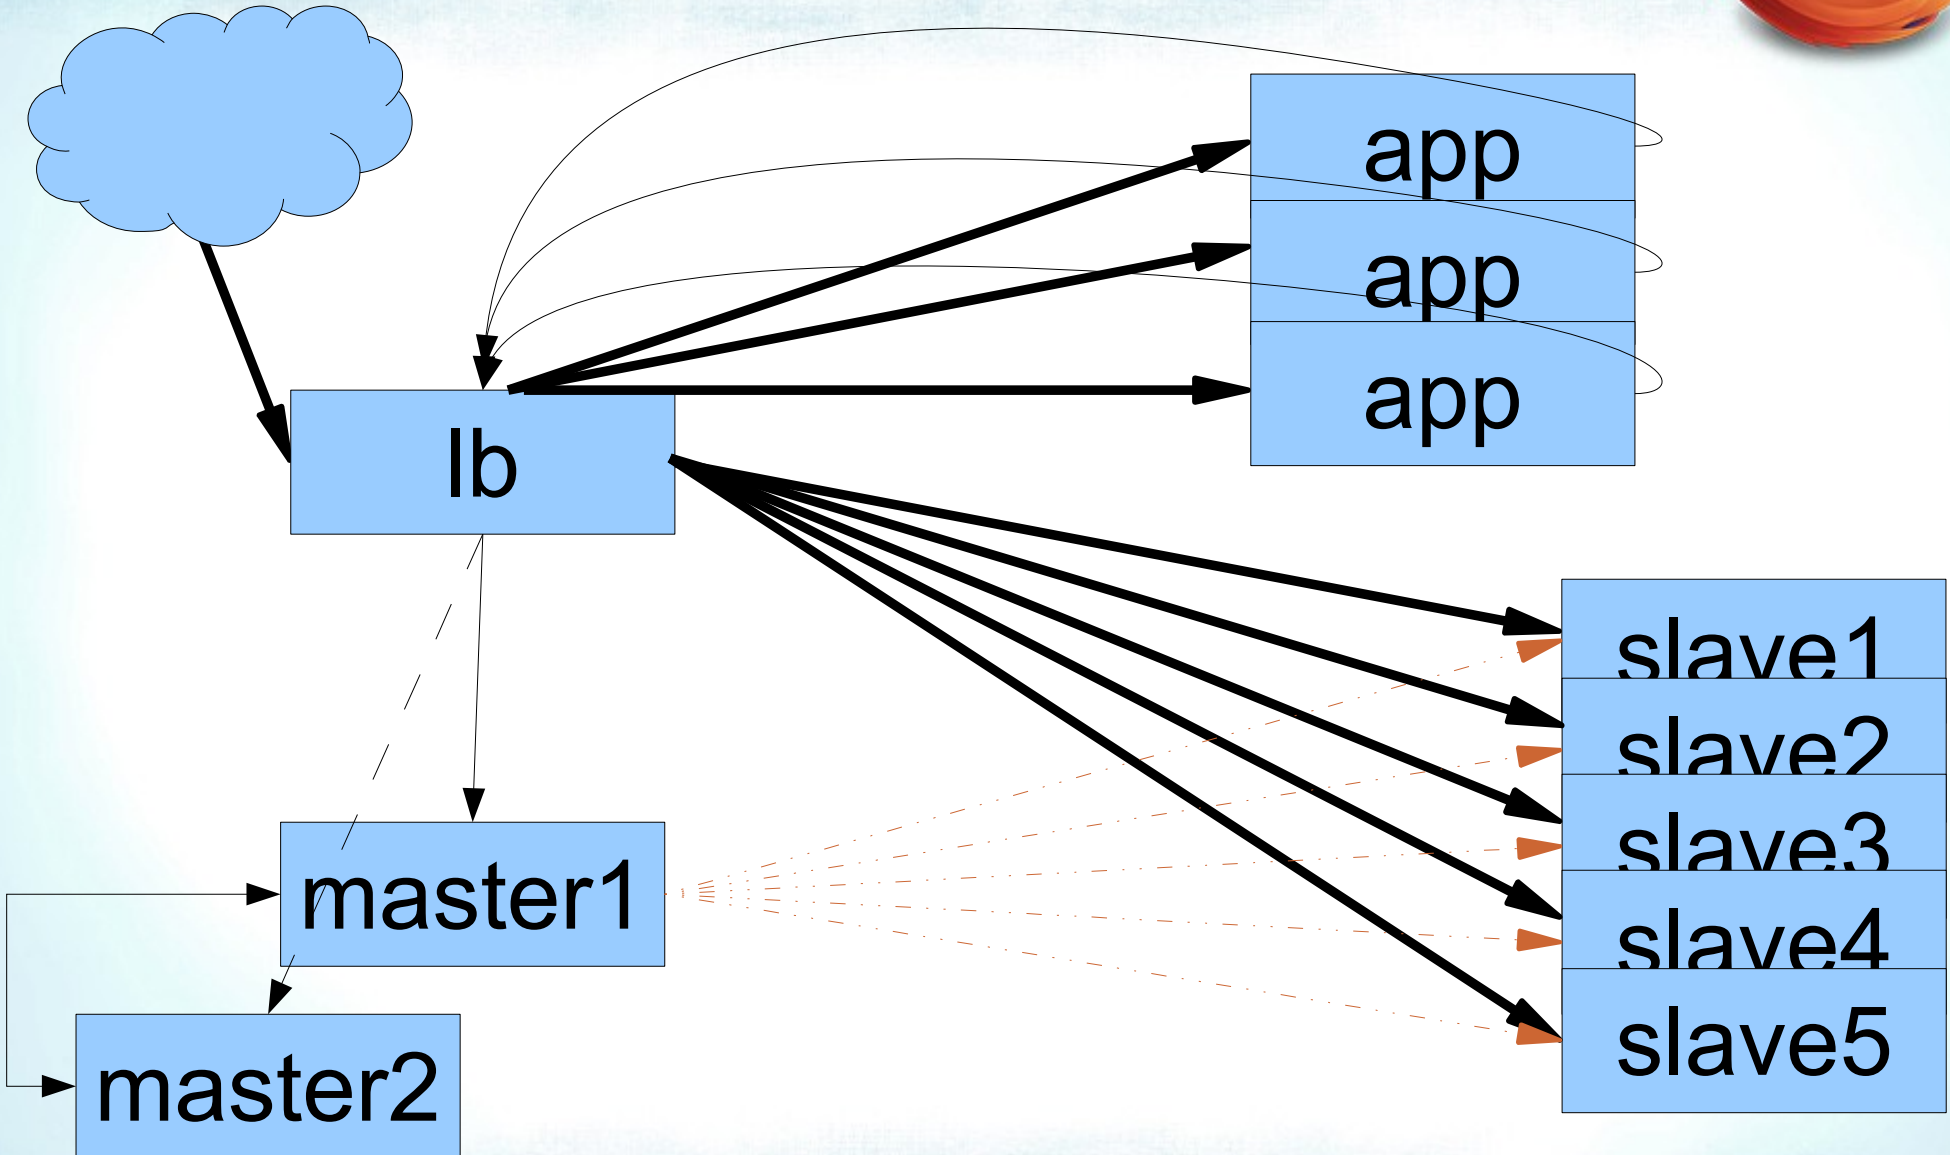
Current Setup



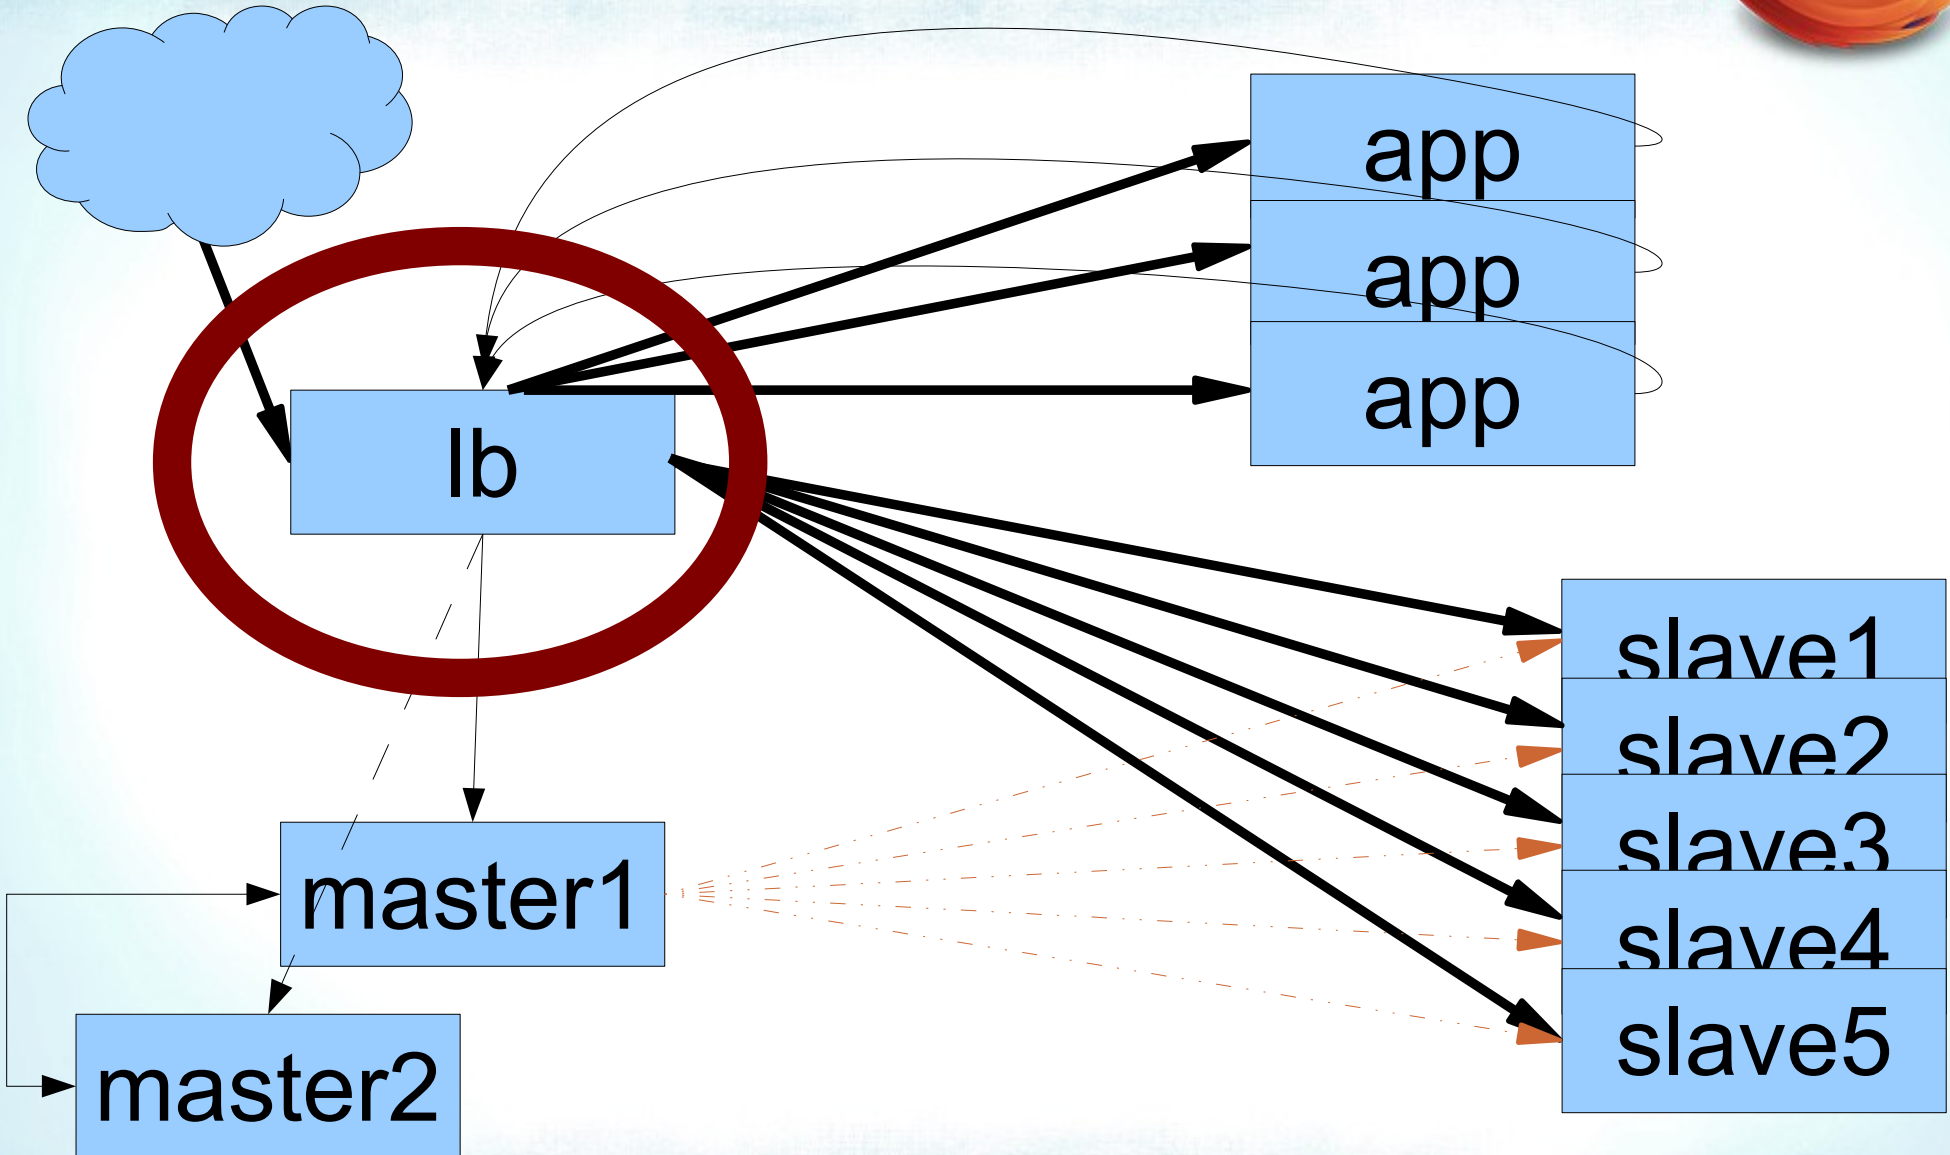
Current Setup



Current Setup



Current Setup



What We Really Needed



...a better load balancer

Our Options



~~Multiple slaves~~ – needed write scaling

~~Partitioning~~ – needed HA

~~Sharding~~ – could not re-architect app

MySQL Cluster

~~Galera Cluster~~ – GA Apr 2012

Our Options



~~Multiple slaves~~ – needed write scaling

~~Partitioning~~ – needed HA

~~Sharding~~ – could not re-architect app

~~MySQL Cluster~~ – not a load-balancing solution

~~Galera Cluster~~ – GA Apr 2012

What We Really Needed



...a better load balancer

to work within our existing framework

What We Really Needed



...a better load balancer

to work within our existing framework

flexible

What We Really Needed



...a better load balancer

to work within our existing framework

flexible

geo-sharding

What We Really Needed



...a better load balancer

to work within our existing framework

flexible

~~geo-sharding~~

Commercial Solutions



Scalarc

Scalebase

Tungsten

Commercial Solutions



~~Scalare~~ – not quite ready in Feb 2012

– Much better now!

Scalebase

Tungsten

Commercial Solutions



~~Scalare~~ – not quite ready in Feb 2012

– Much better now!

Scalebase

~~Tungsten~~ – late-comer, then ran out of time

– But does do what we need!

Scalebase



Good load balancing

Flexible

- Avoid the “zero-day load” problem
- Can change architecture easily

Is also highly available (redundancy)

Scalebase

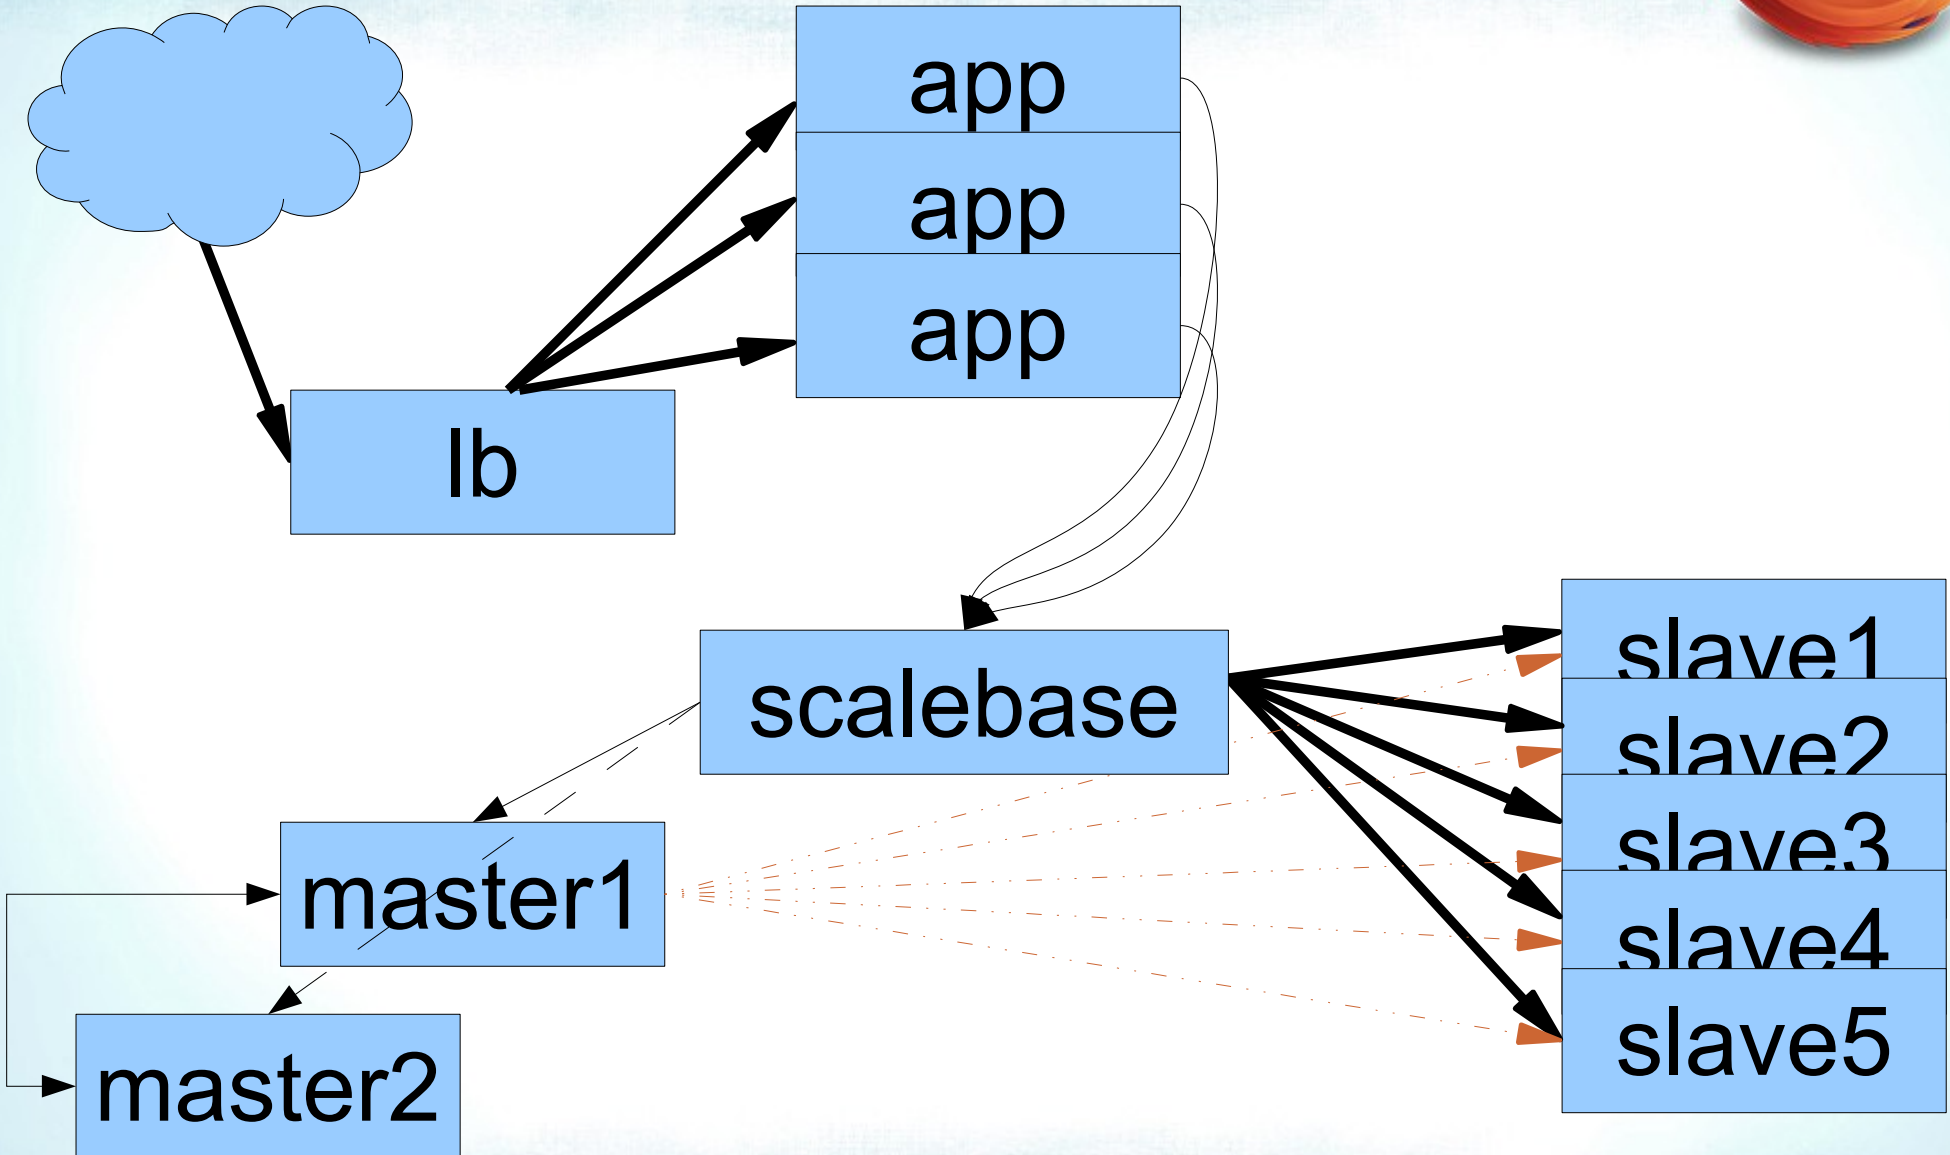


R/W splitting

- Eliminates need for separate VIPs
- Just go to scalebase1.db and scalebase2.db
- Eliminate load balancer dependency

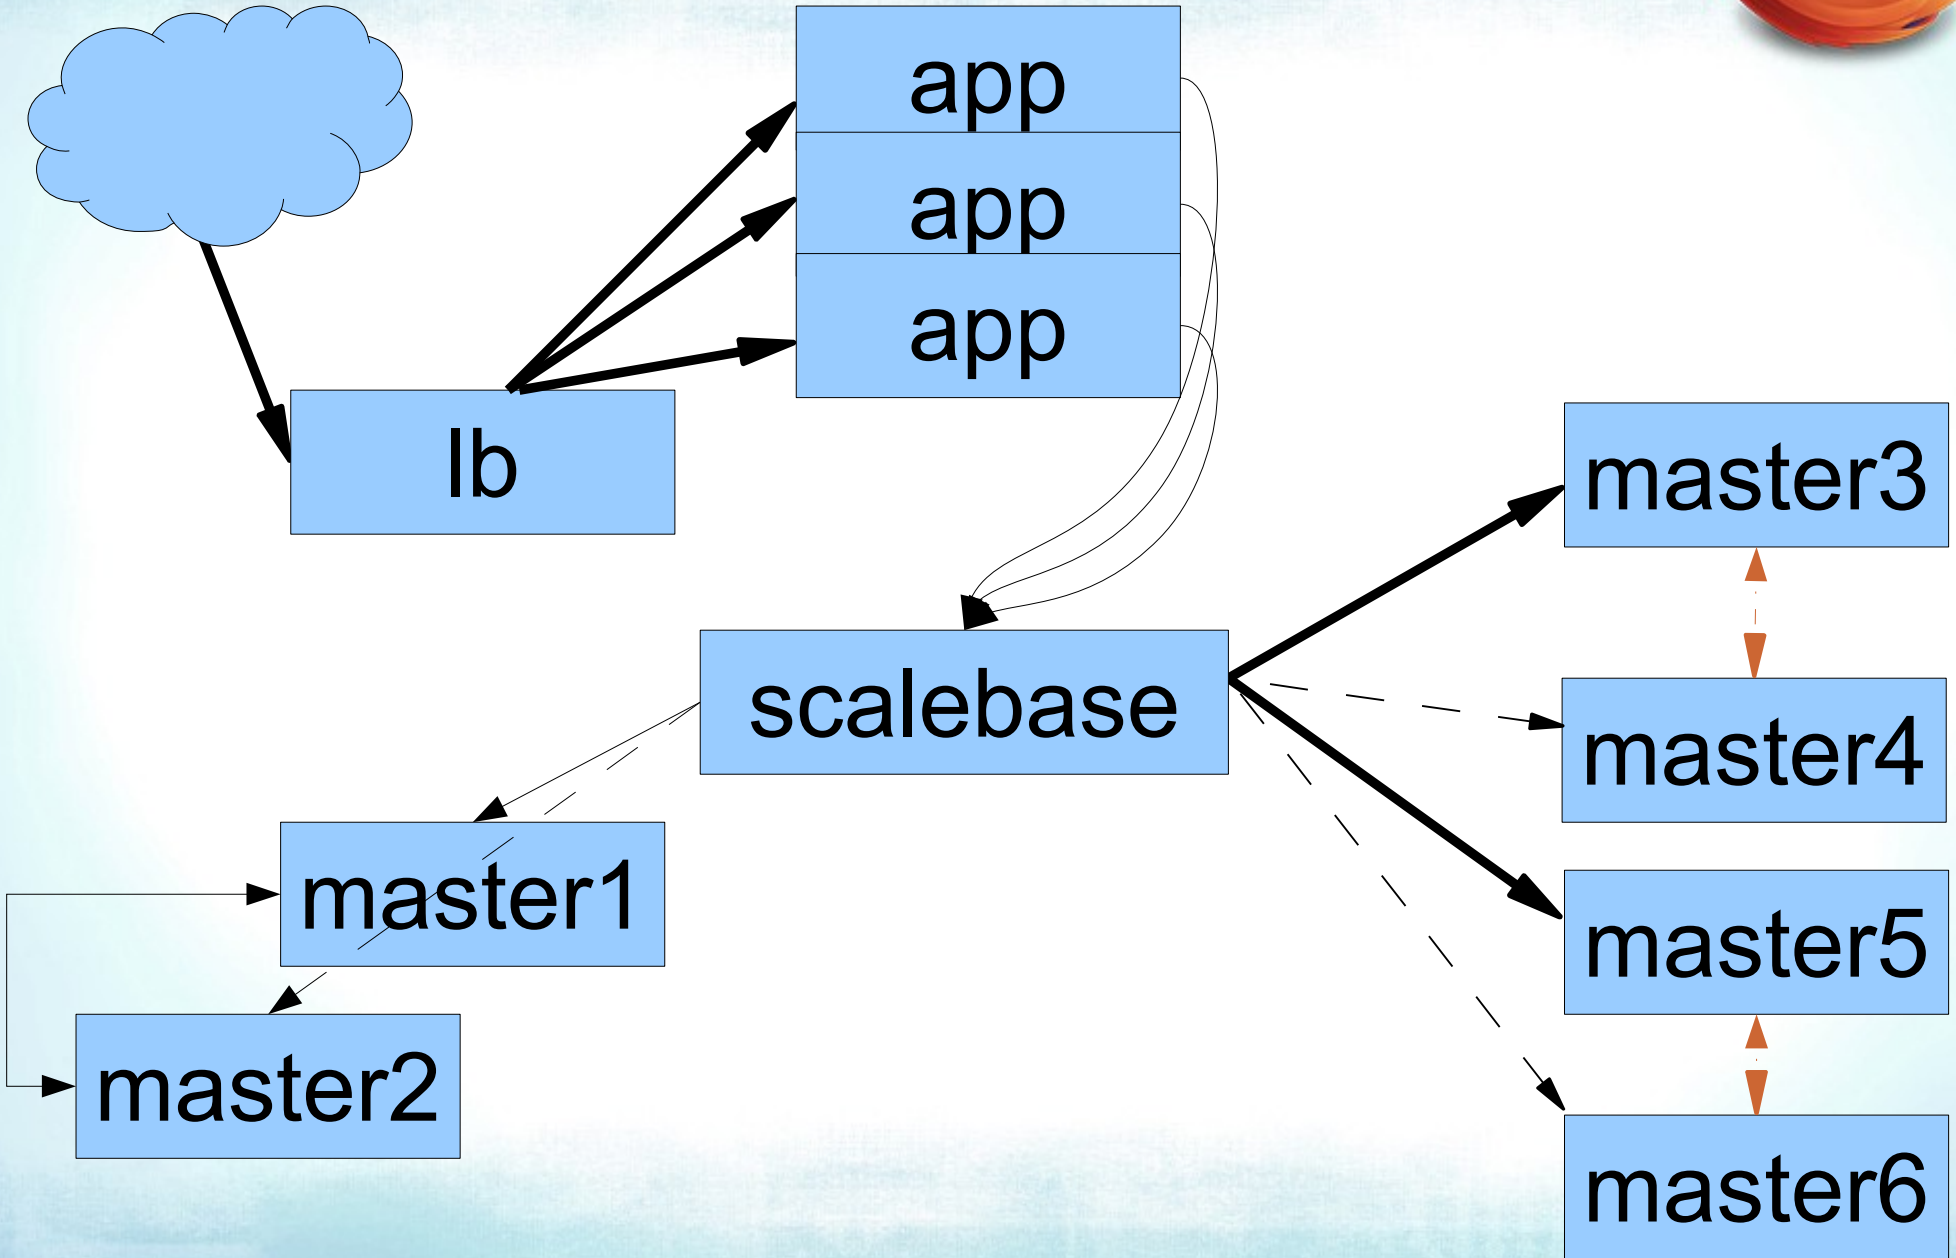


Current Setup





Current Setup



Other Features



Query redirection

Transparent geo-sharding

- You pick the sharding key

Data flow management

- Query redirecting

Other Features



- Eliminates need for separate VIPs
- Just go to scalebase1.db and scalebase2.db
- Eliminate load balancer dependency

Final Note



- It works for us
- Based on our needs
- YMMV

Mozilla is Hiring!



<http://careers.mozilla.org/>

DBA – MySQL, Postgres, NoSQL (some or all)

<http://careers.mozilla.org/en-US/position/oTqLWfwK>

SysAdmin – RHEL, CentOS, Fedora

<http://careers.mozilla.org/en-US/position/o97xWfwt>

US, Canada, UK, France, Spain, Germany, Netherlands,
Sweden, Denmark, Poland, China, Japan, New Zealand


In Mountain View, **Senior Windows/Unix SysAdmin**

<http://careers.mozilla.org/en-US/position/oZmJWfwK>

Want to Learn MySQL?



12 weeks, 1 book. MySQL Marinate!



Soak it in!

www.mysqlmarinate.com

Questions? Comments?




OurSQL podcast

- www.oursql.com

MySQL Administrator's Bible

- tinyurl.com/mysqlbible



bit.ly/ghackmysql

kimtag.com/mysql

planet.mysql.com