



# LinkedIn's Active/Active Evolution


Erran Berger

Head of Engineering, Content Products, LinkedIn

[www.linkedin.com/in/erranberger](http://www.linkedin.com/in/erranberger)


@eberger45





Time was  
running out





Time was  
running out



Multicolor from the get-go



Who is this Erran Berger  
guy anyway?



What were the circumstances  
of our decision?



Why should you care  
about any of this  
stuff?



# Multicolor From the Get-Go





# Multicolor From the Get-Go

1

Stream  
Processing

2

Personal Data

3

Cache Invalidation  
& Replication  
Conflicts



1

Stream  
processing



8 billion

---

kafka messages

175 TB

---

of data



The great thing about kafka is  
that its data is replicated like  
everything else



WORST

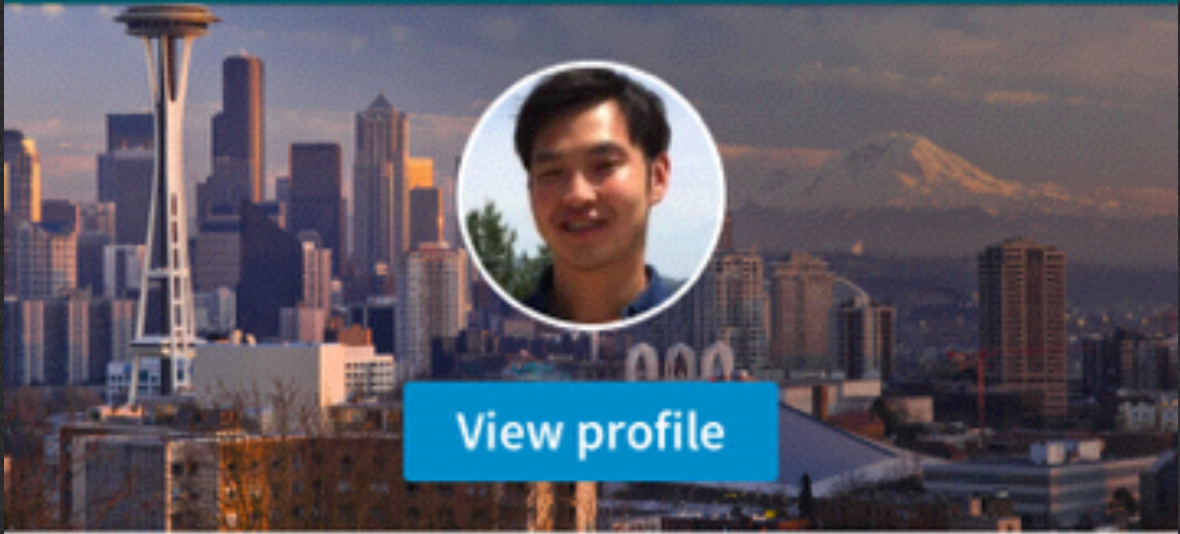
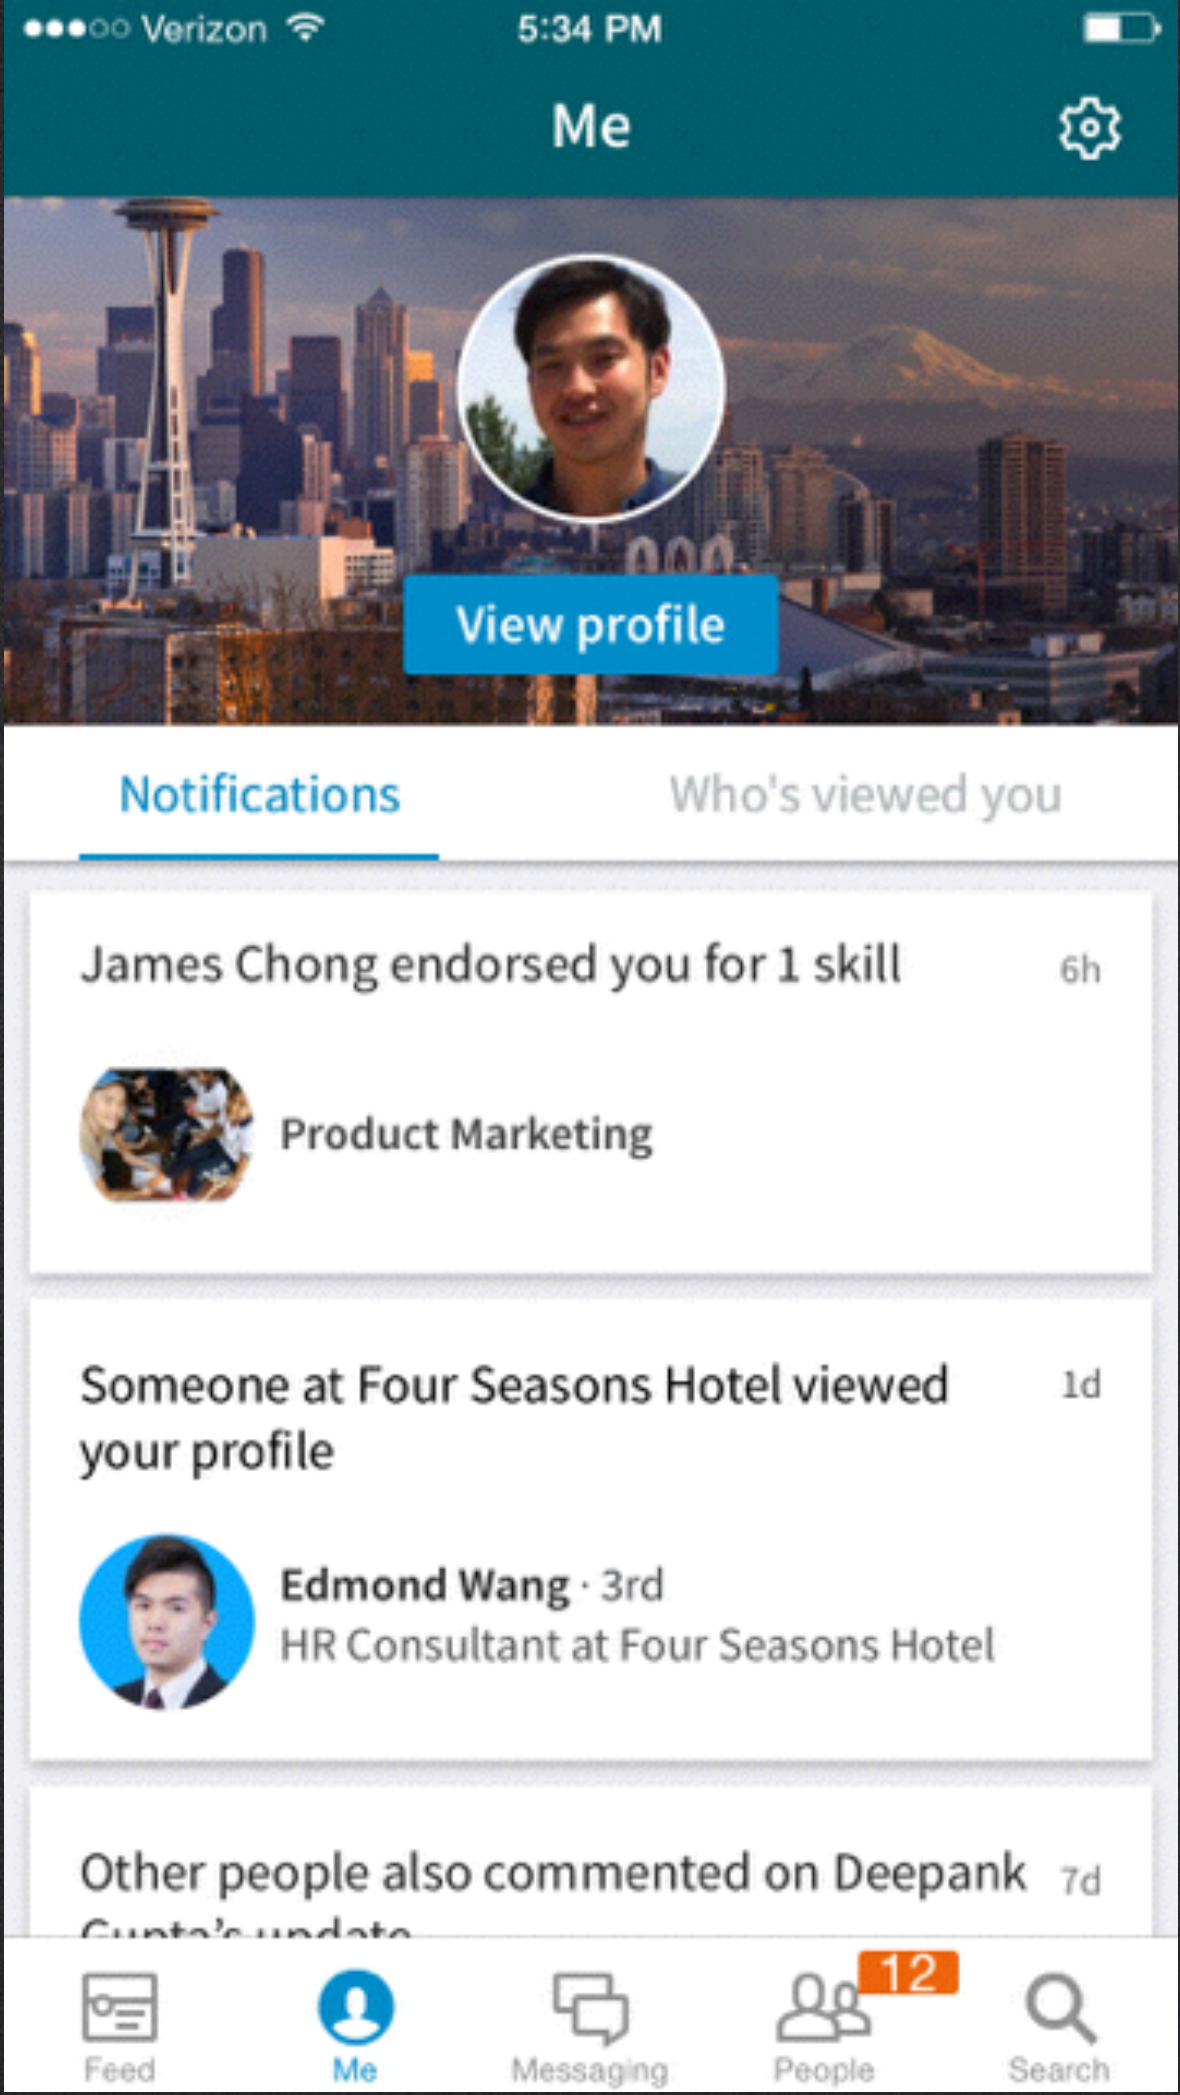
The ~~great~~ thing about kafka is  
that its data is replicated like  
everything else



Only Process Once







Notifications Who's viewed you

James Chong endorsed you for 1 skill 6h



Someone at Four Seasons Hotel viewed your profile 1d

Edmond Wang · 3rd  
HR Consultant at Four Seasons Hotel

Other people also commented on Deepank 7d  
Cun... update

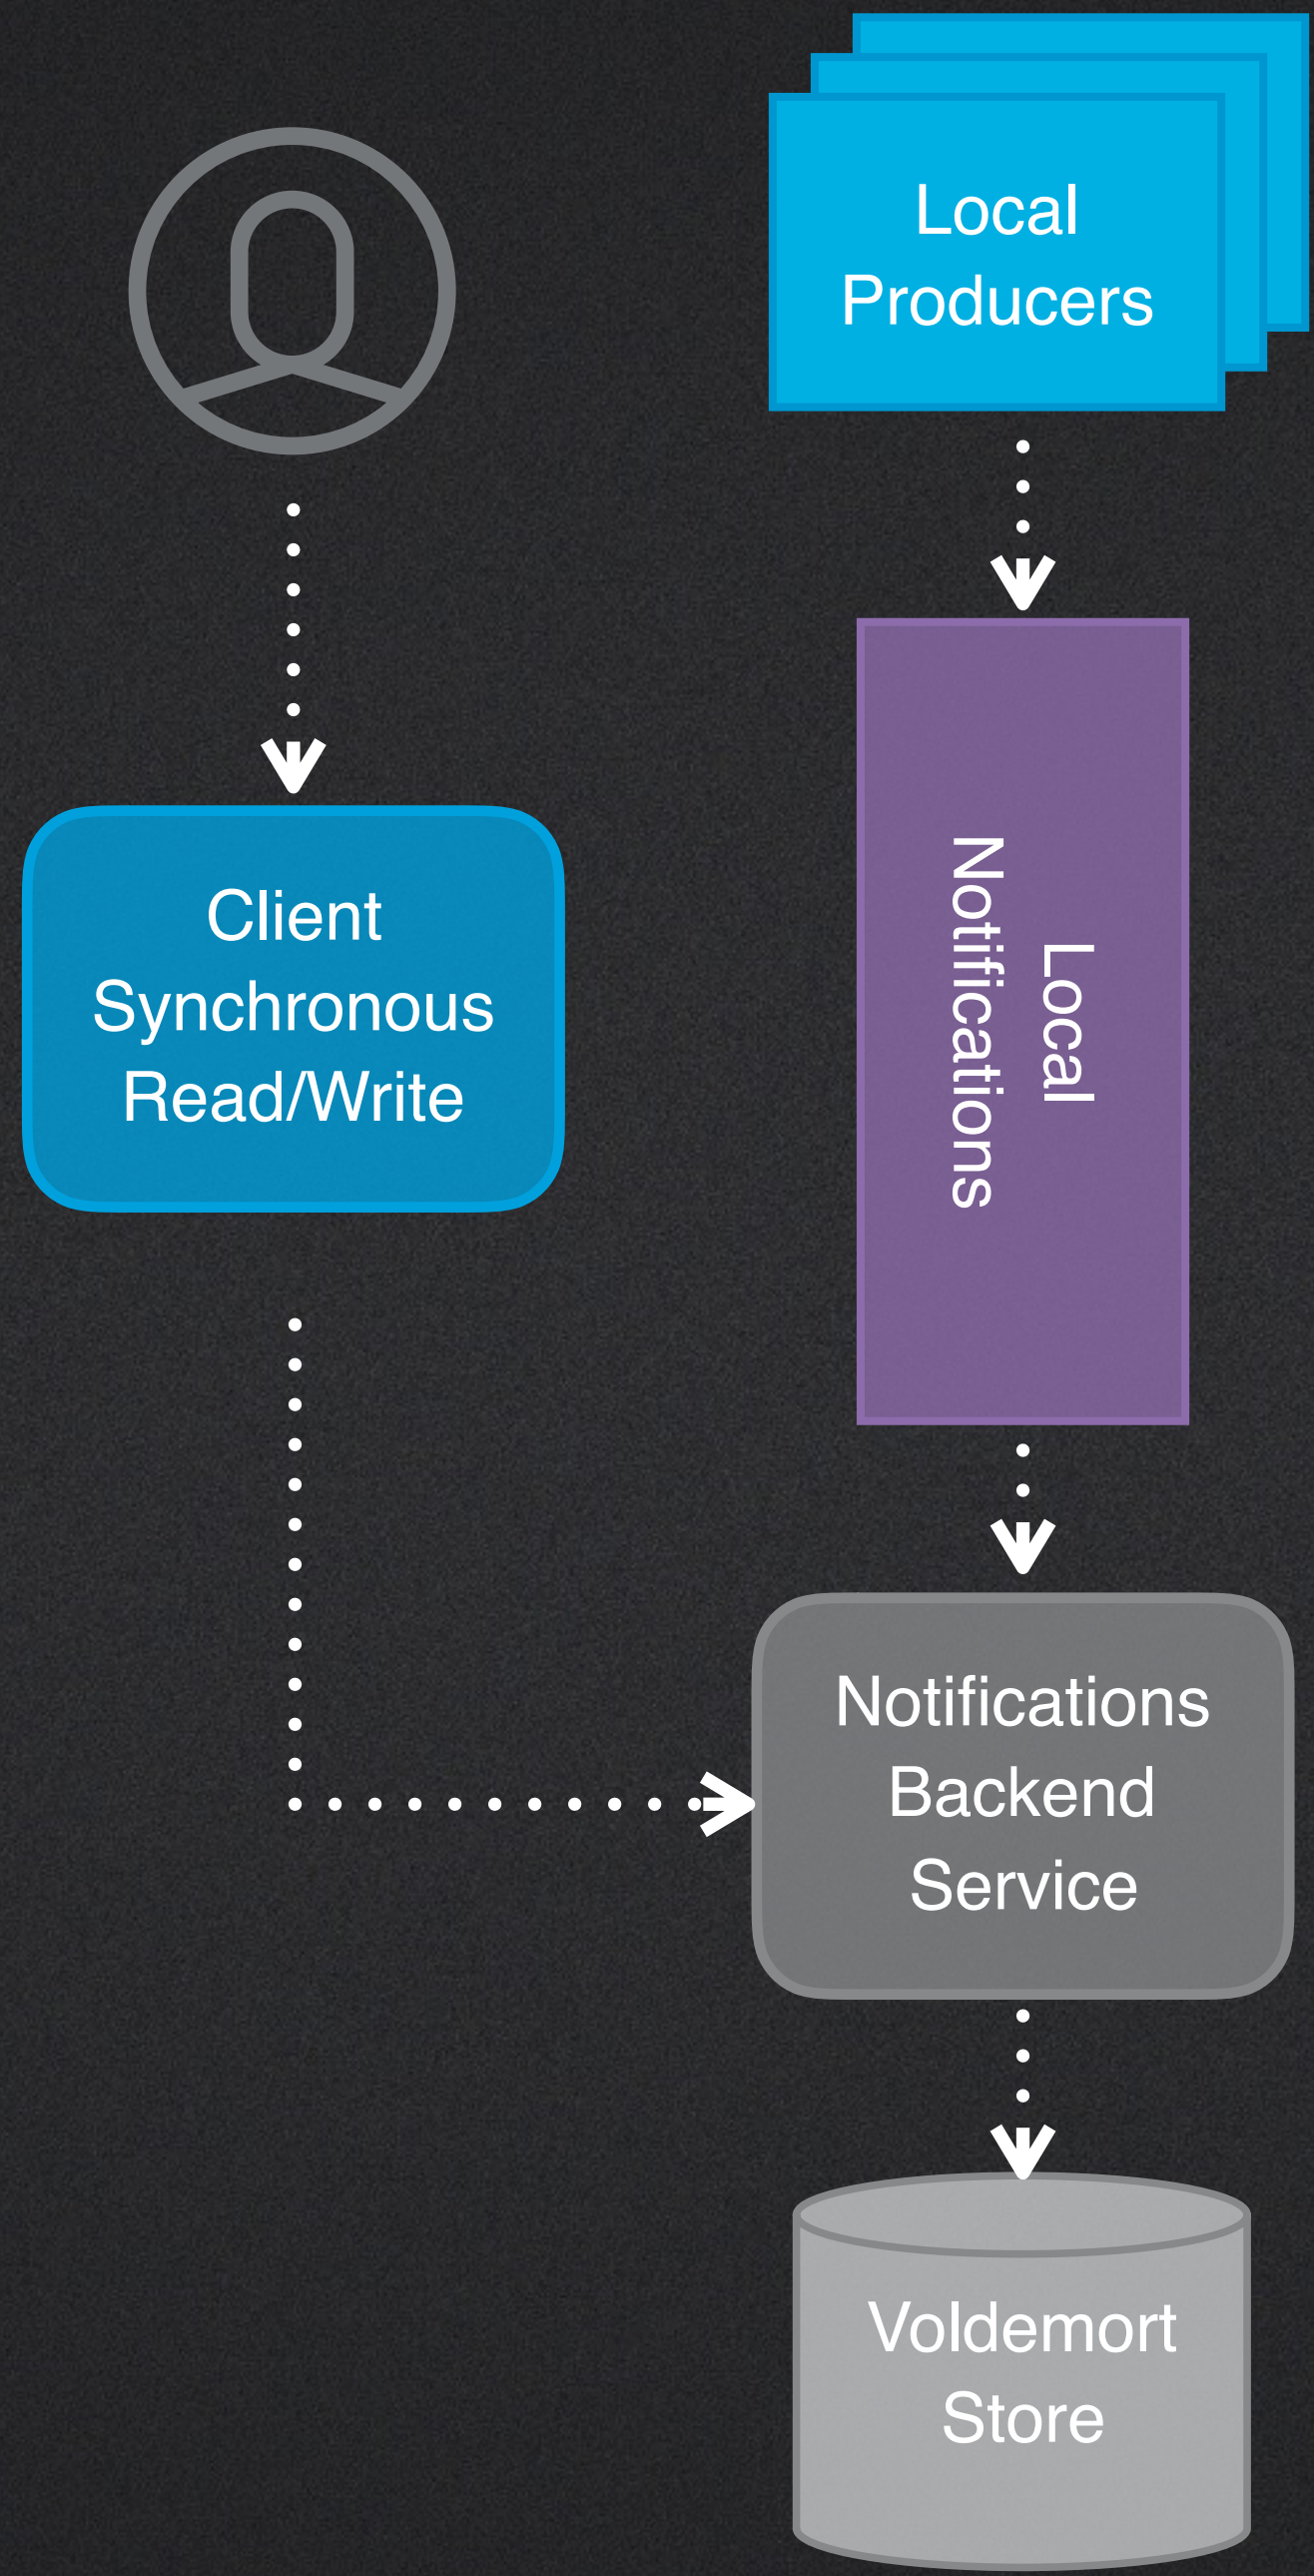


100M

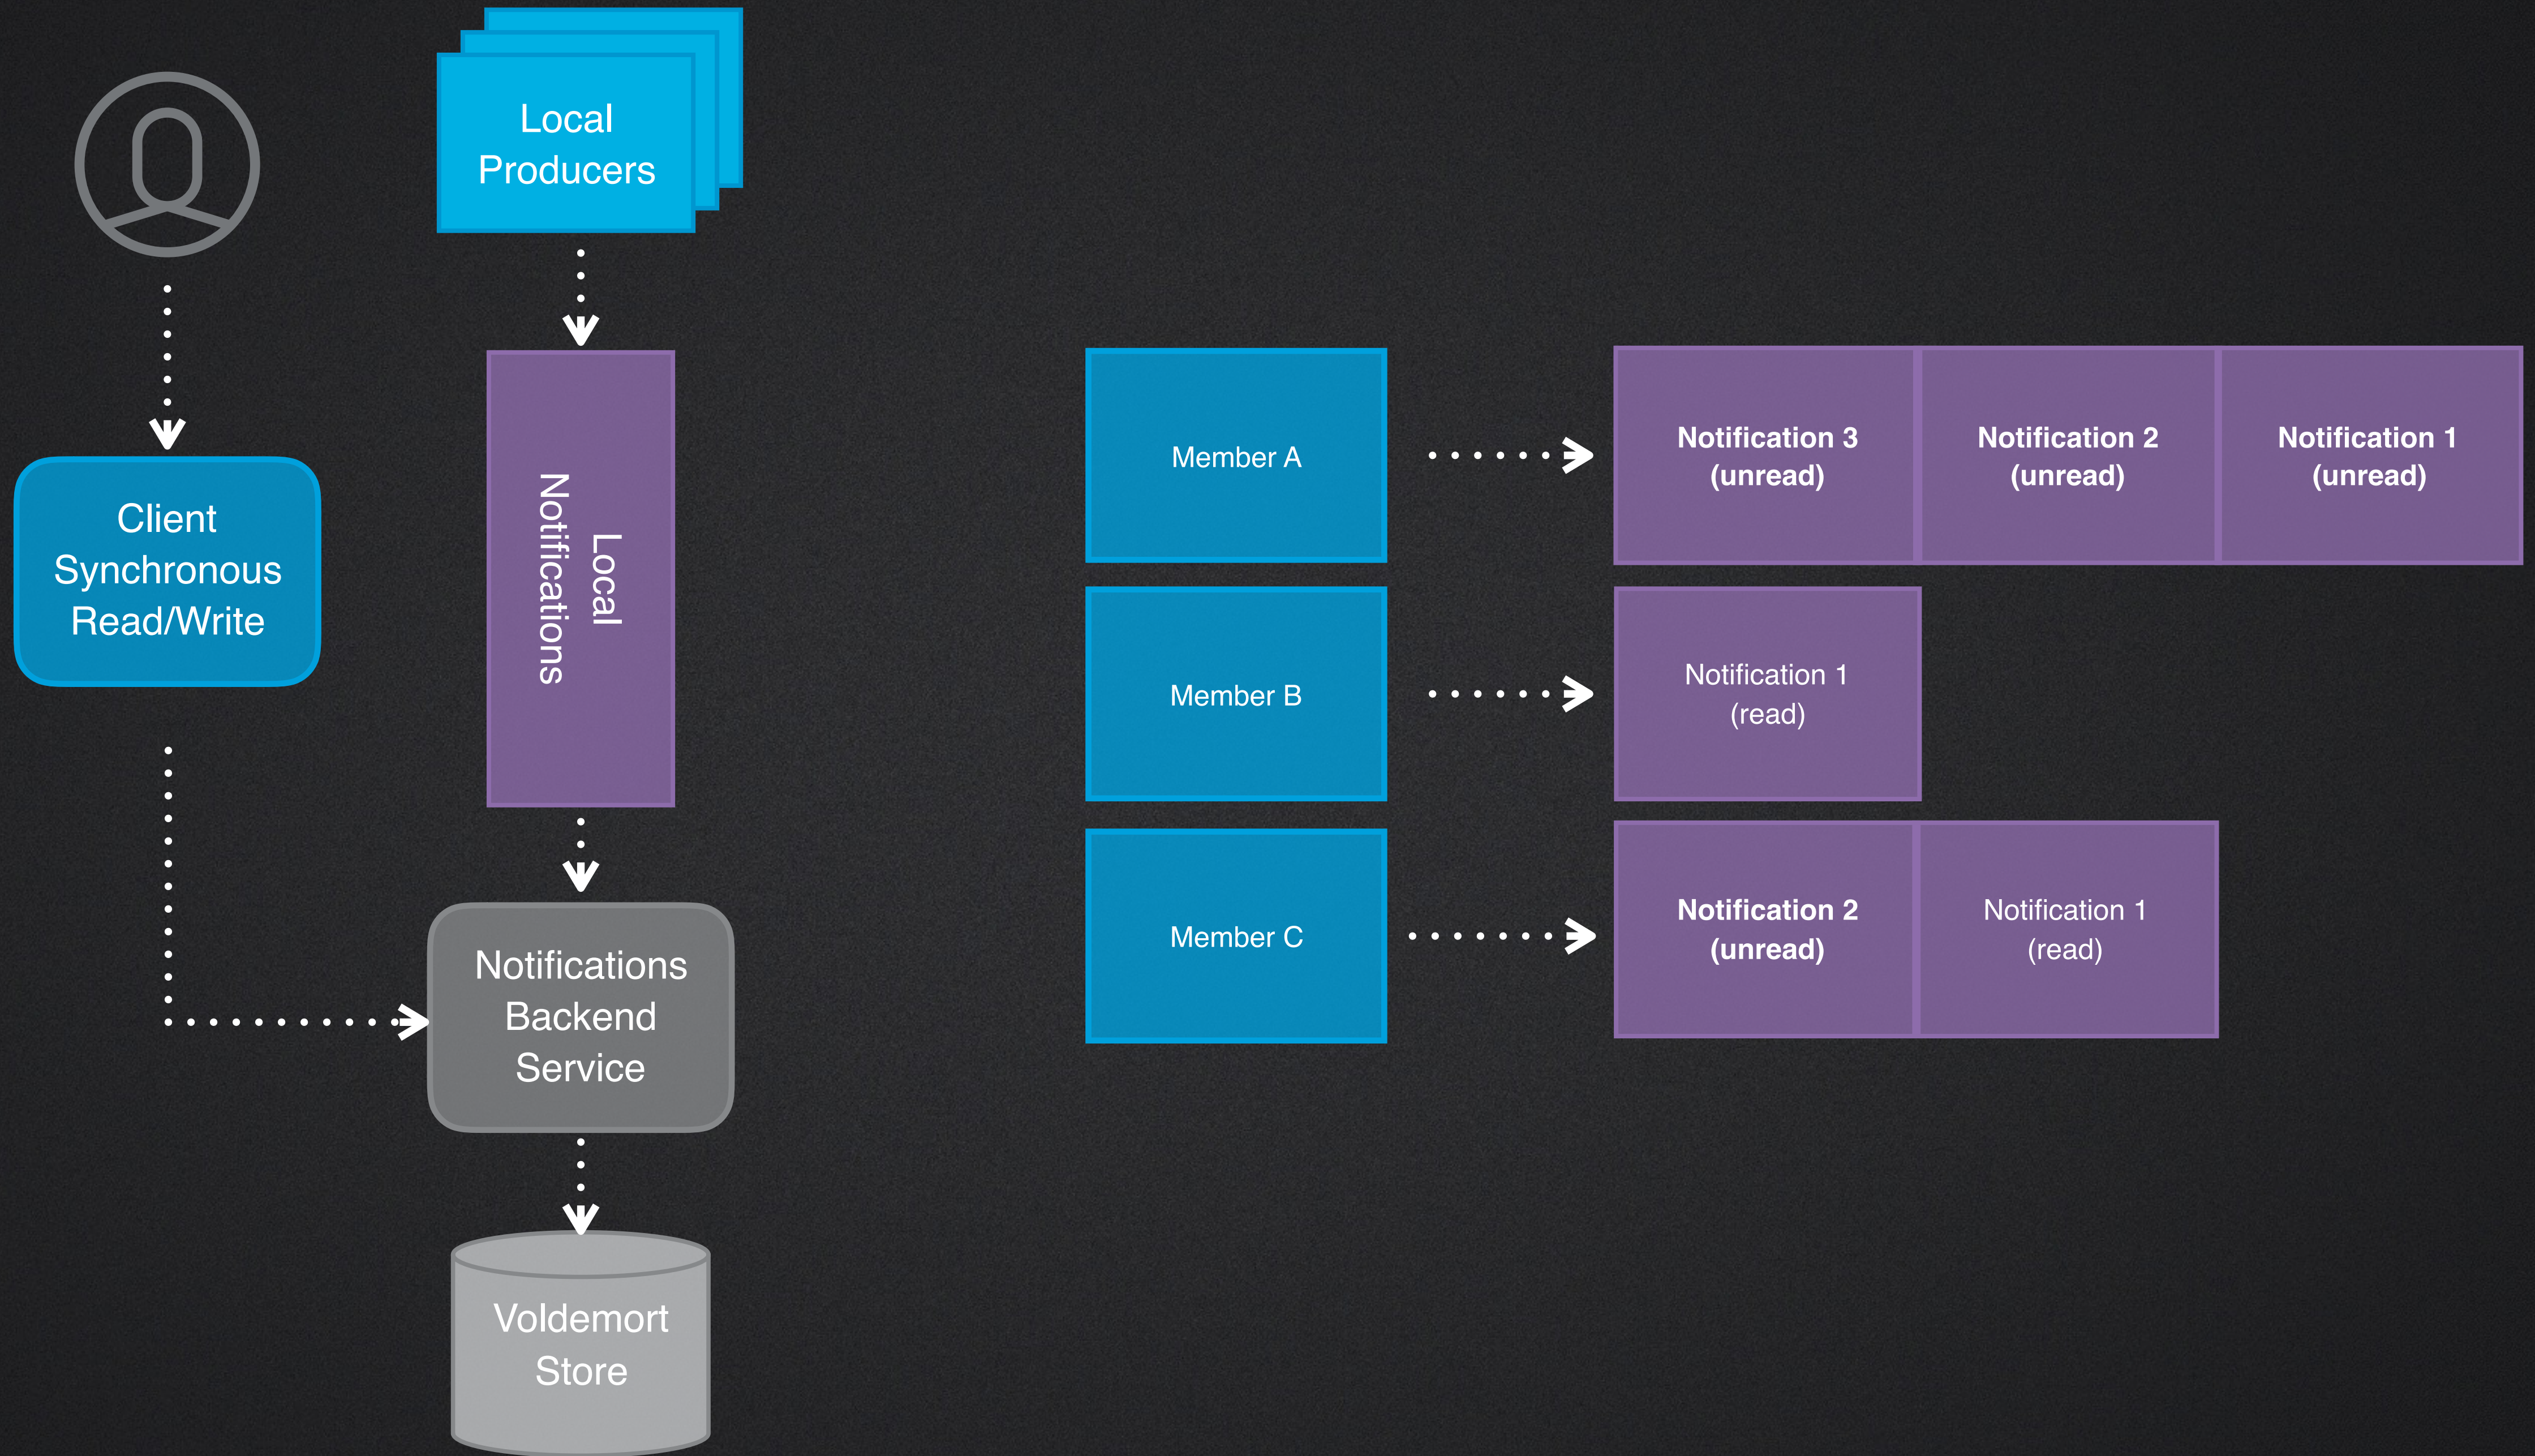
---

Notifications per day



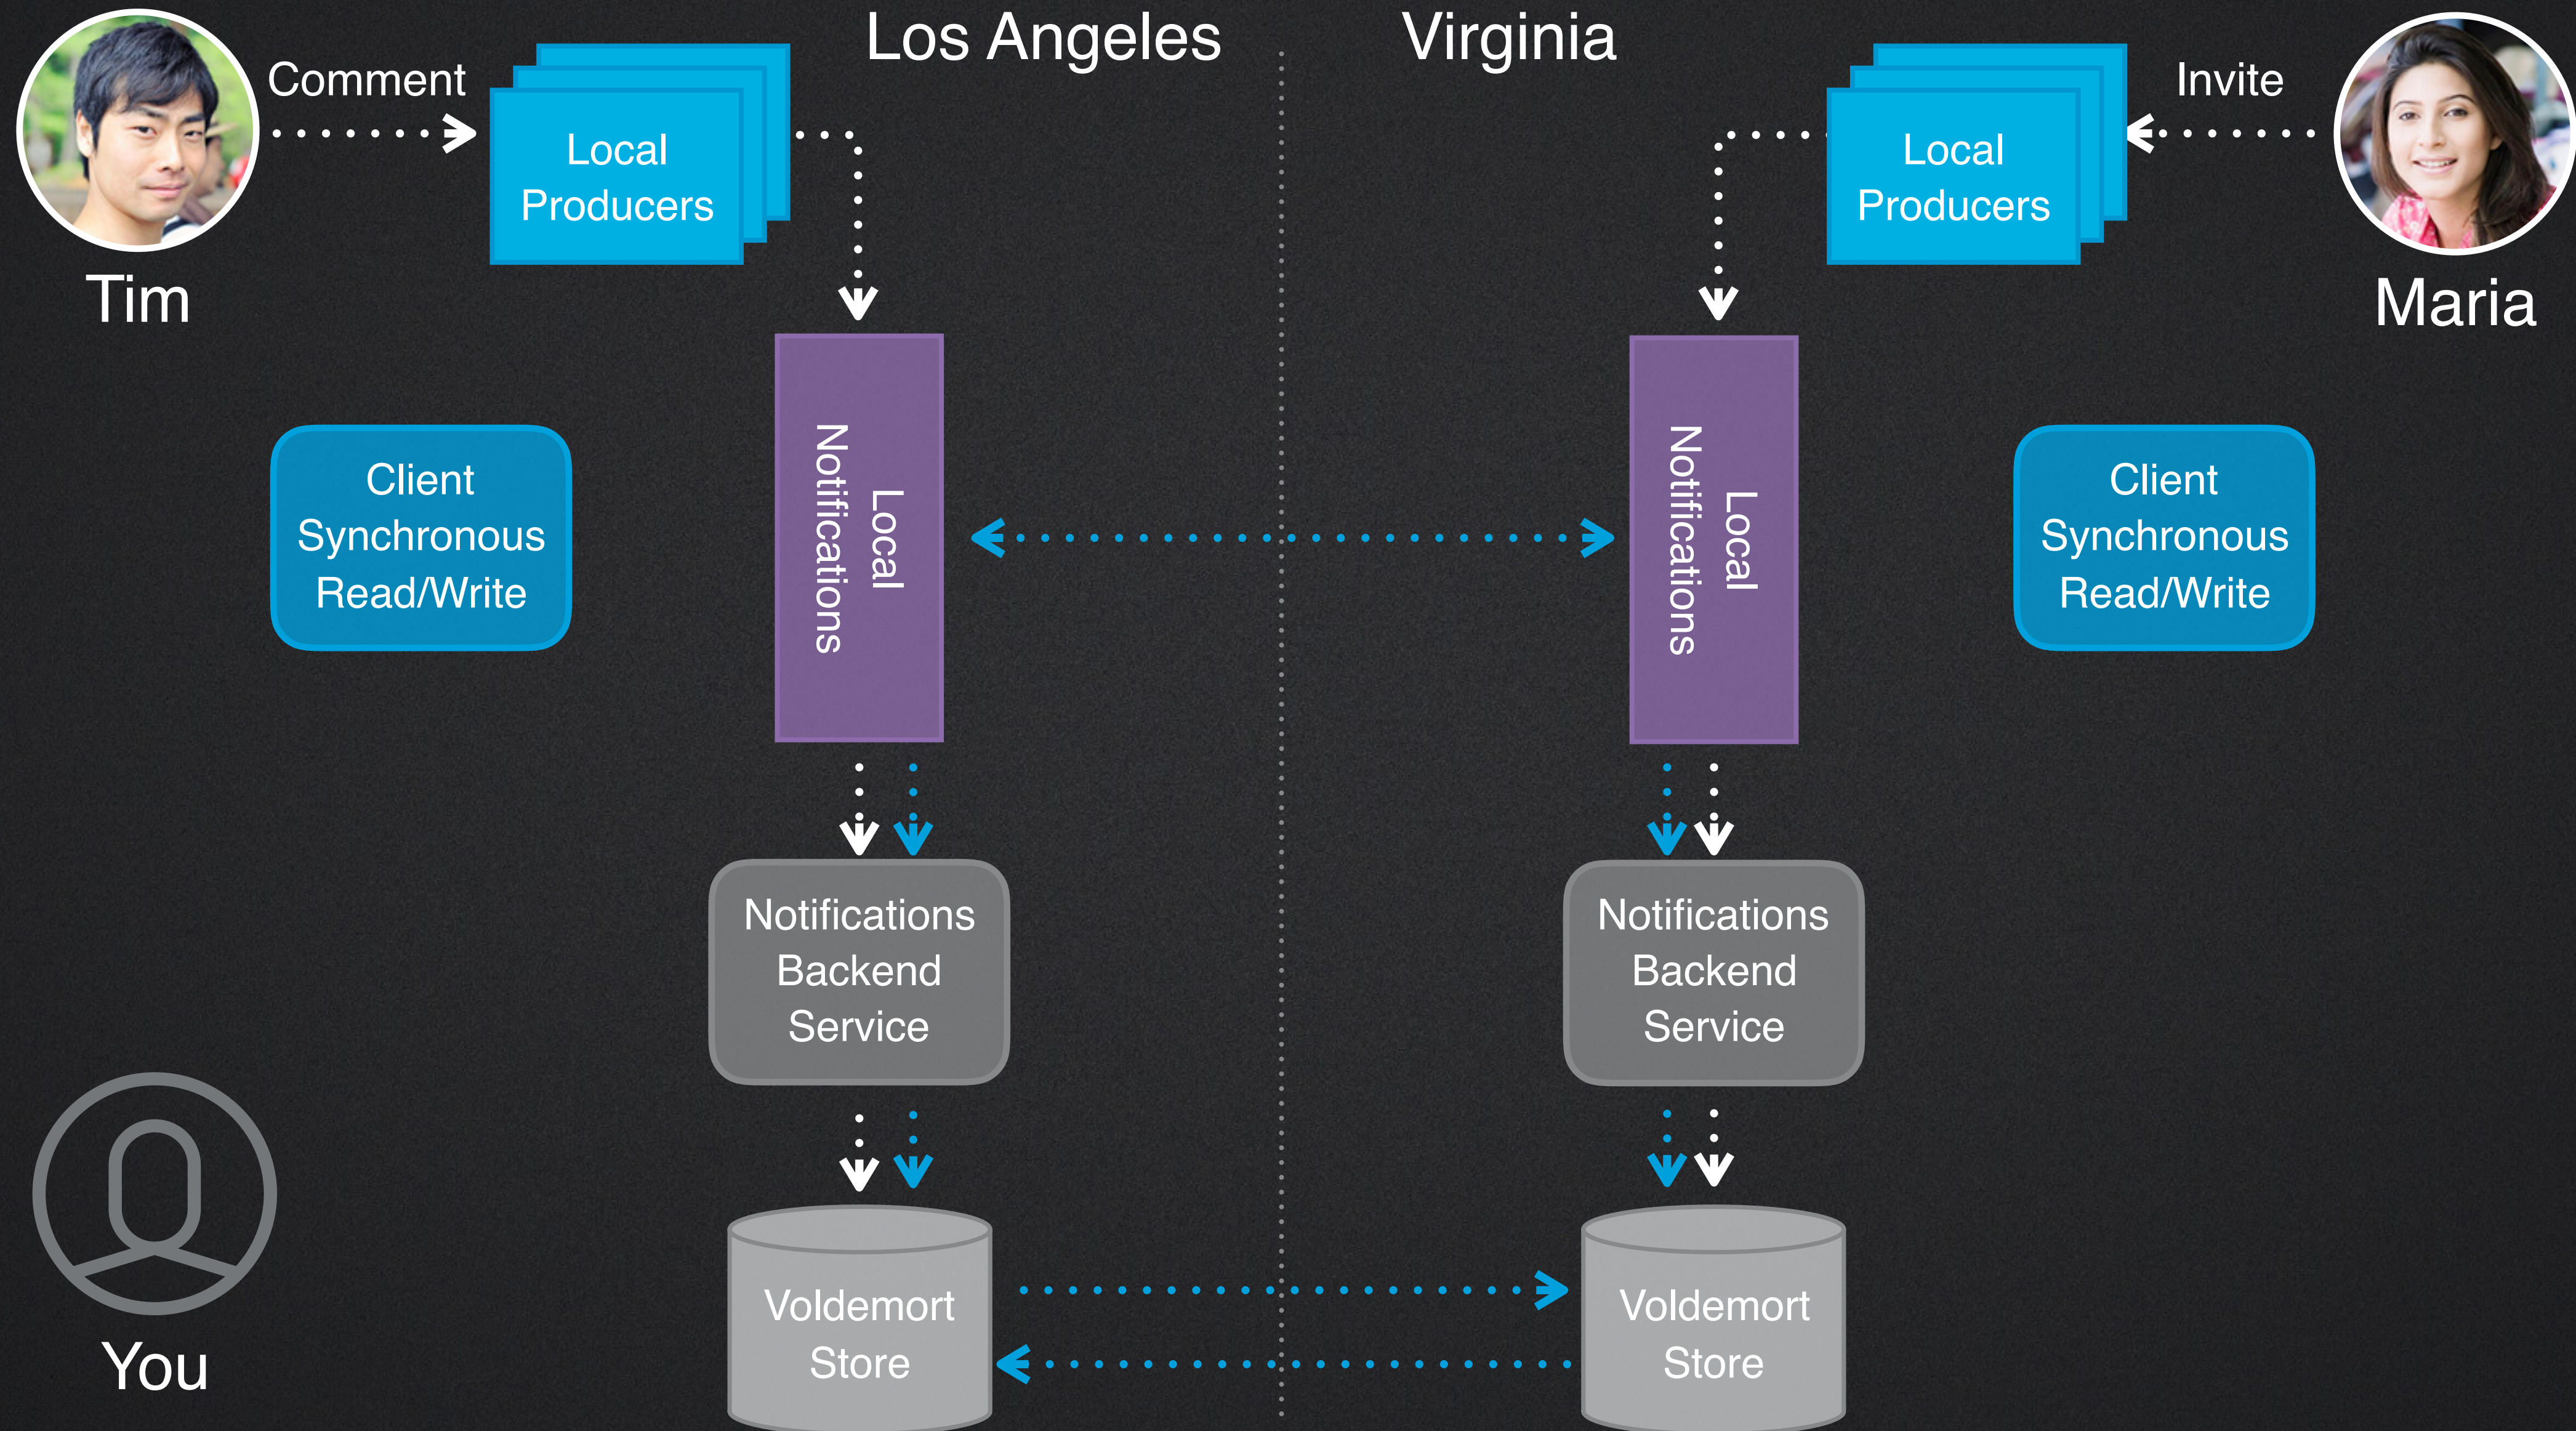






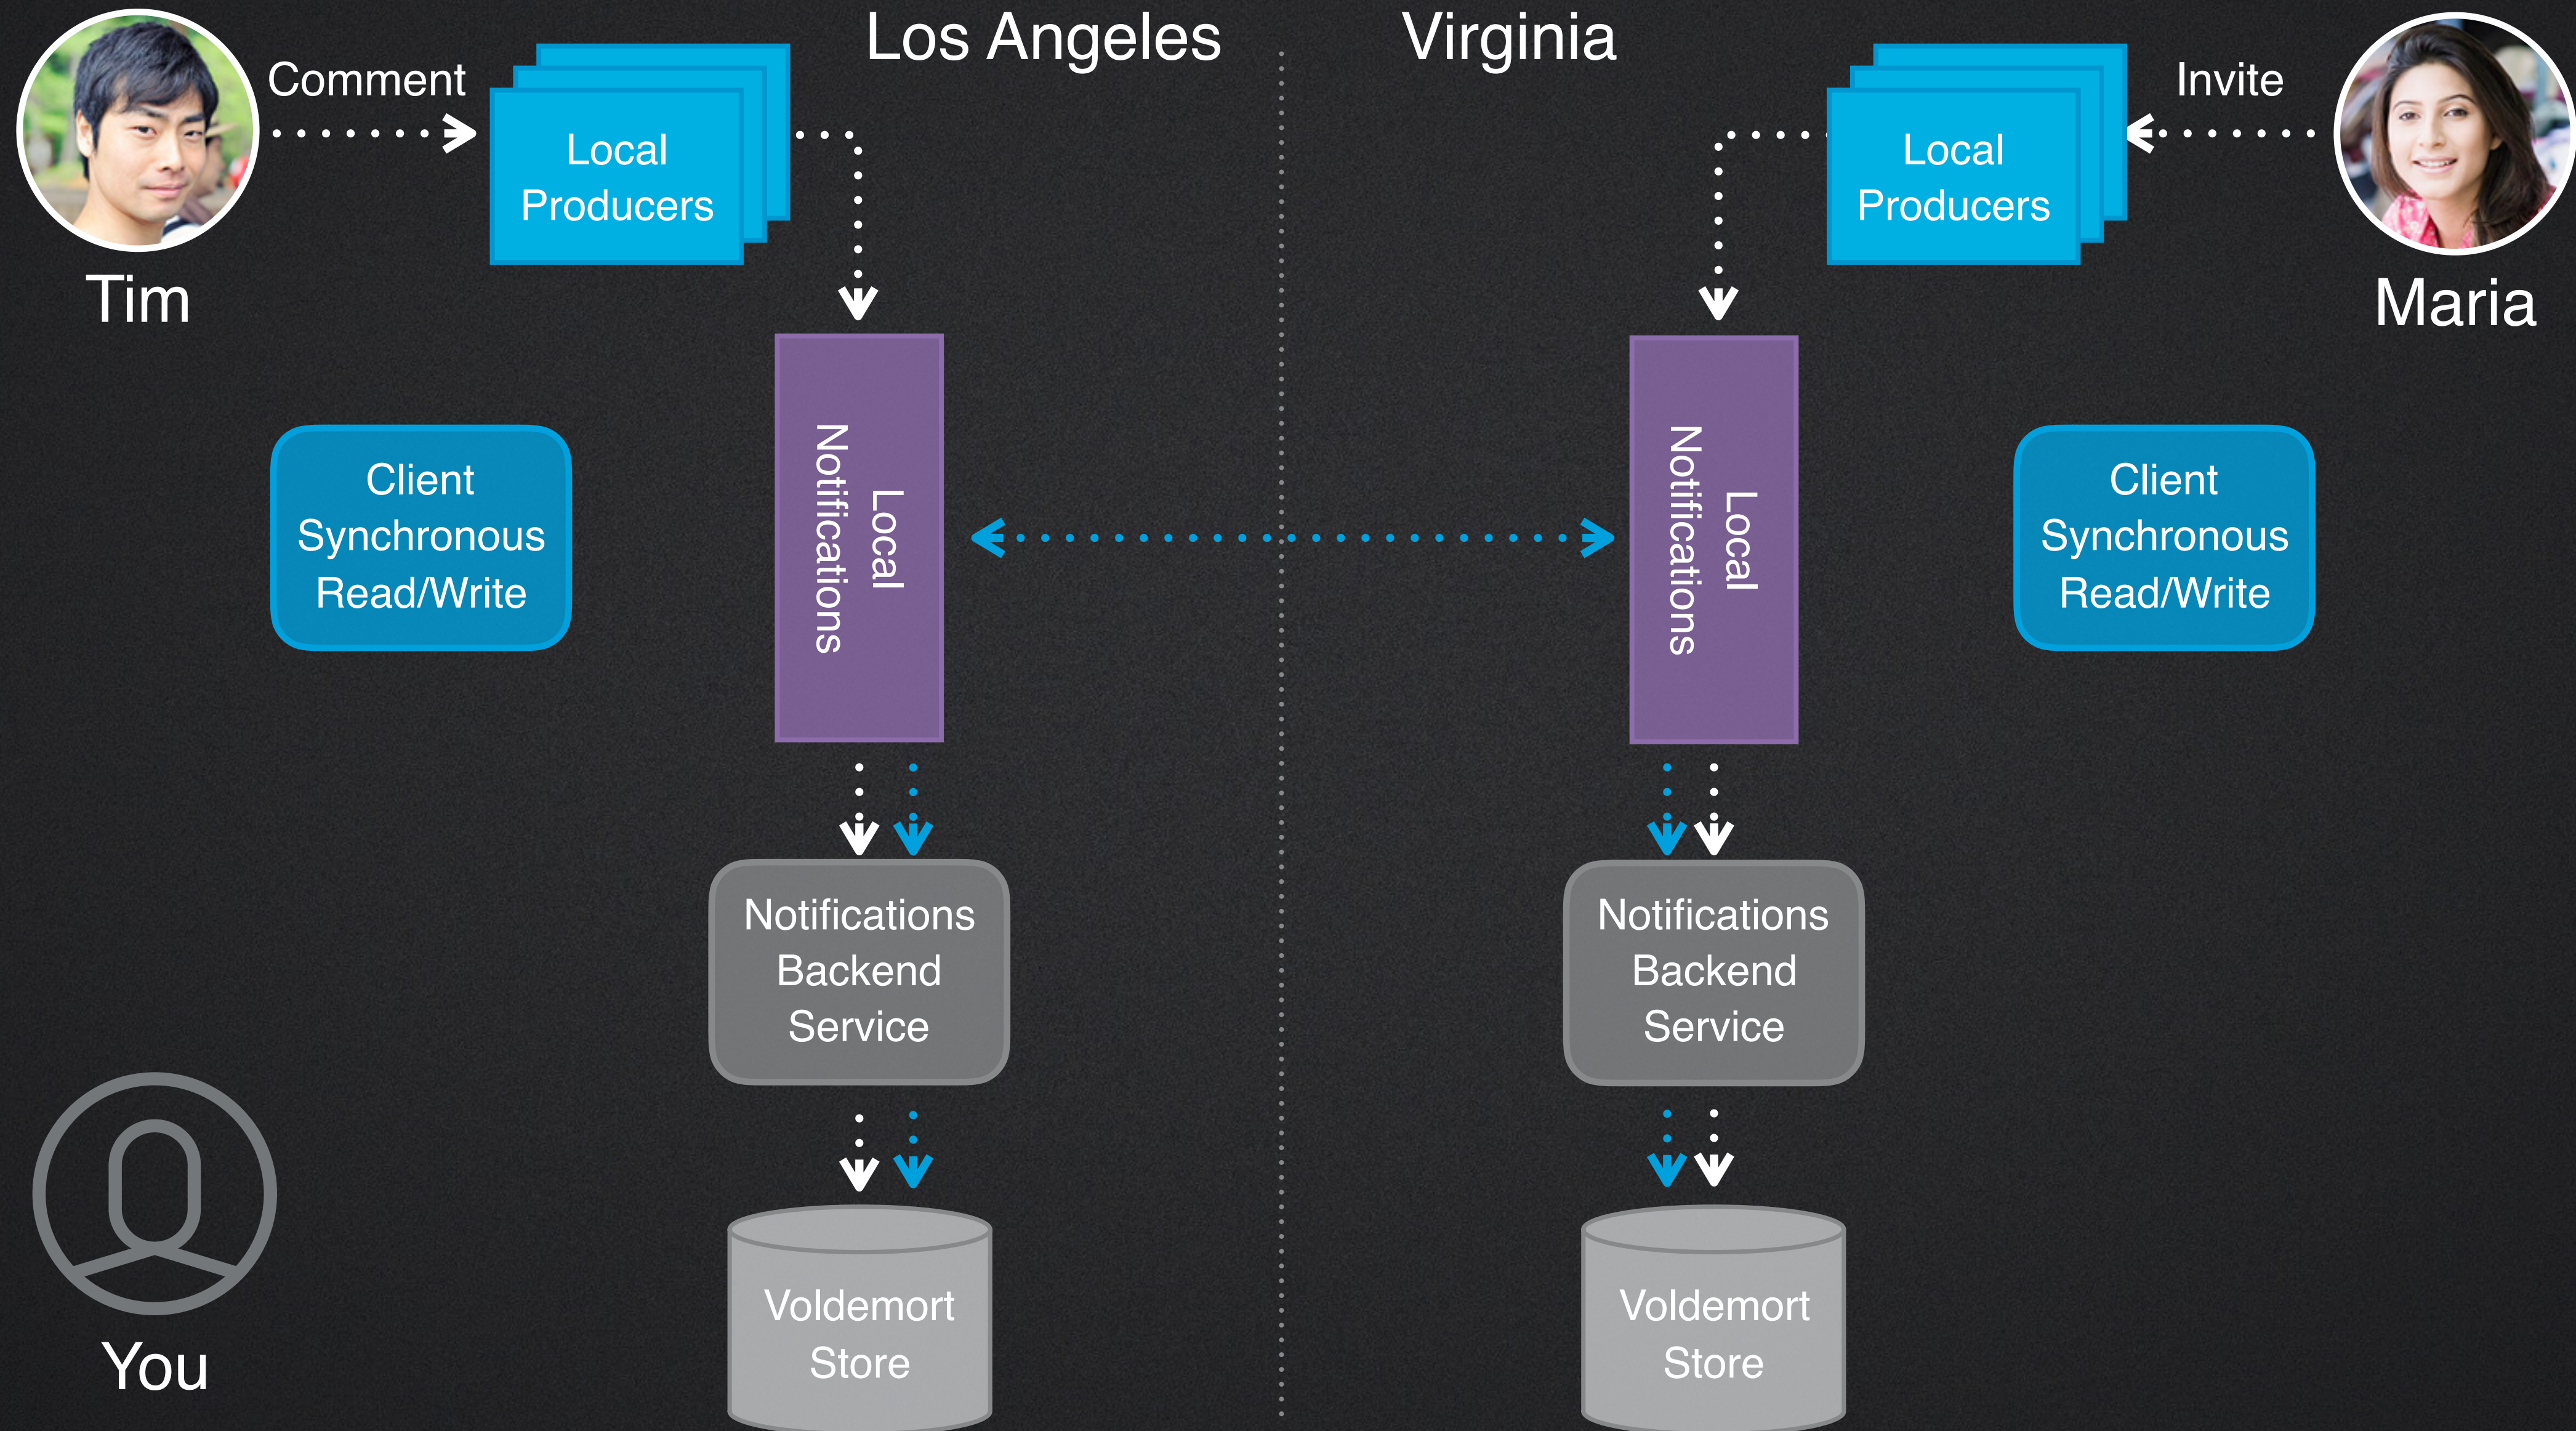


# Problem: Double Processing



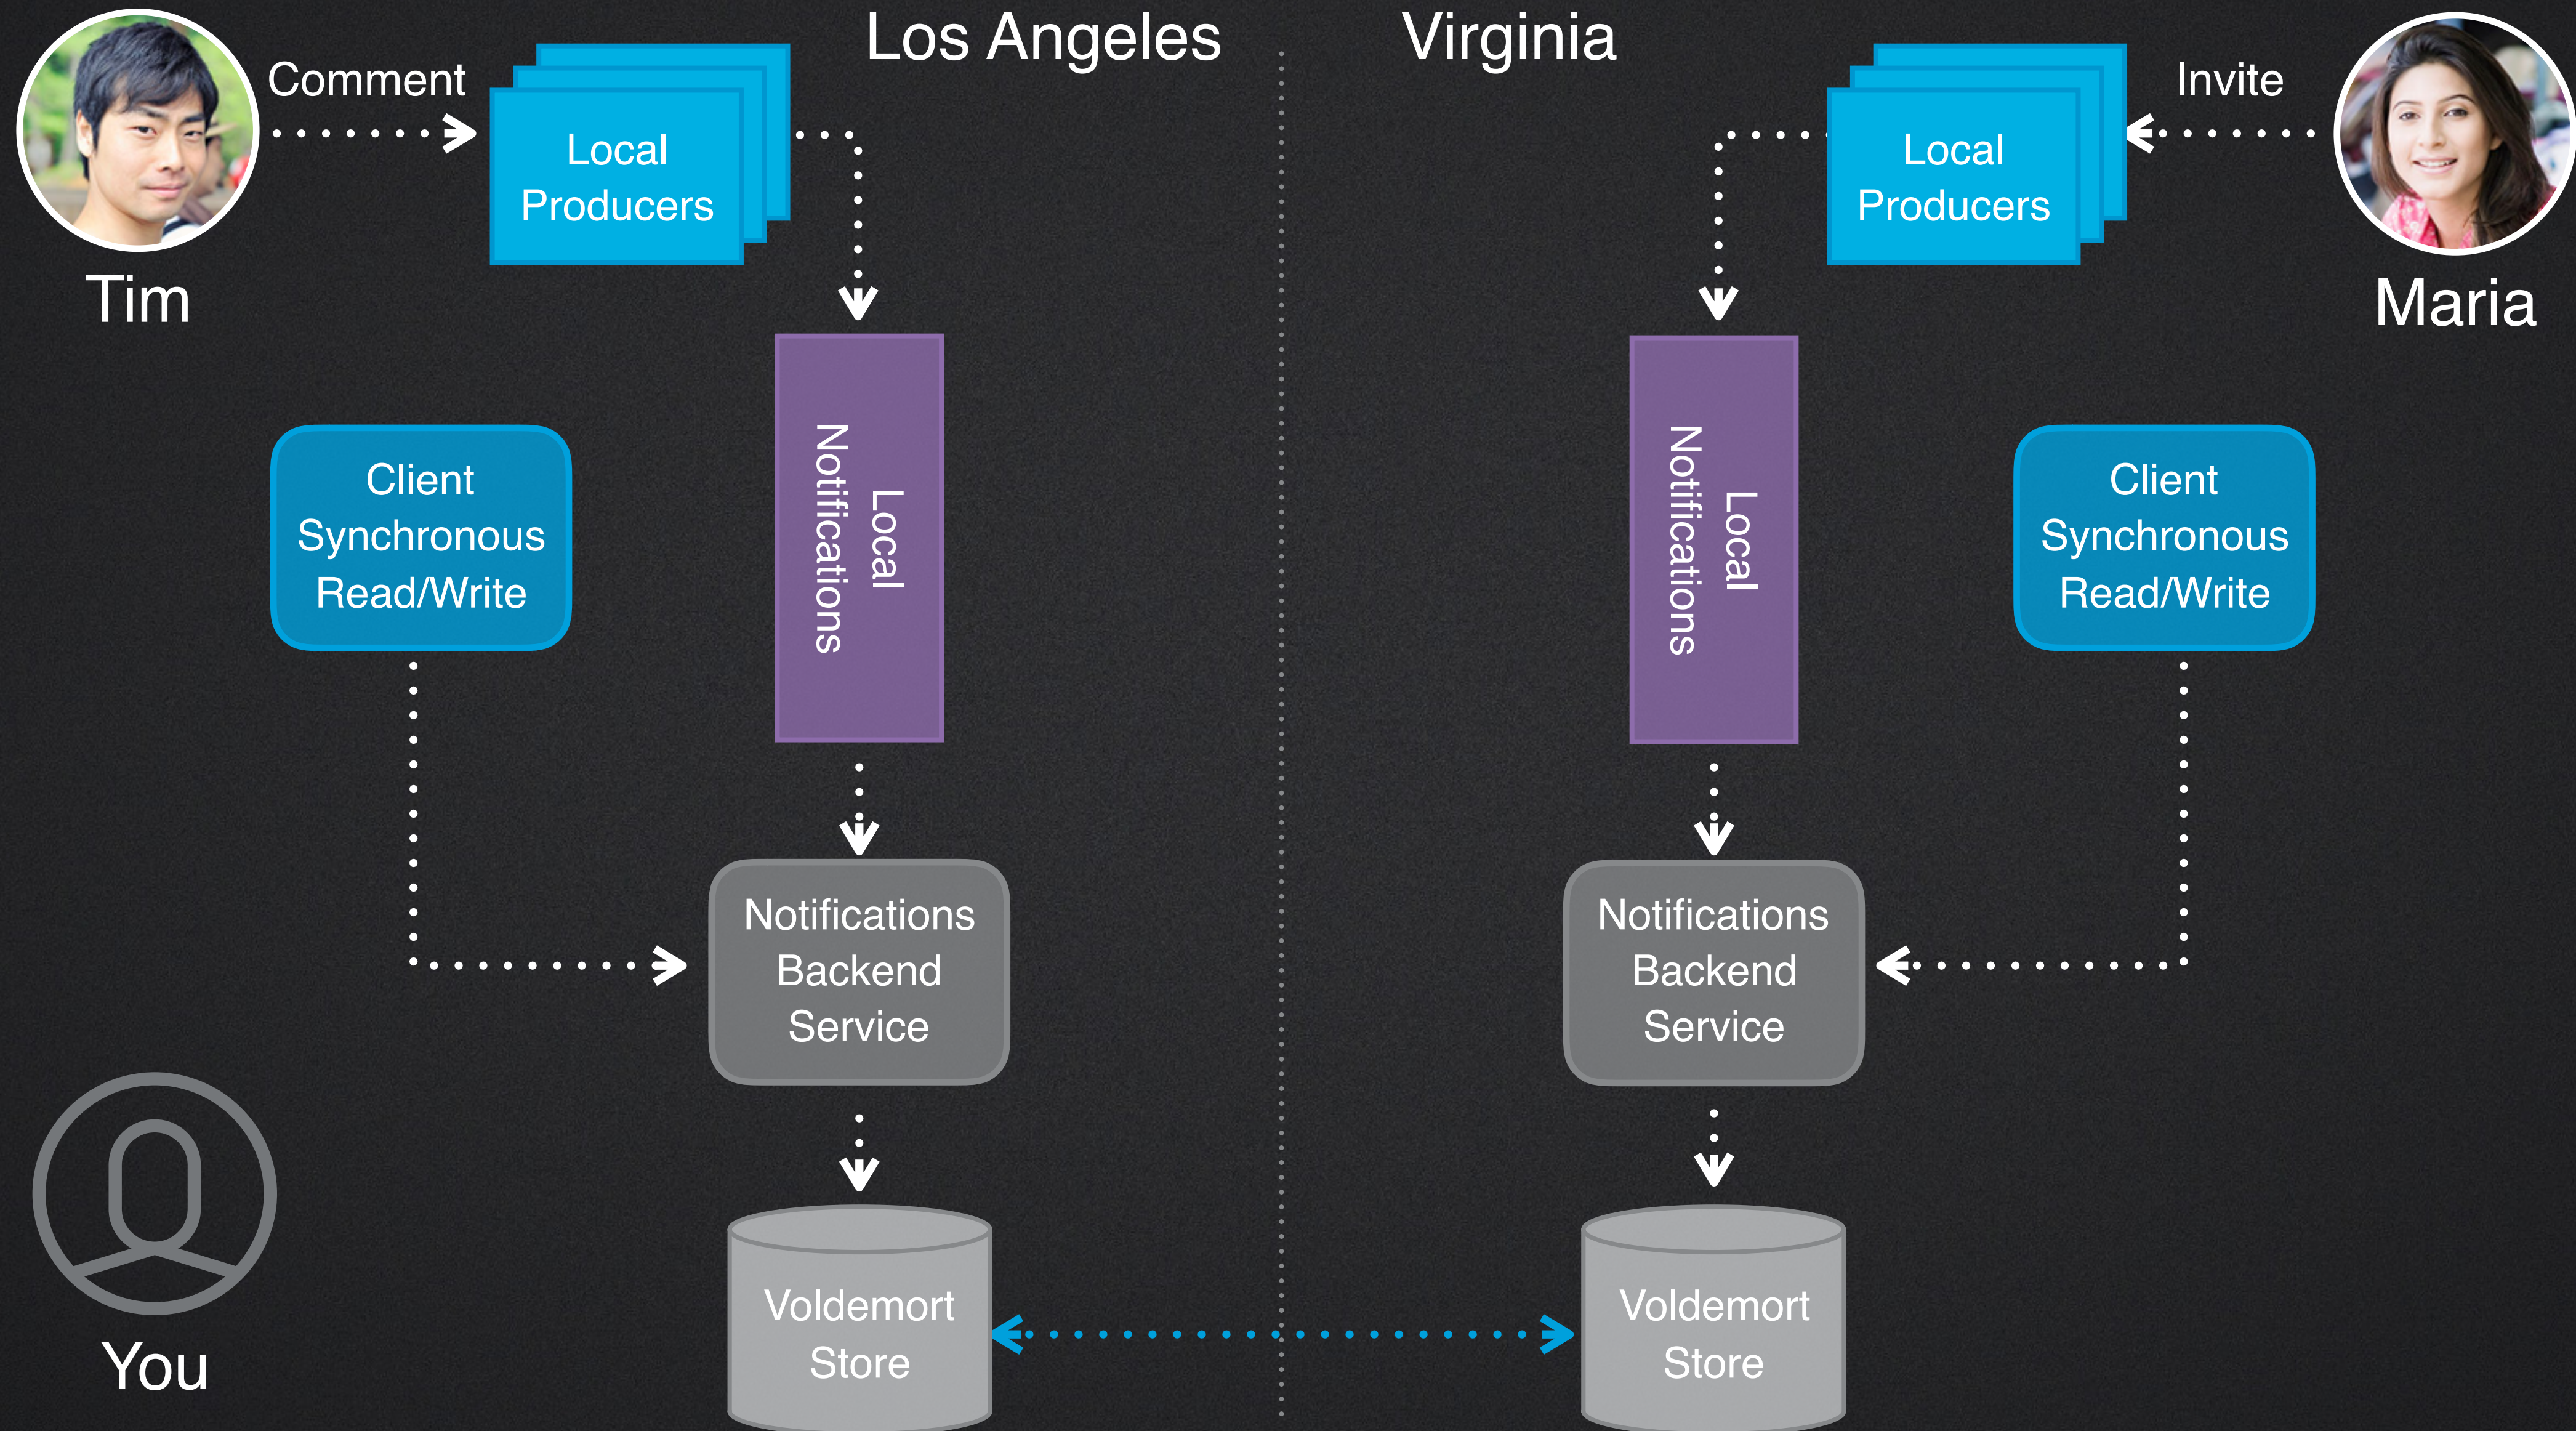


# Option 1: Replicate Kafka only





# Option 2: Replicate Voldemort only





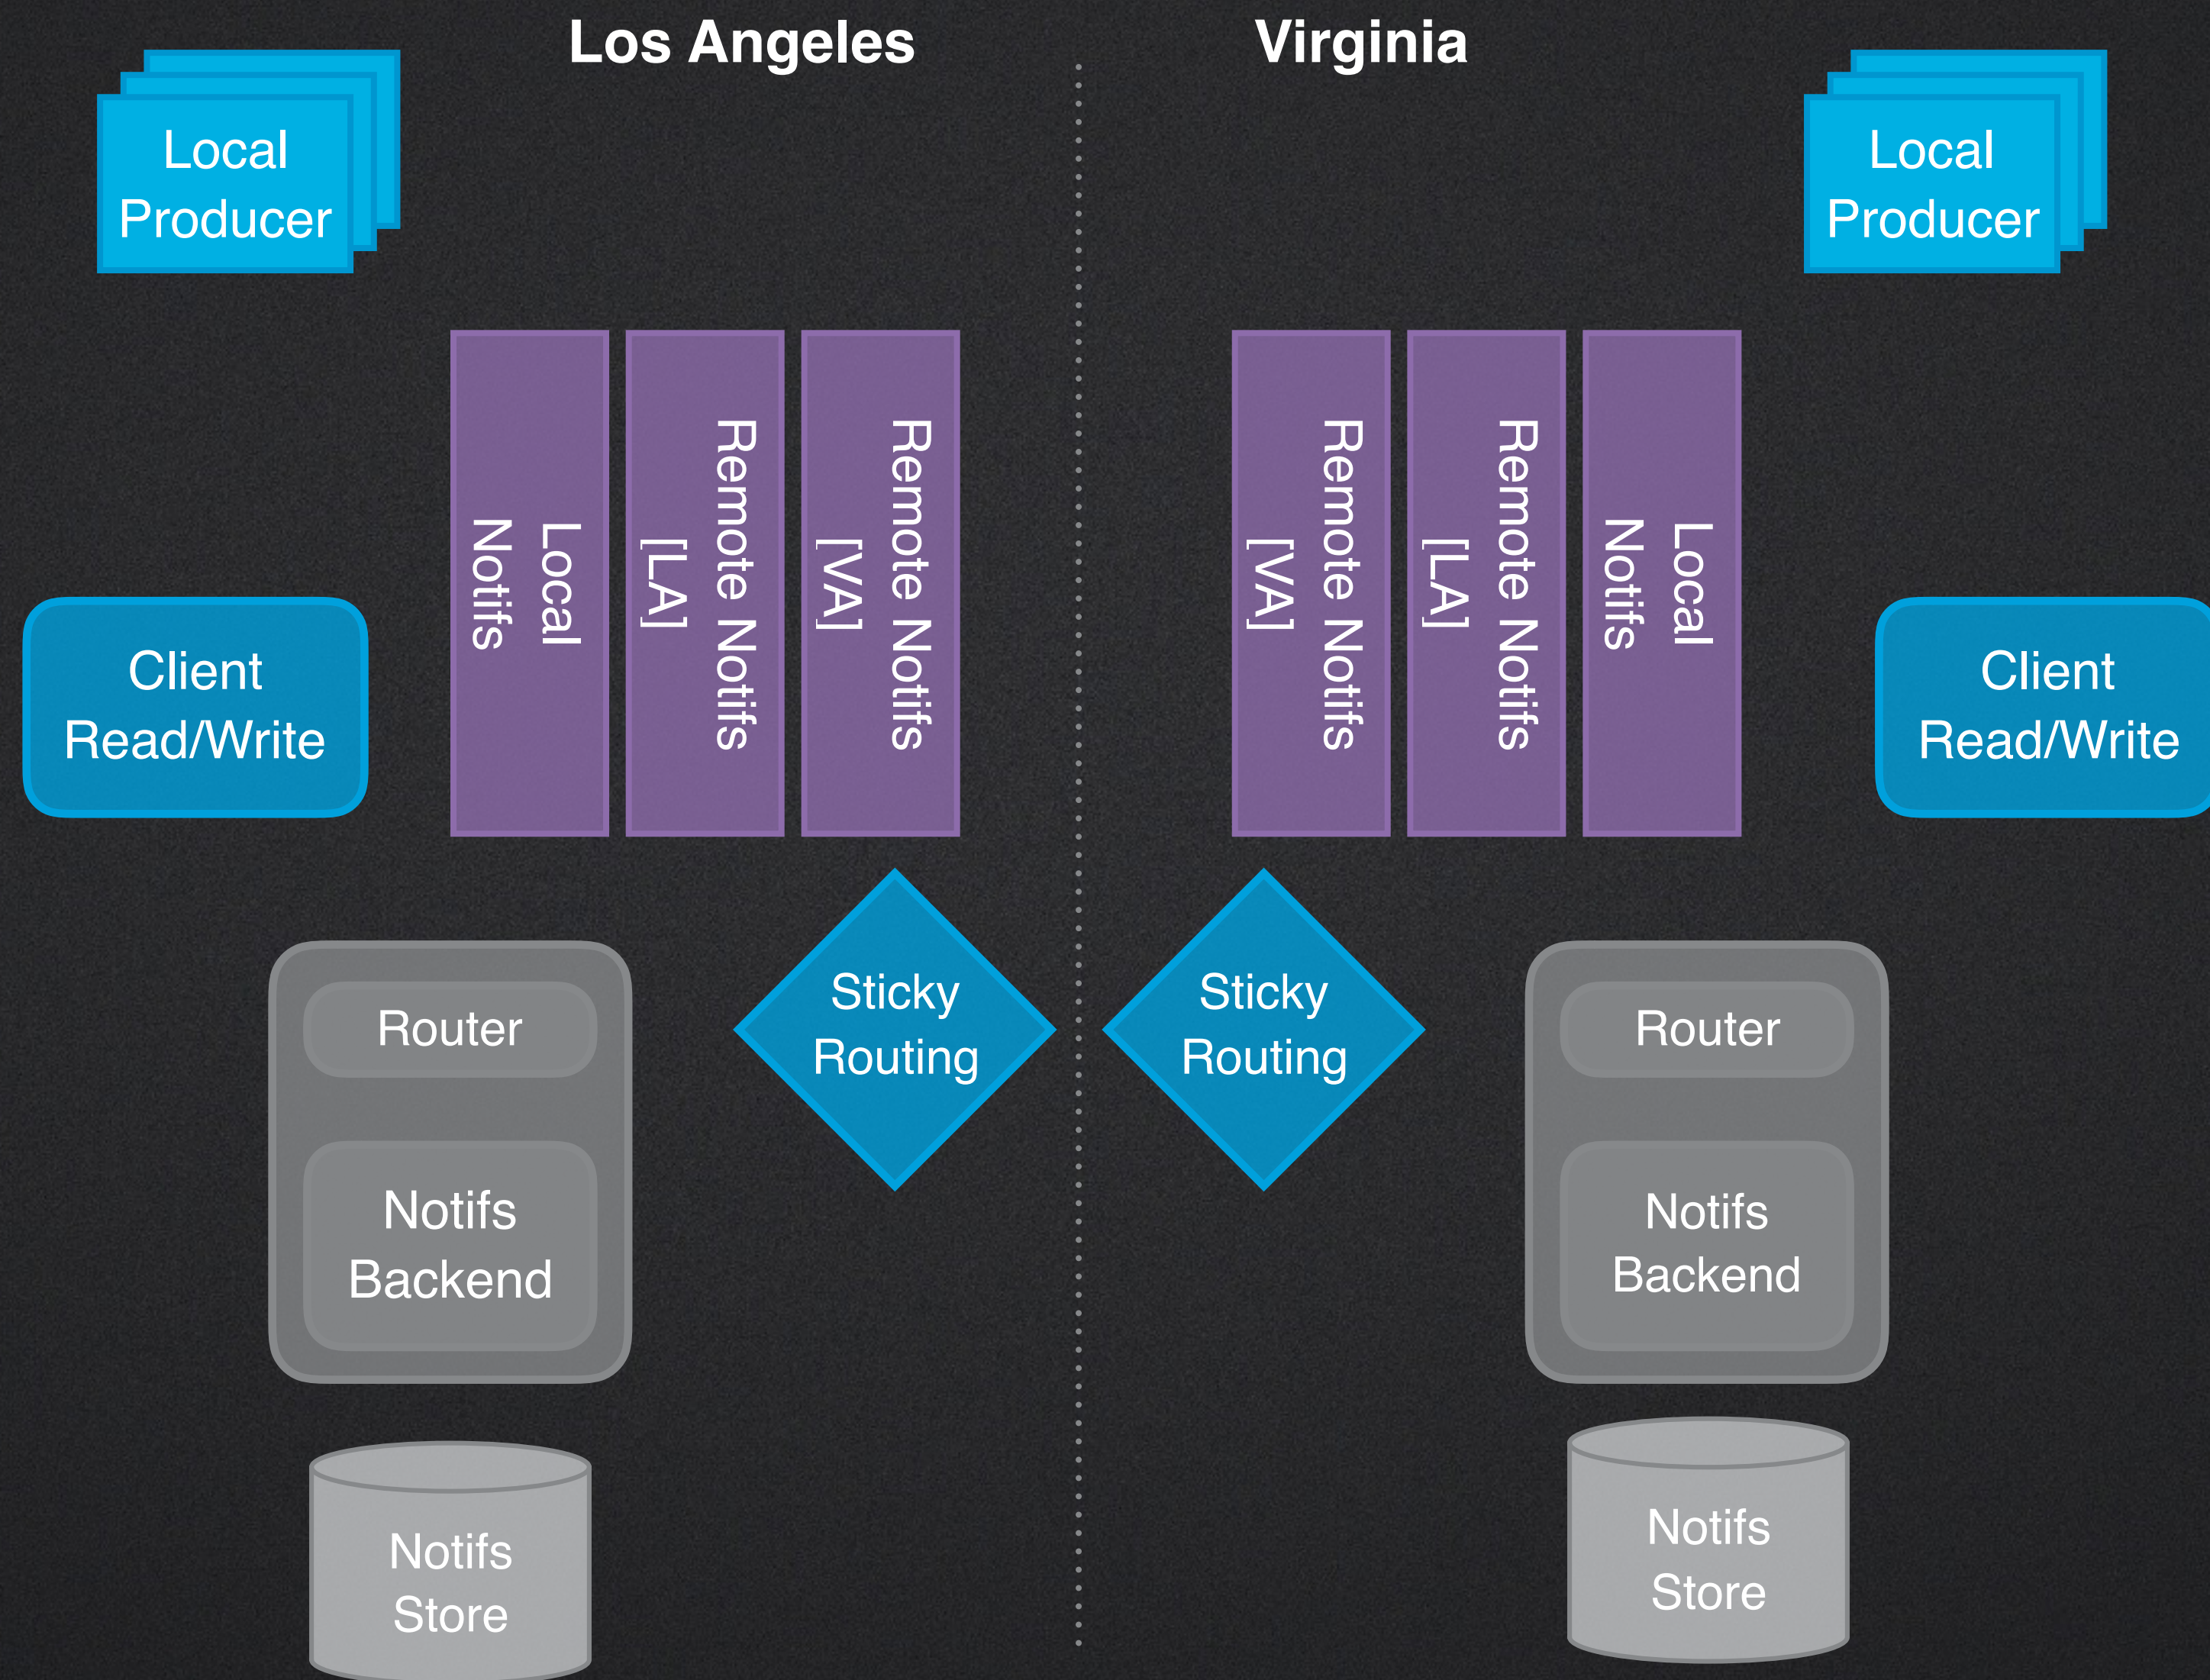
# The Solution



Tim



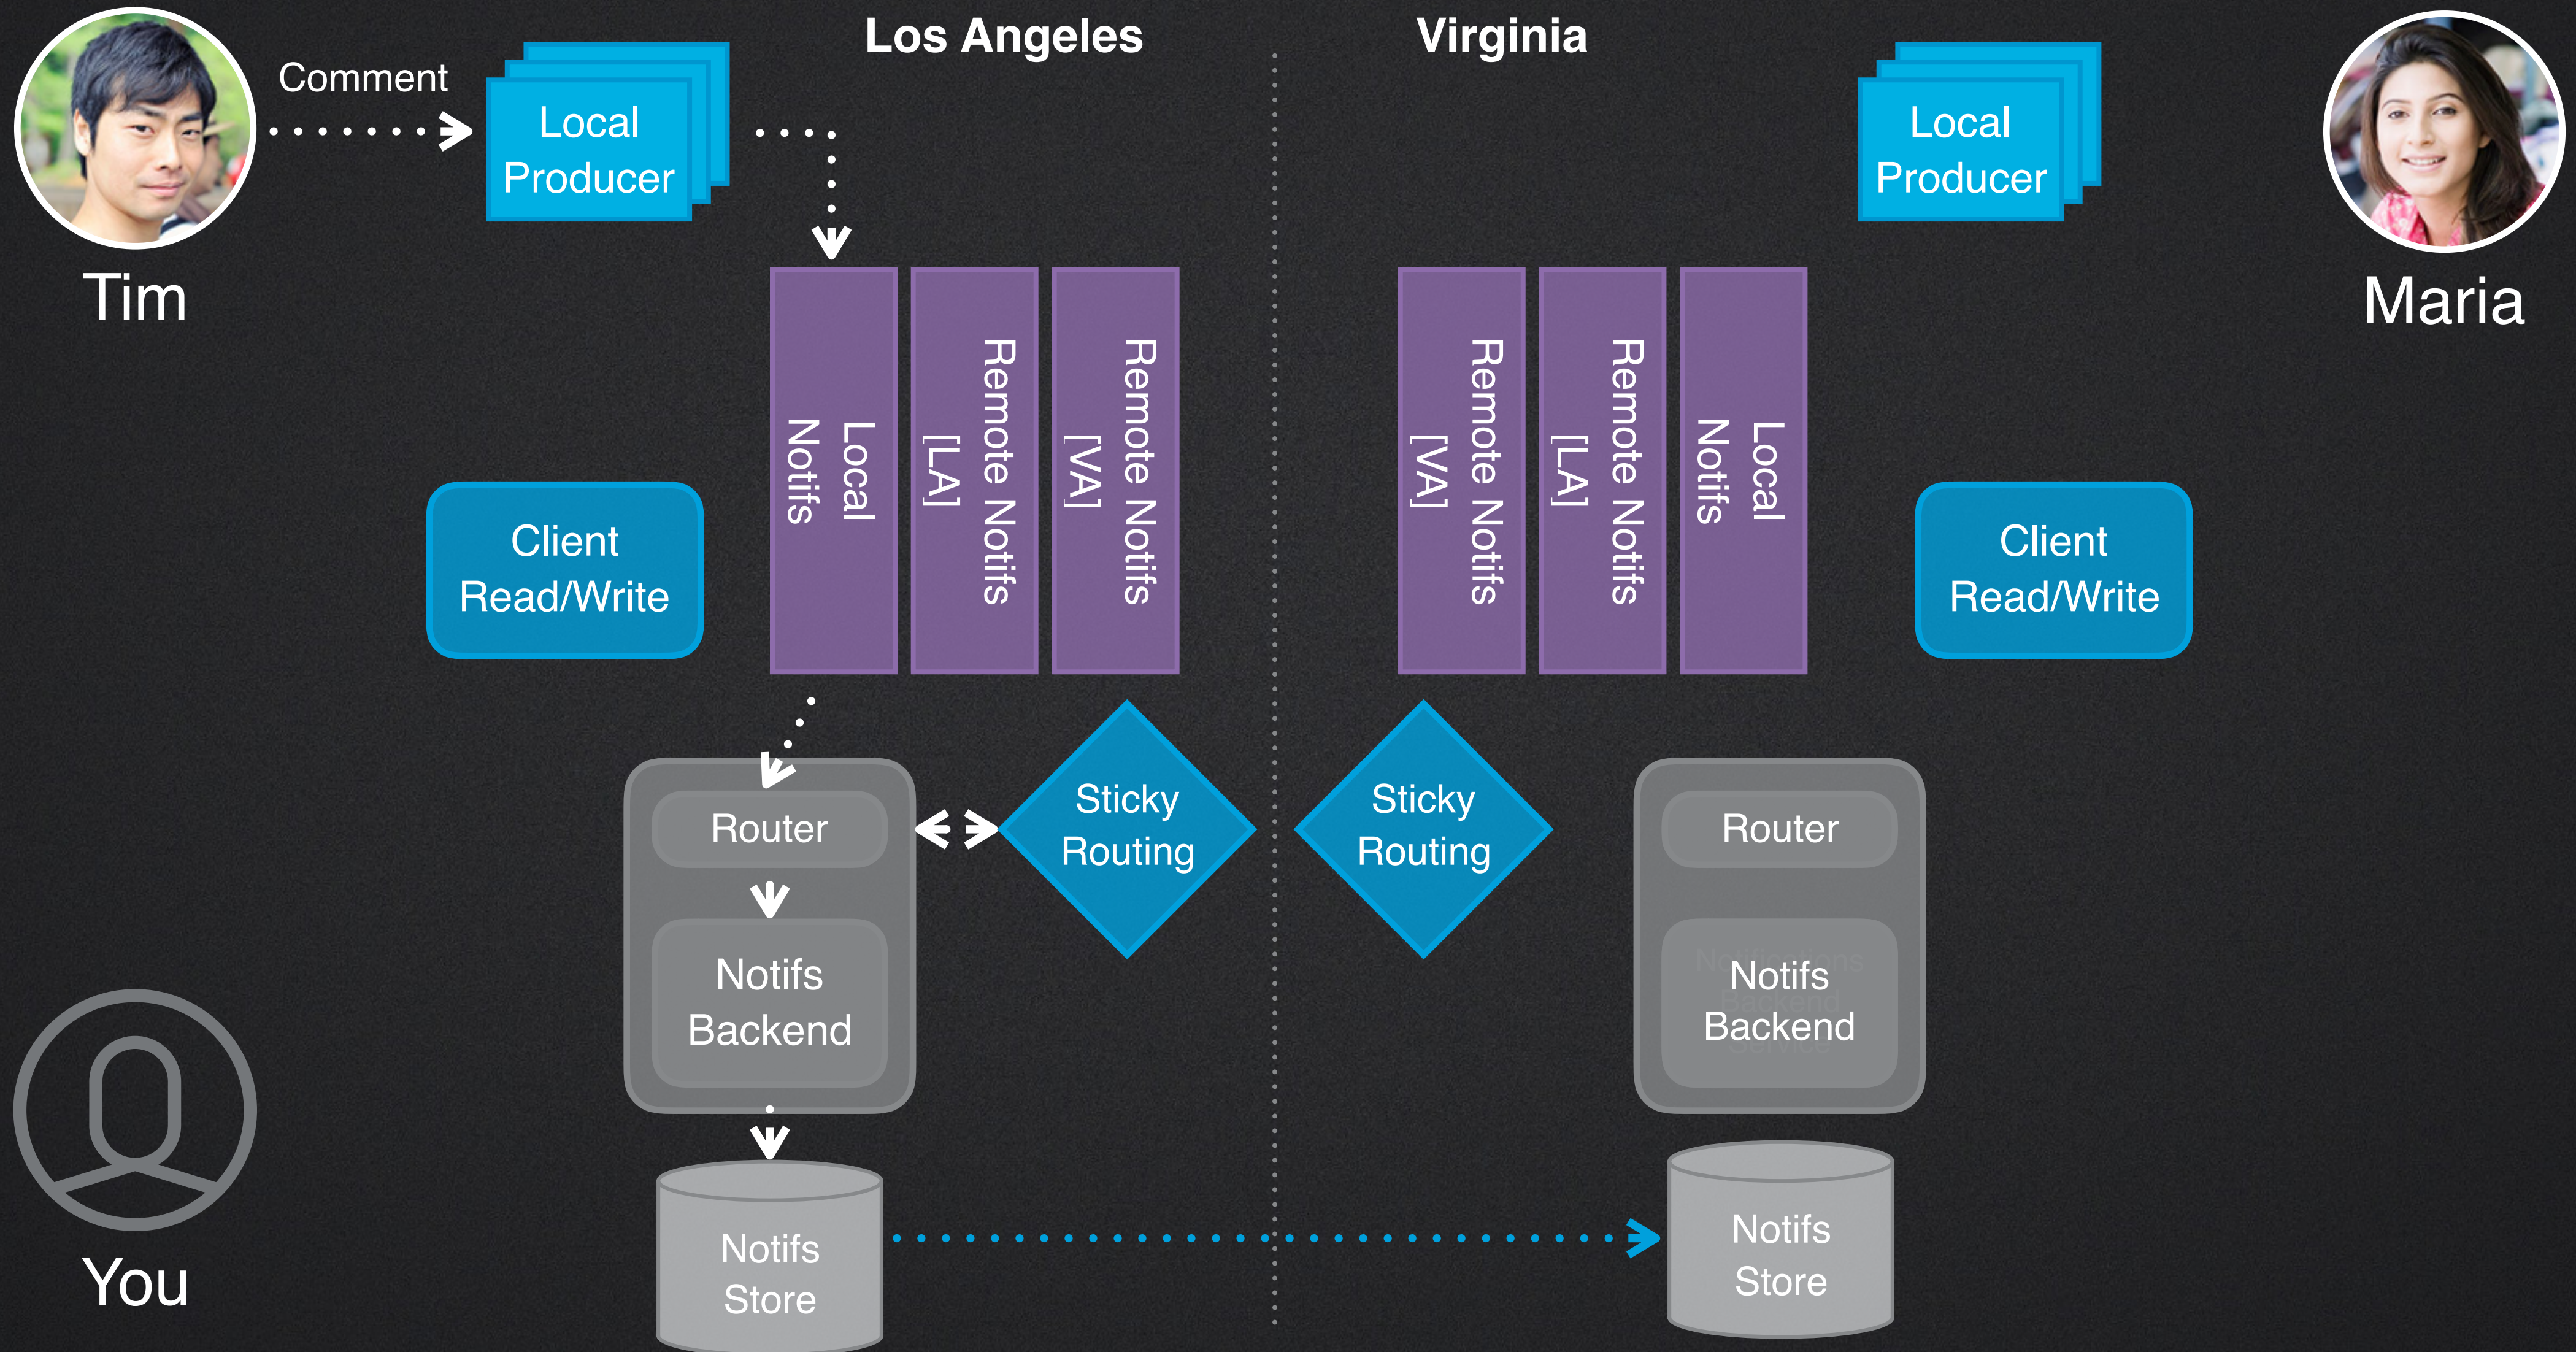
Maria



You

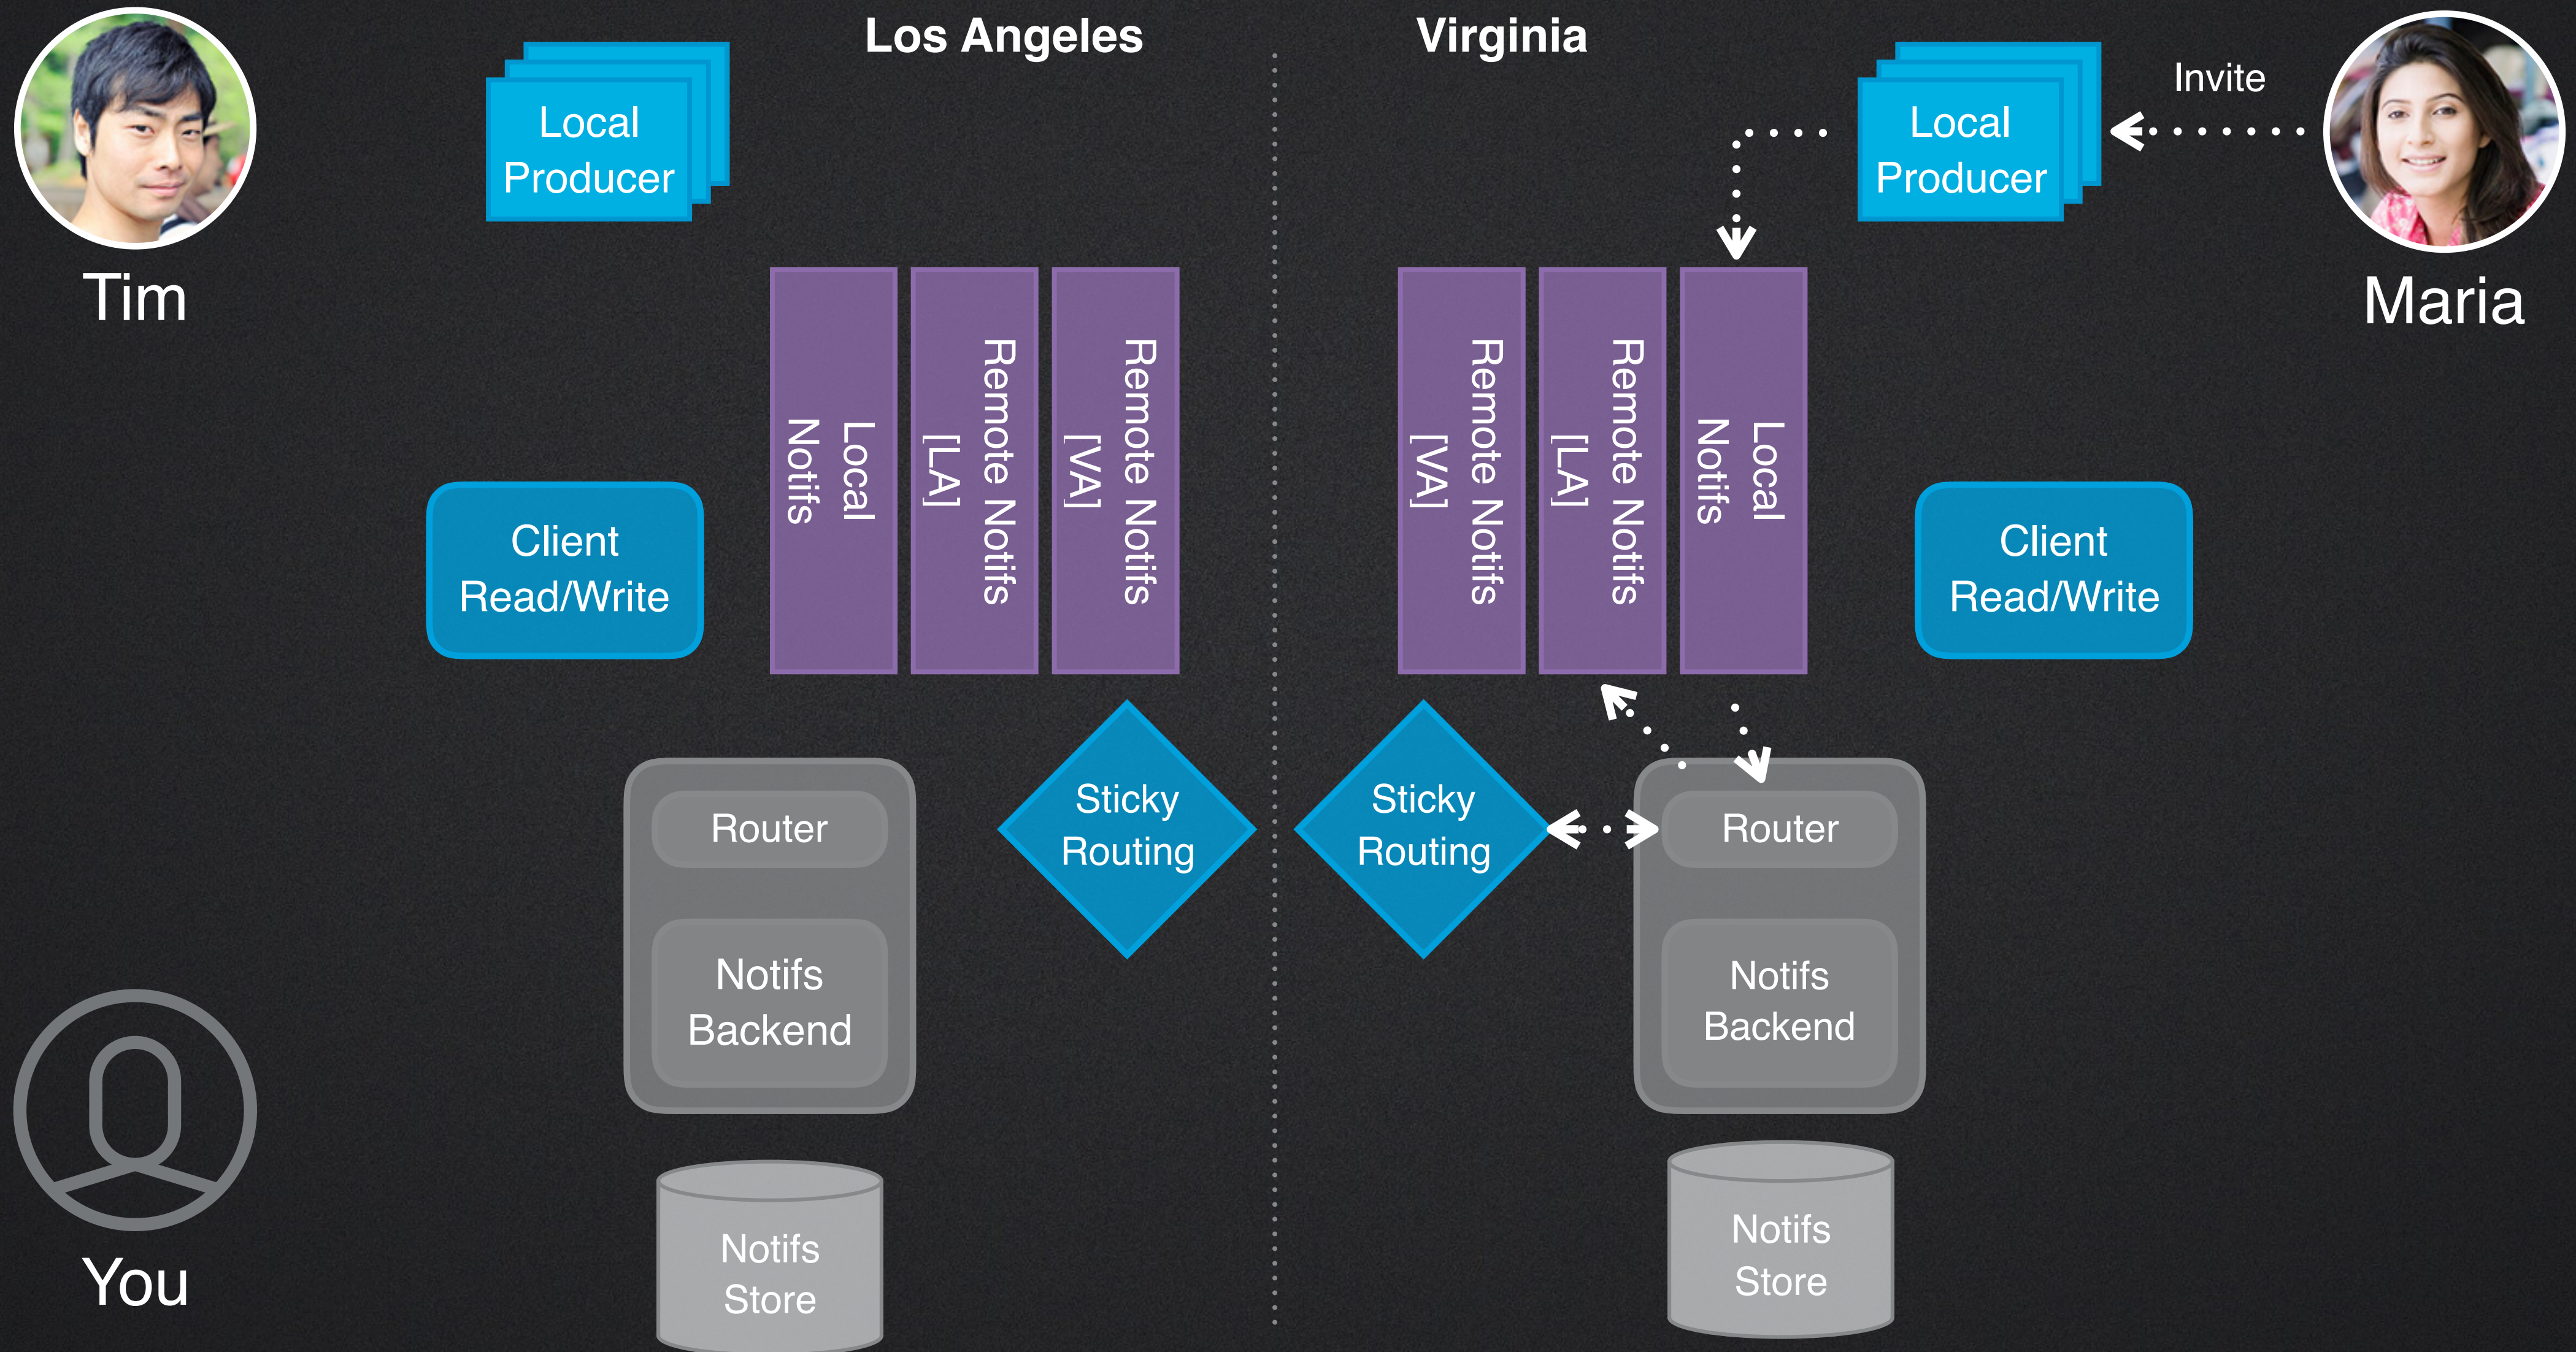


# Example



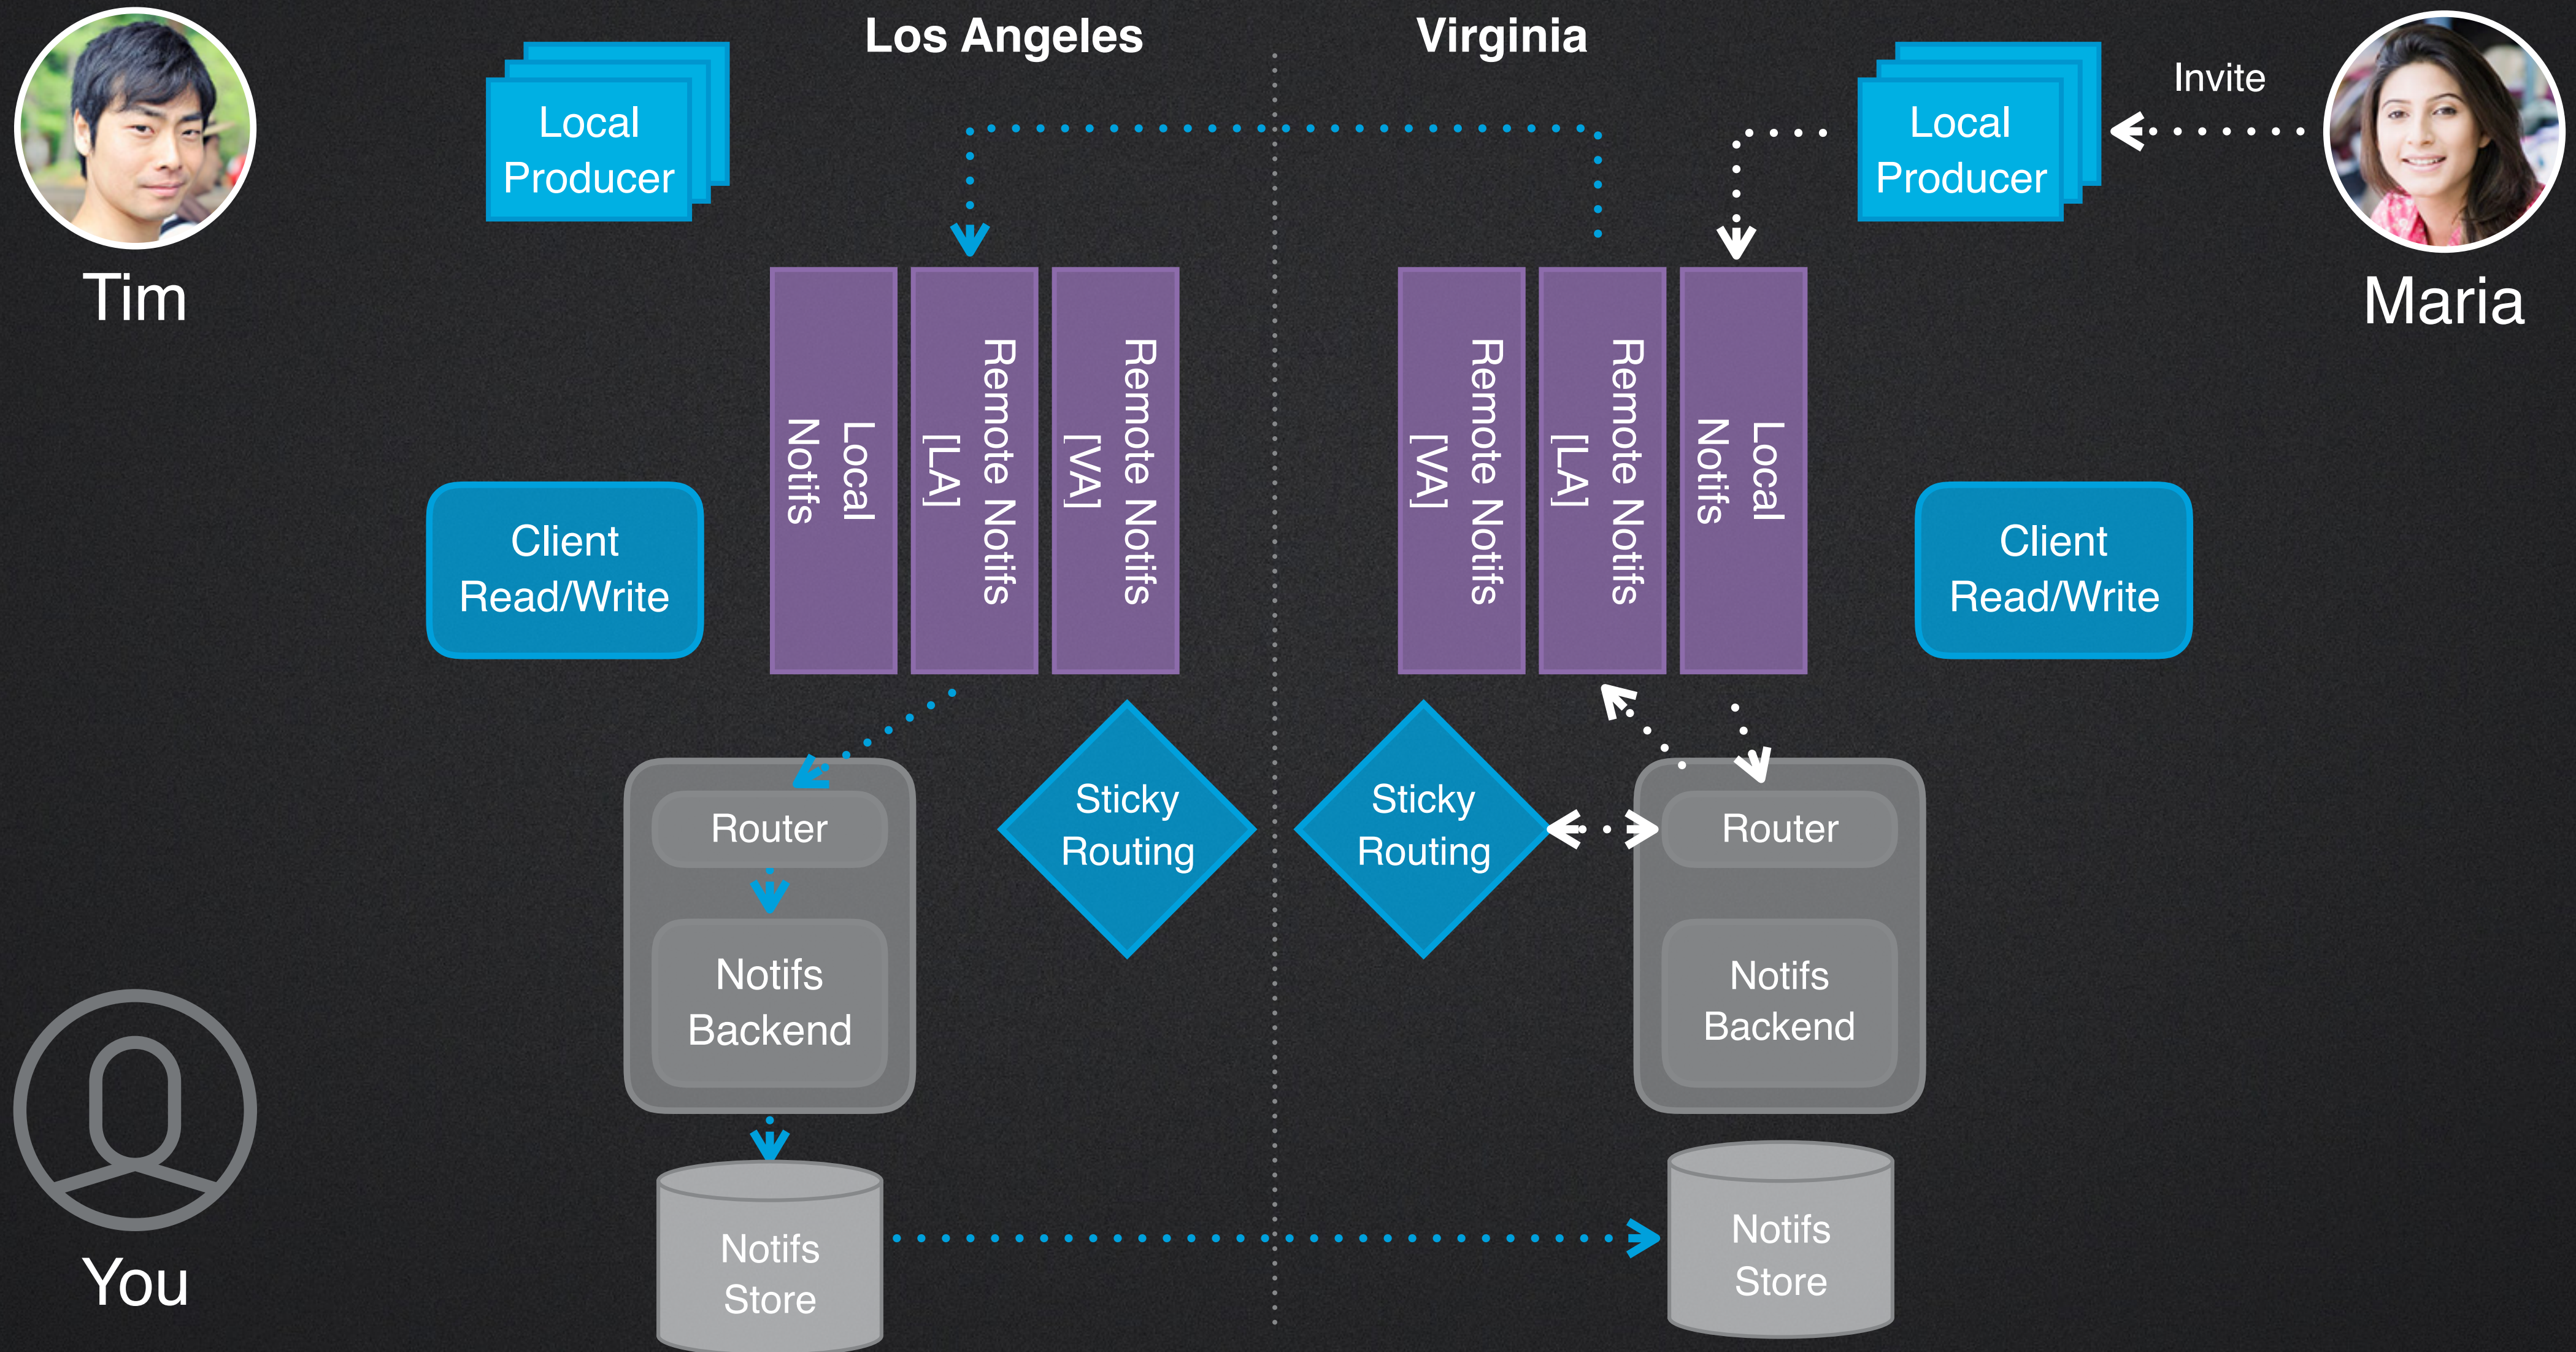


# Example



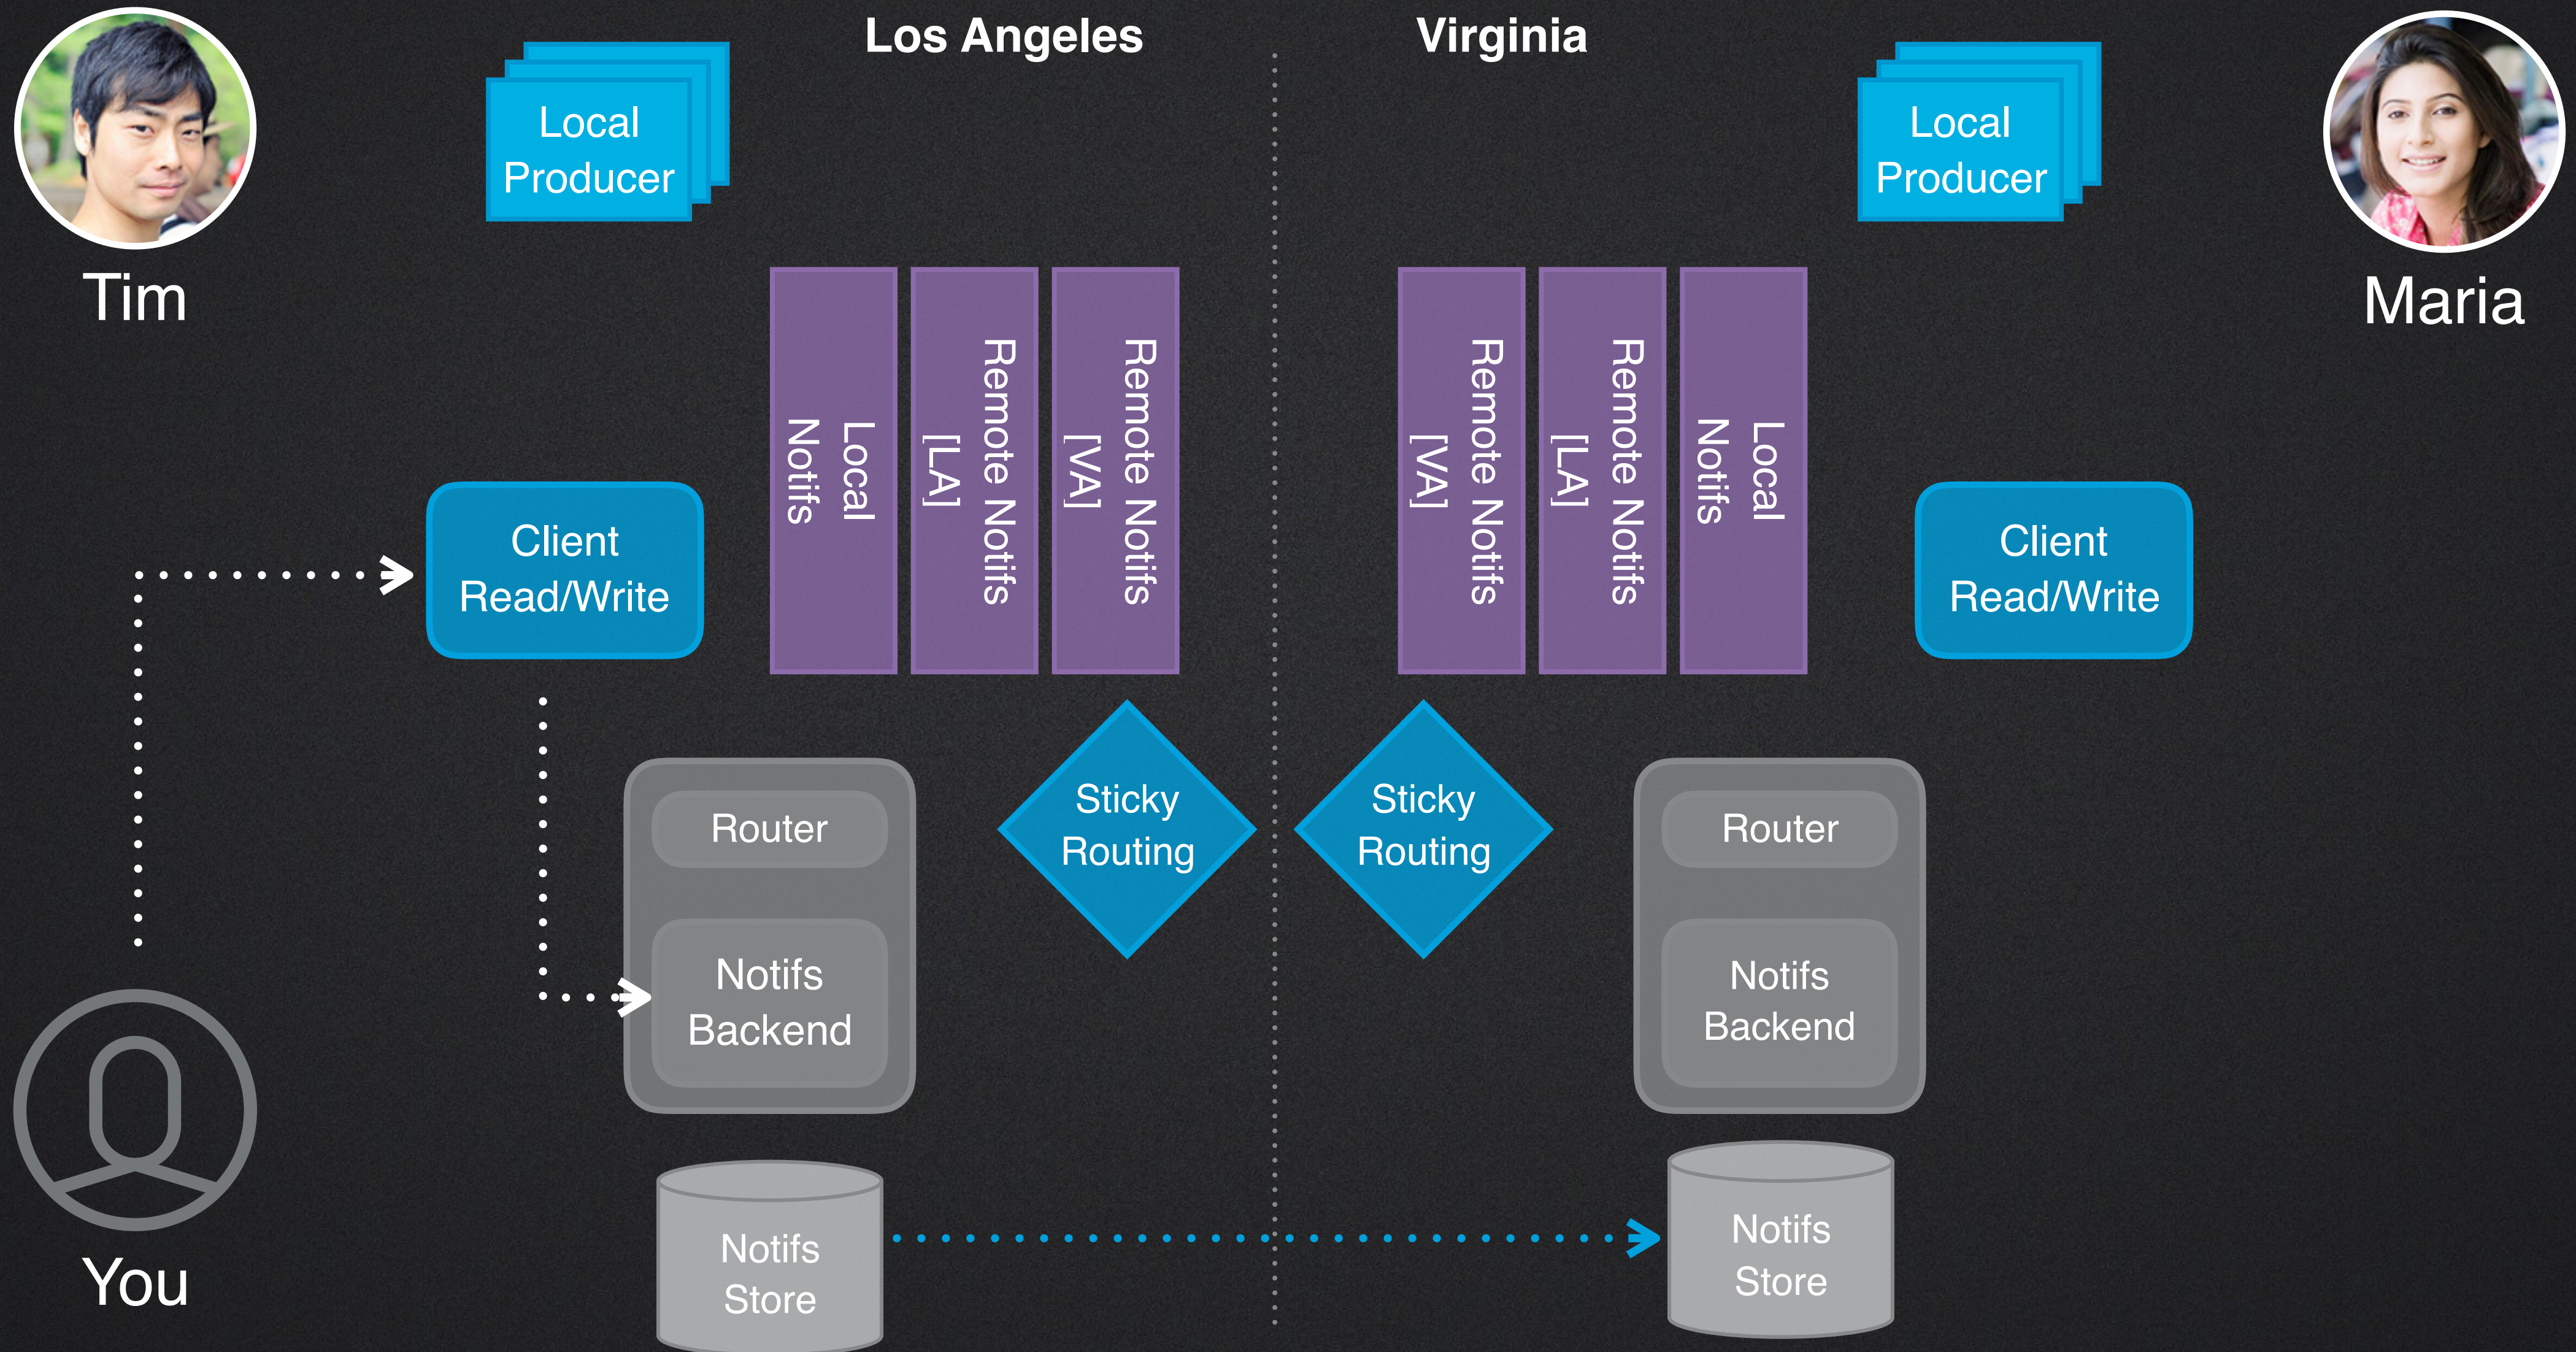


# Example





# Example



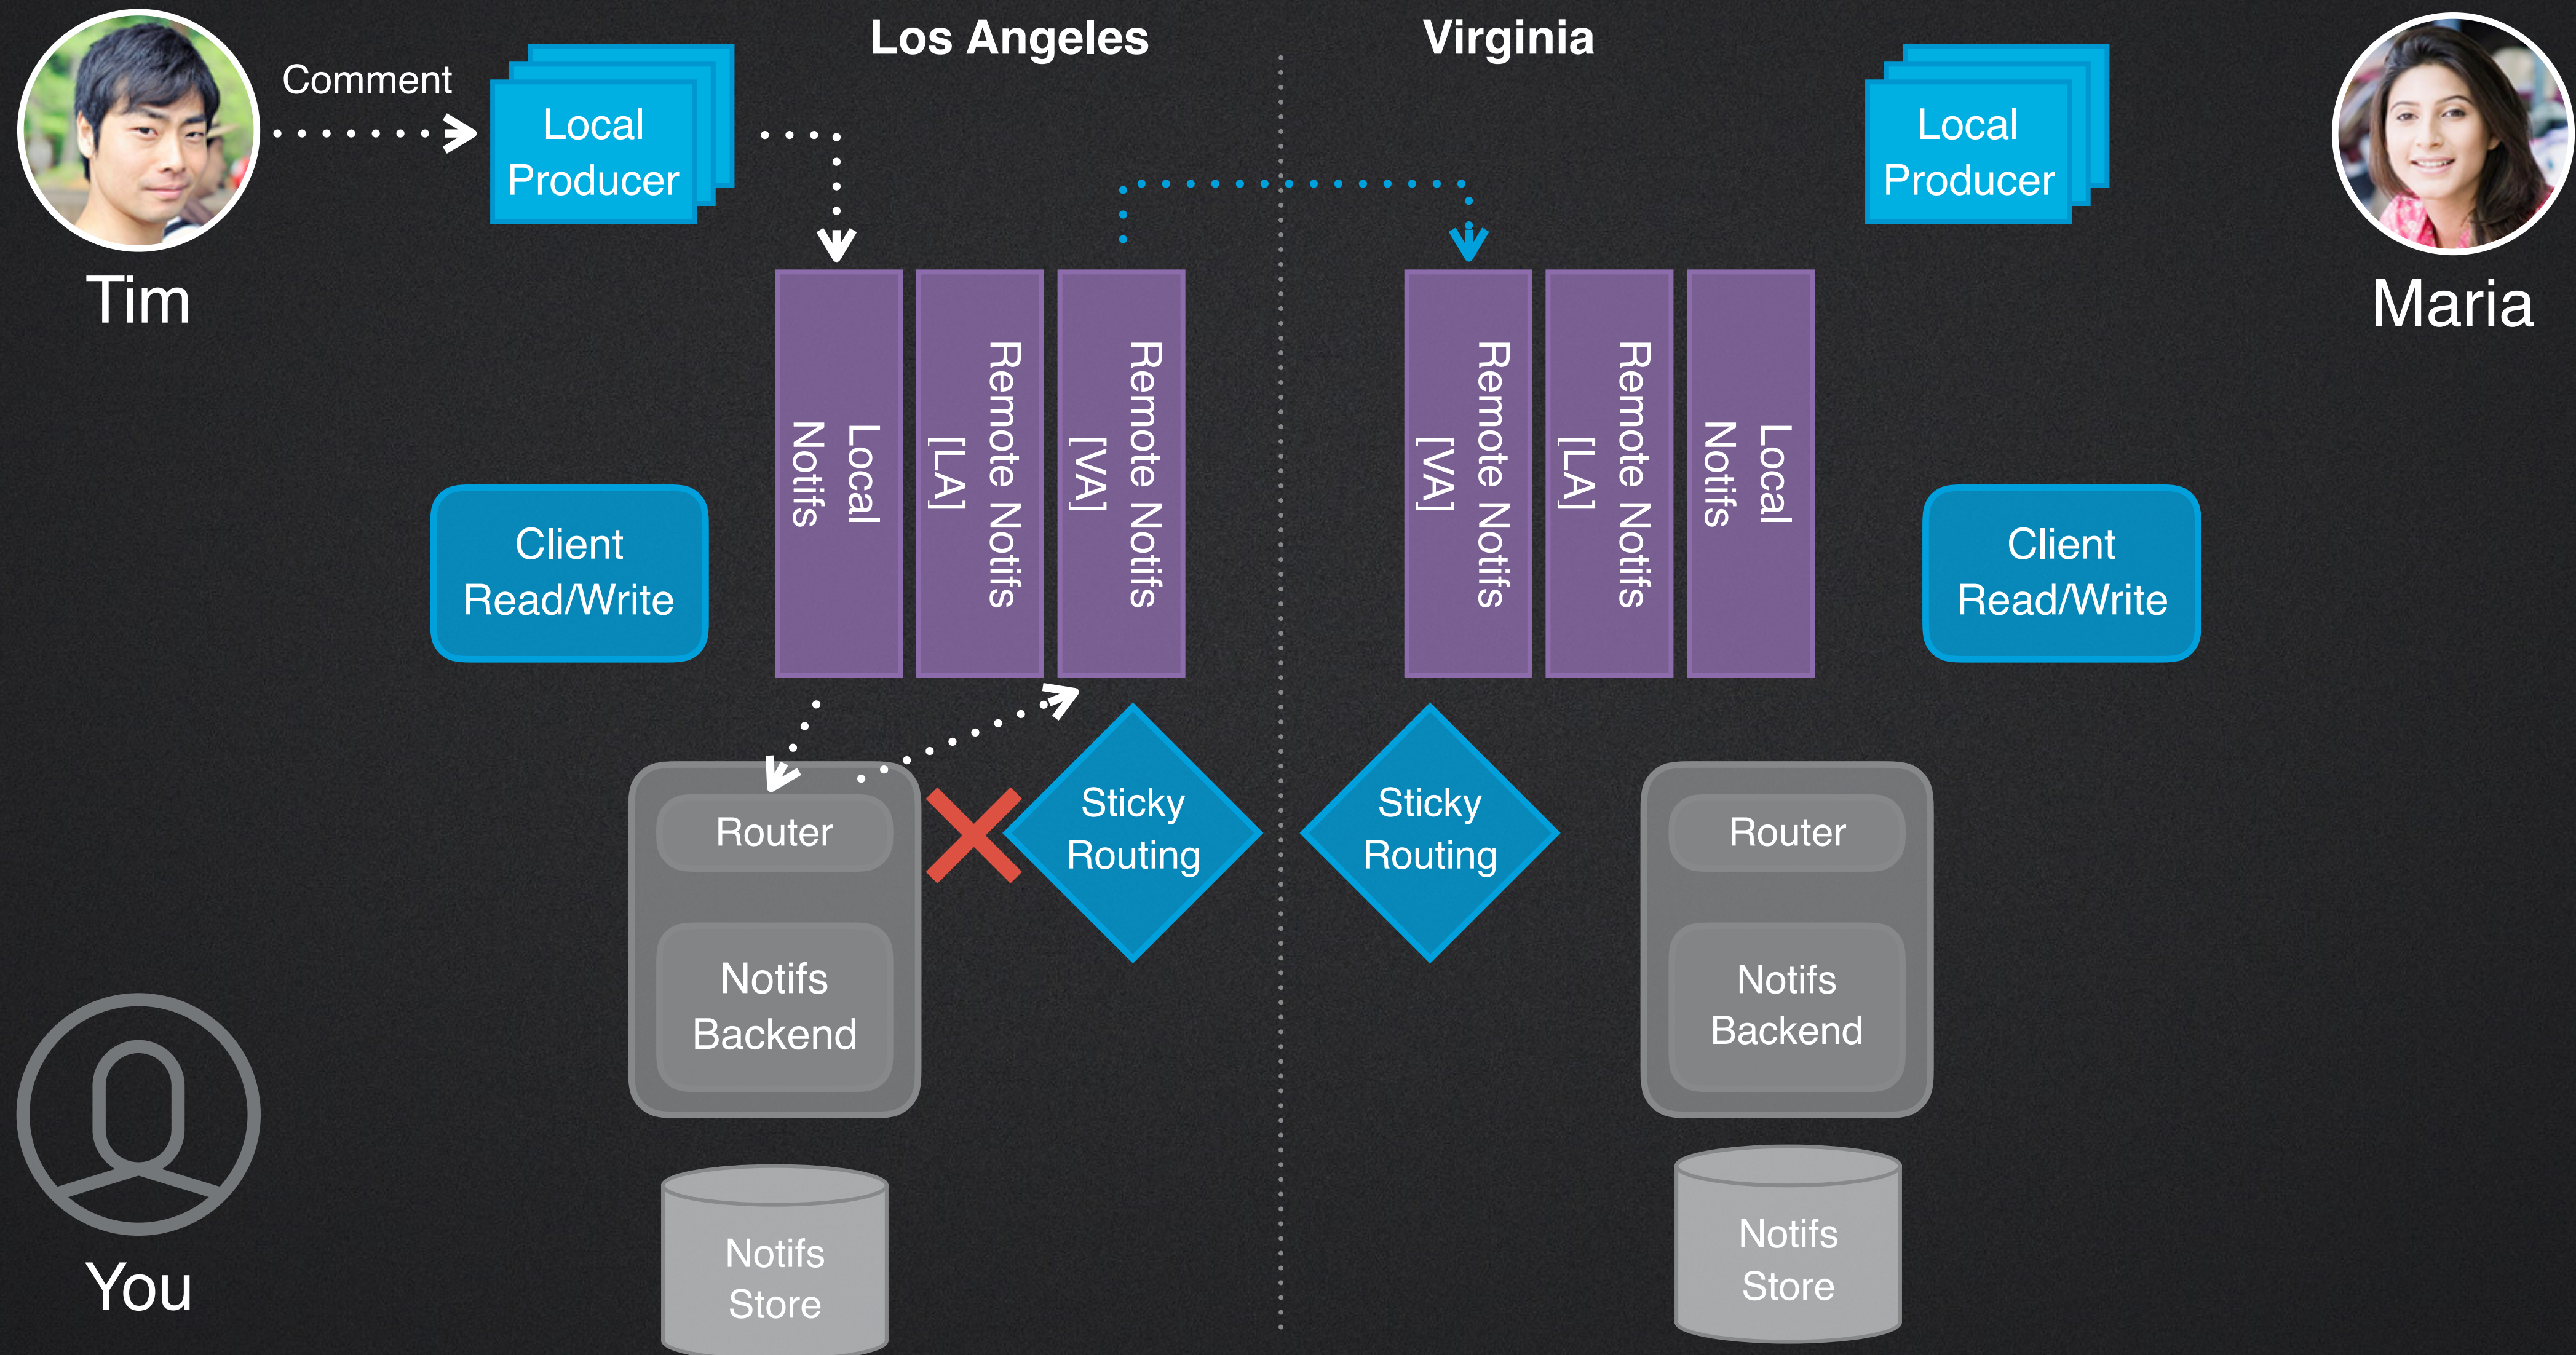


# Edge Cases!



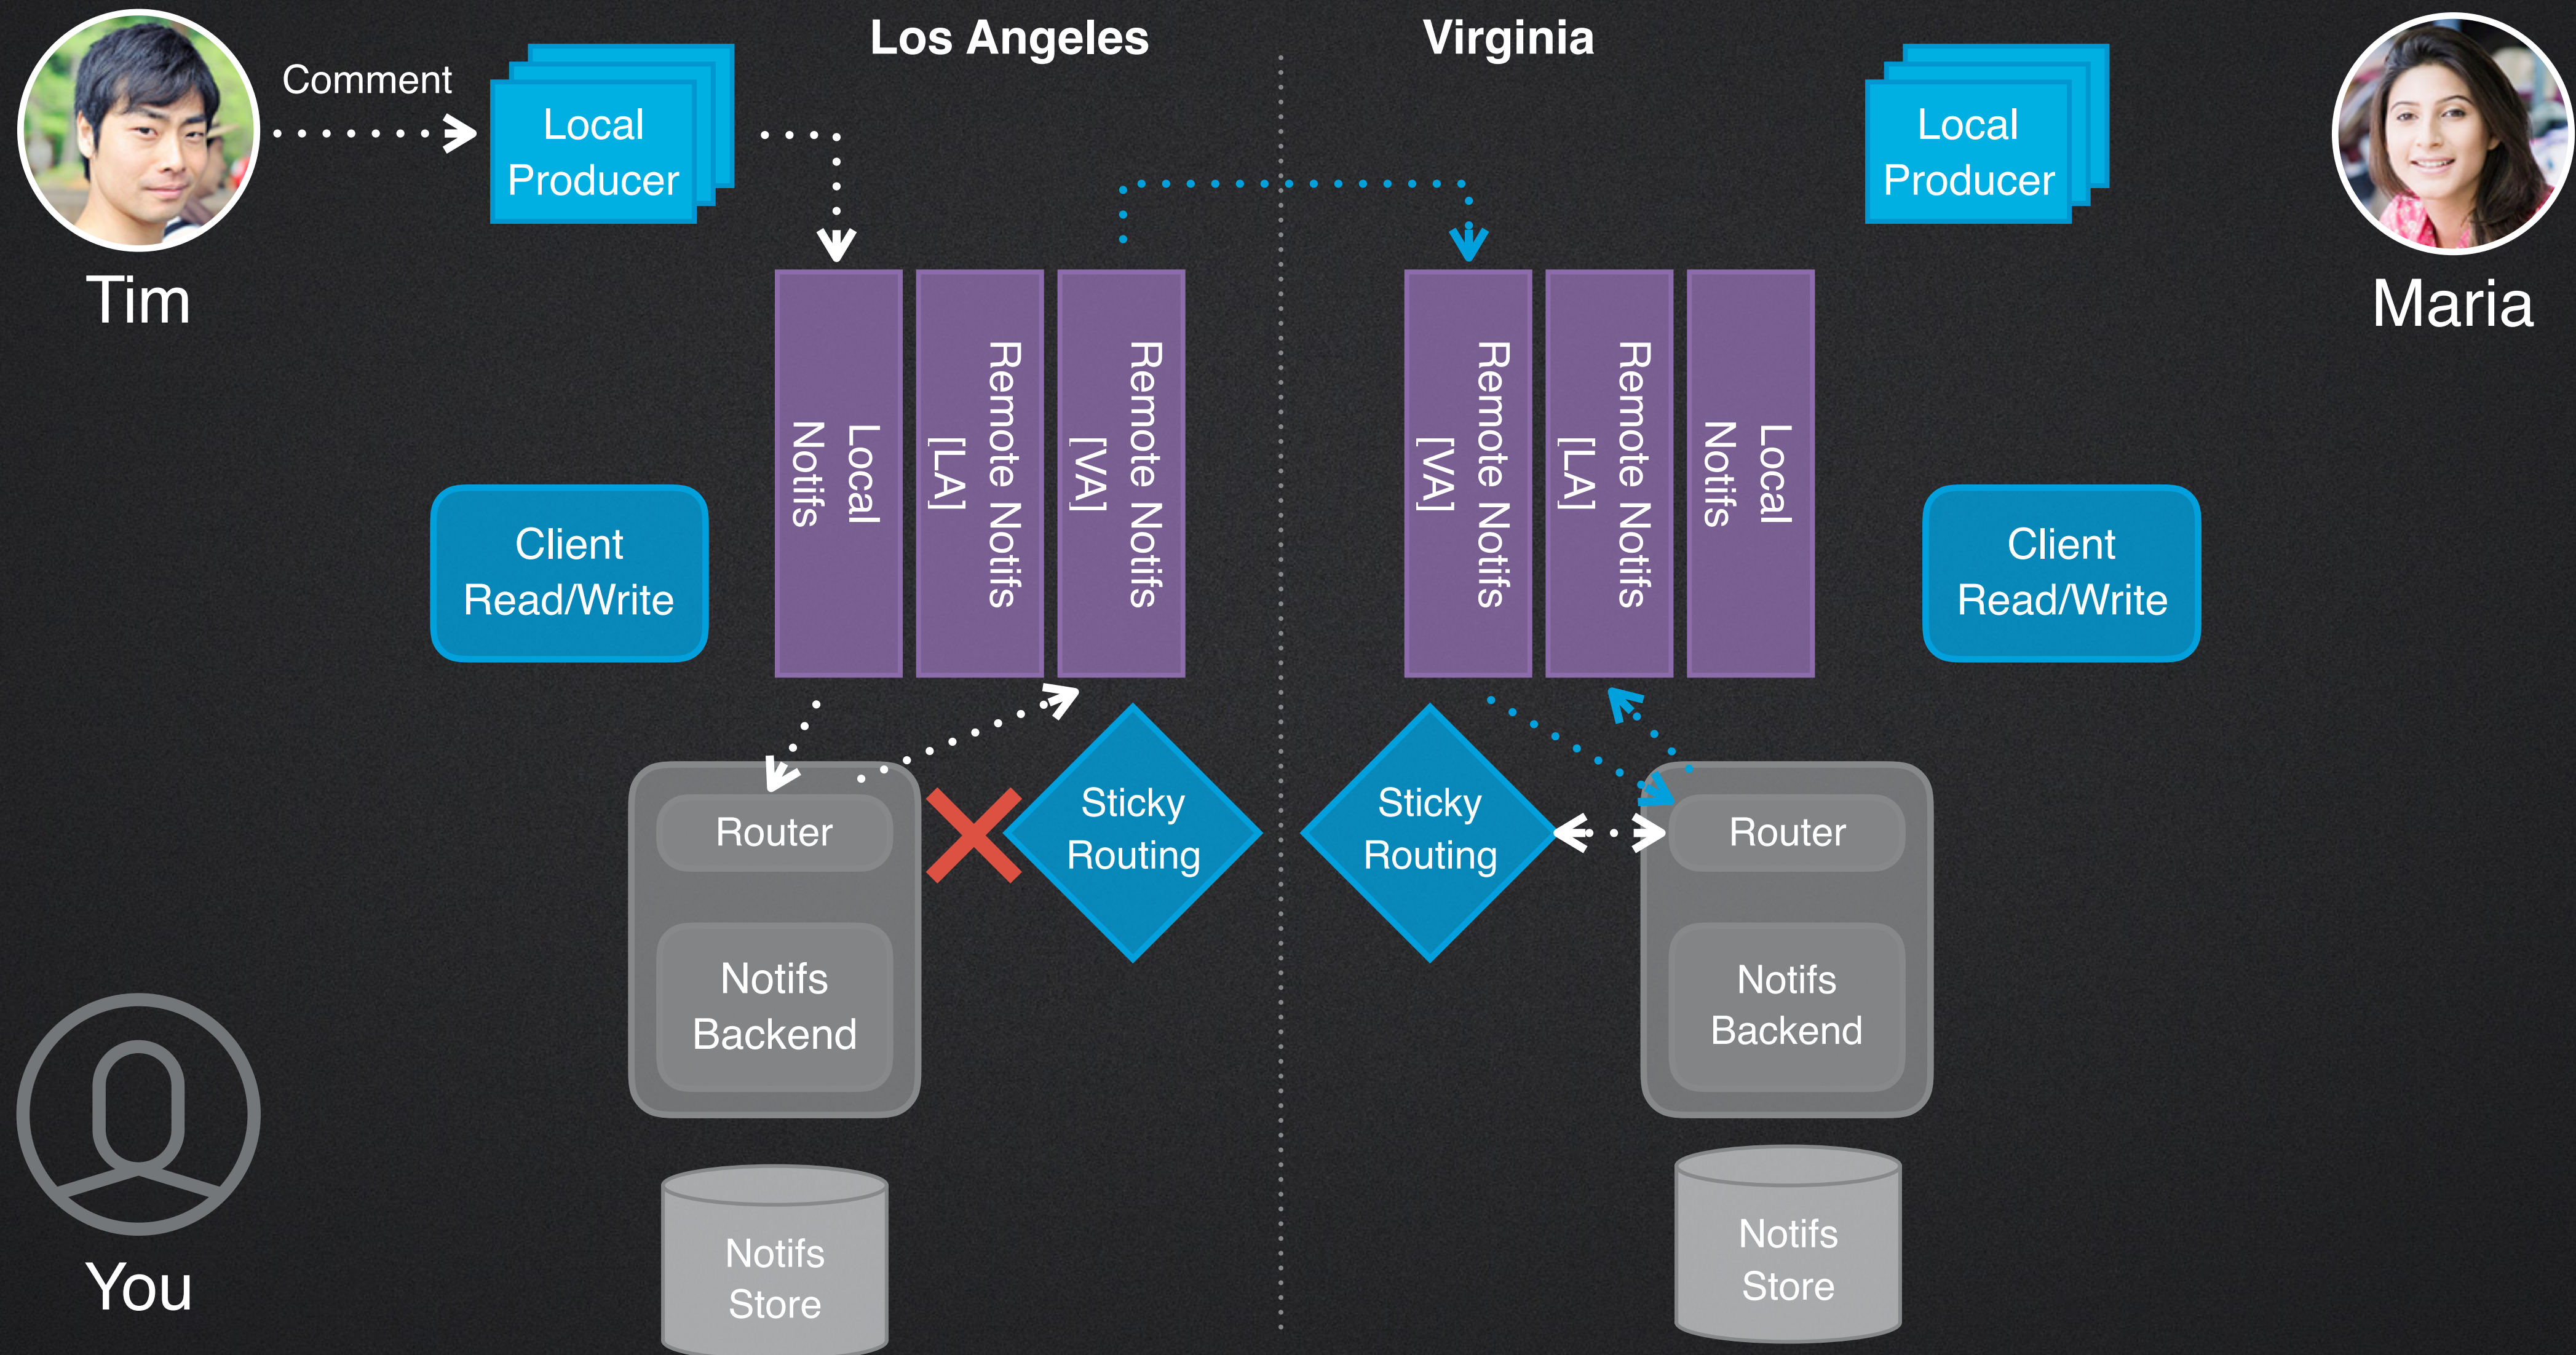


# Edge Case 1: Sticky routing is down



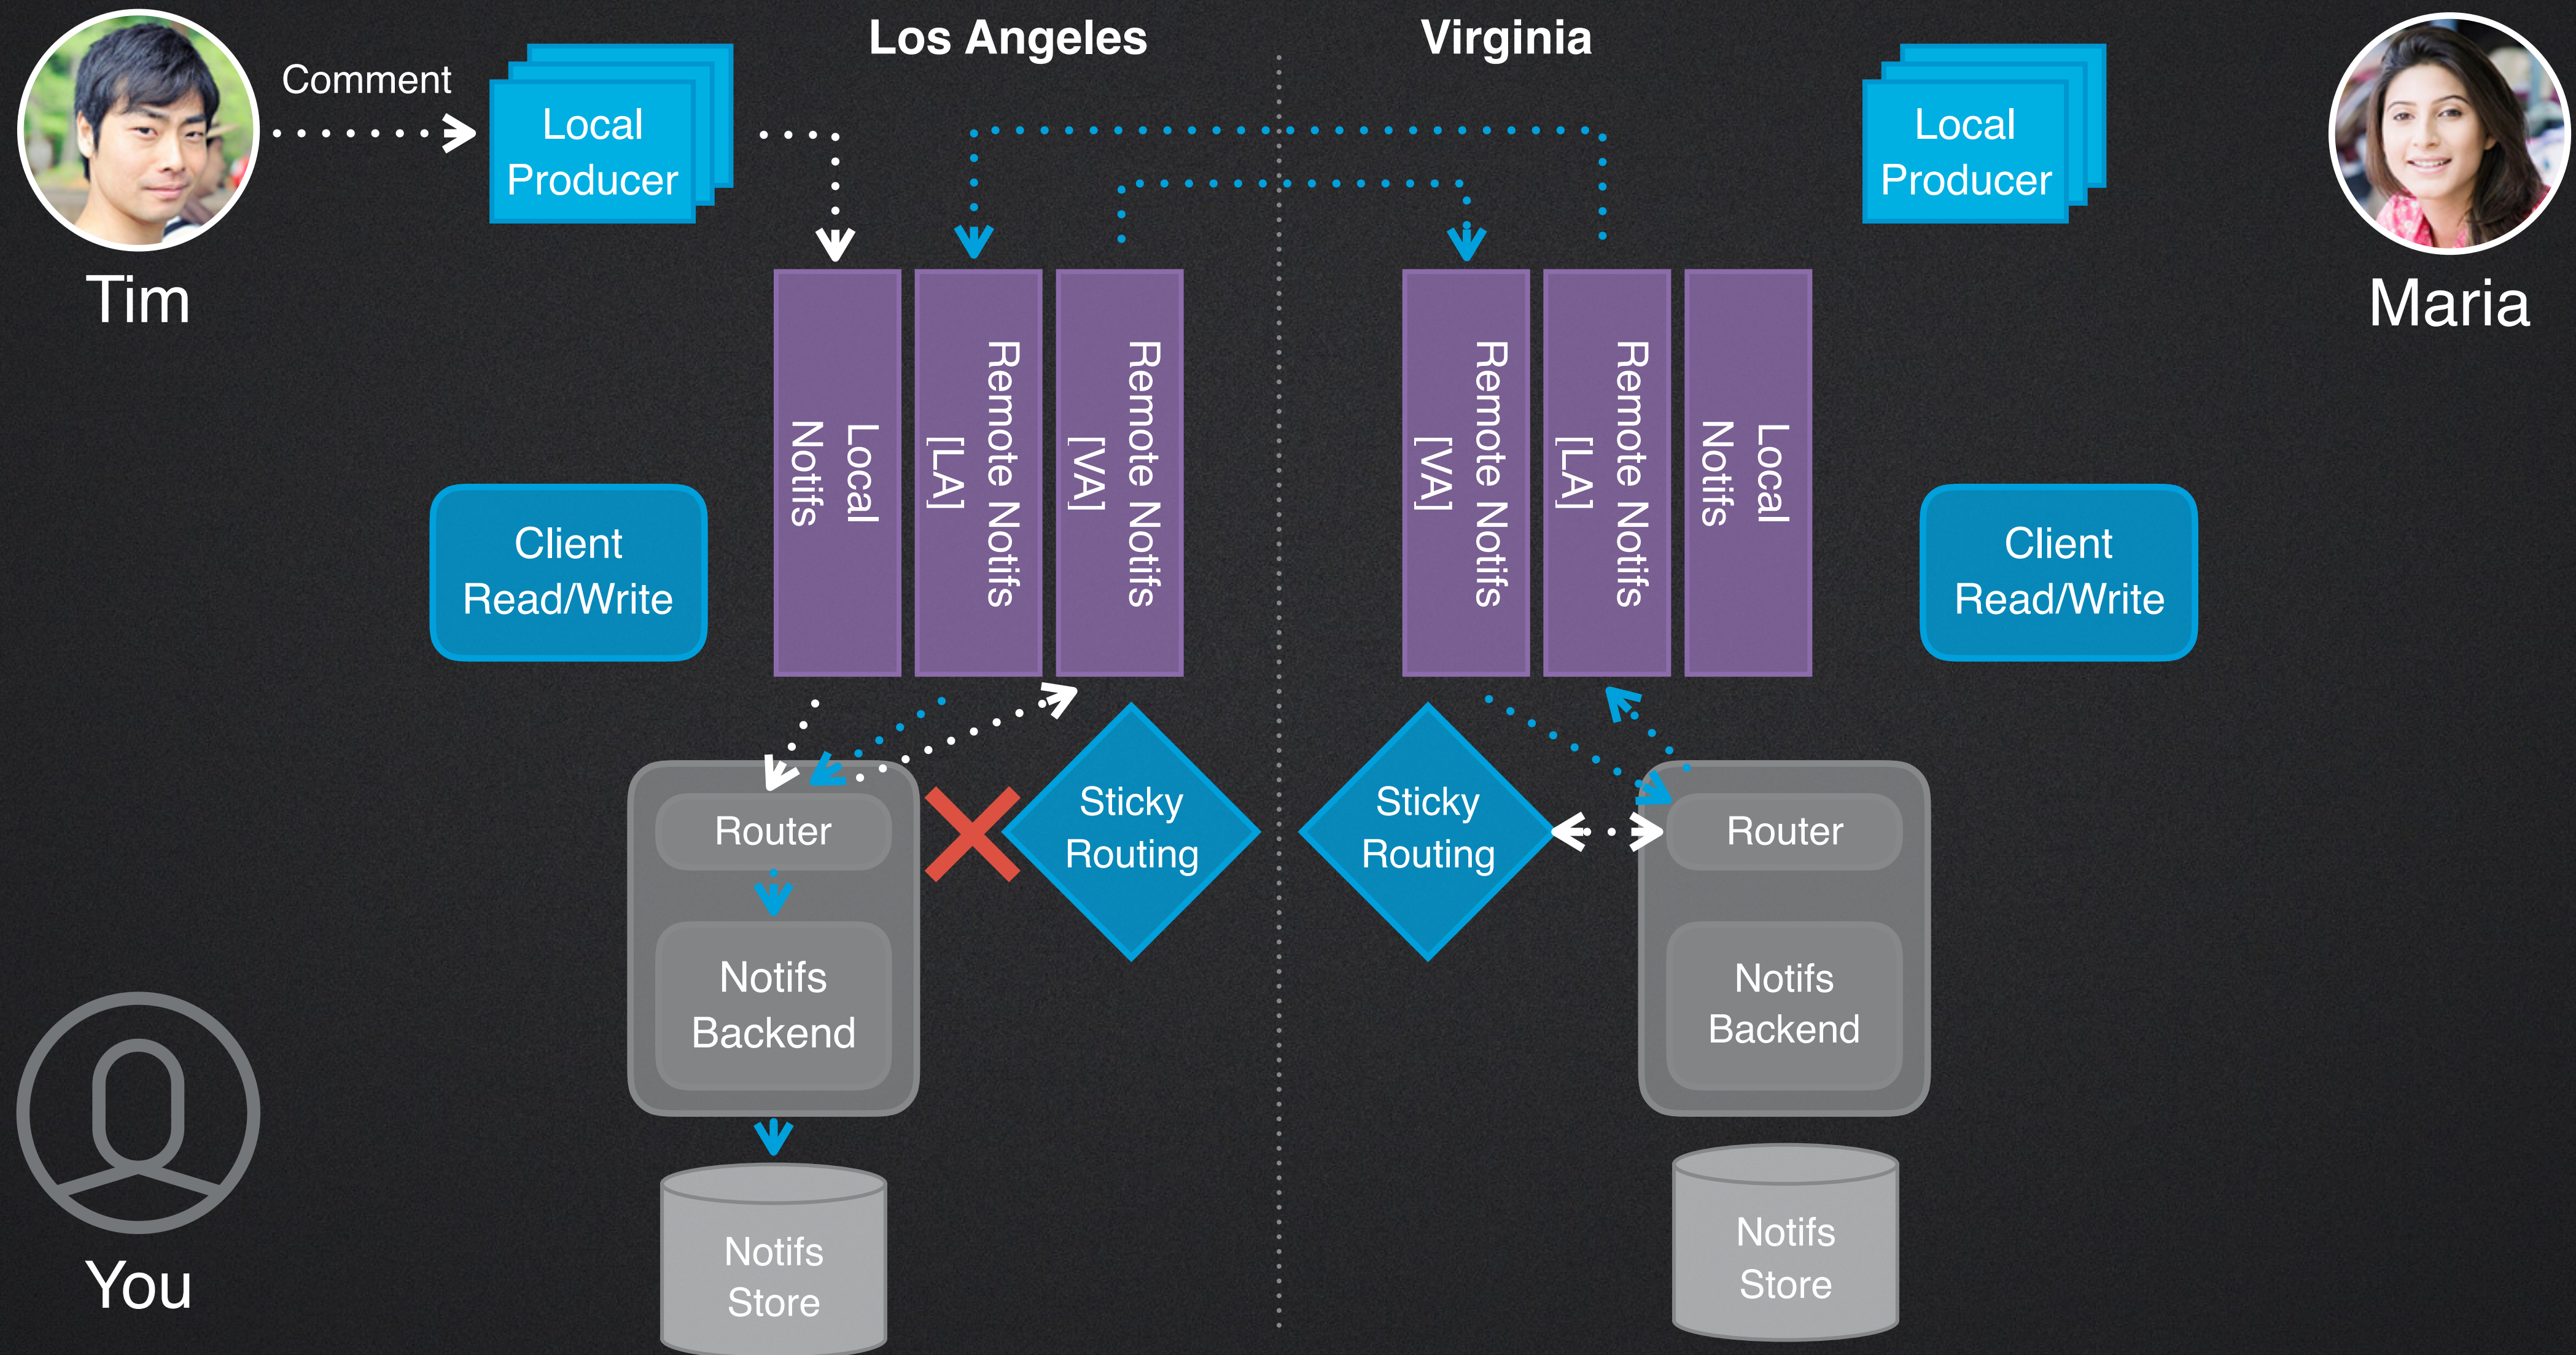


# Edge Case 1: Sticky routing is down



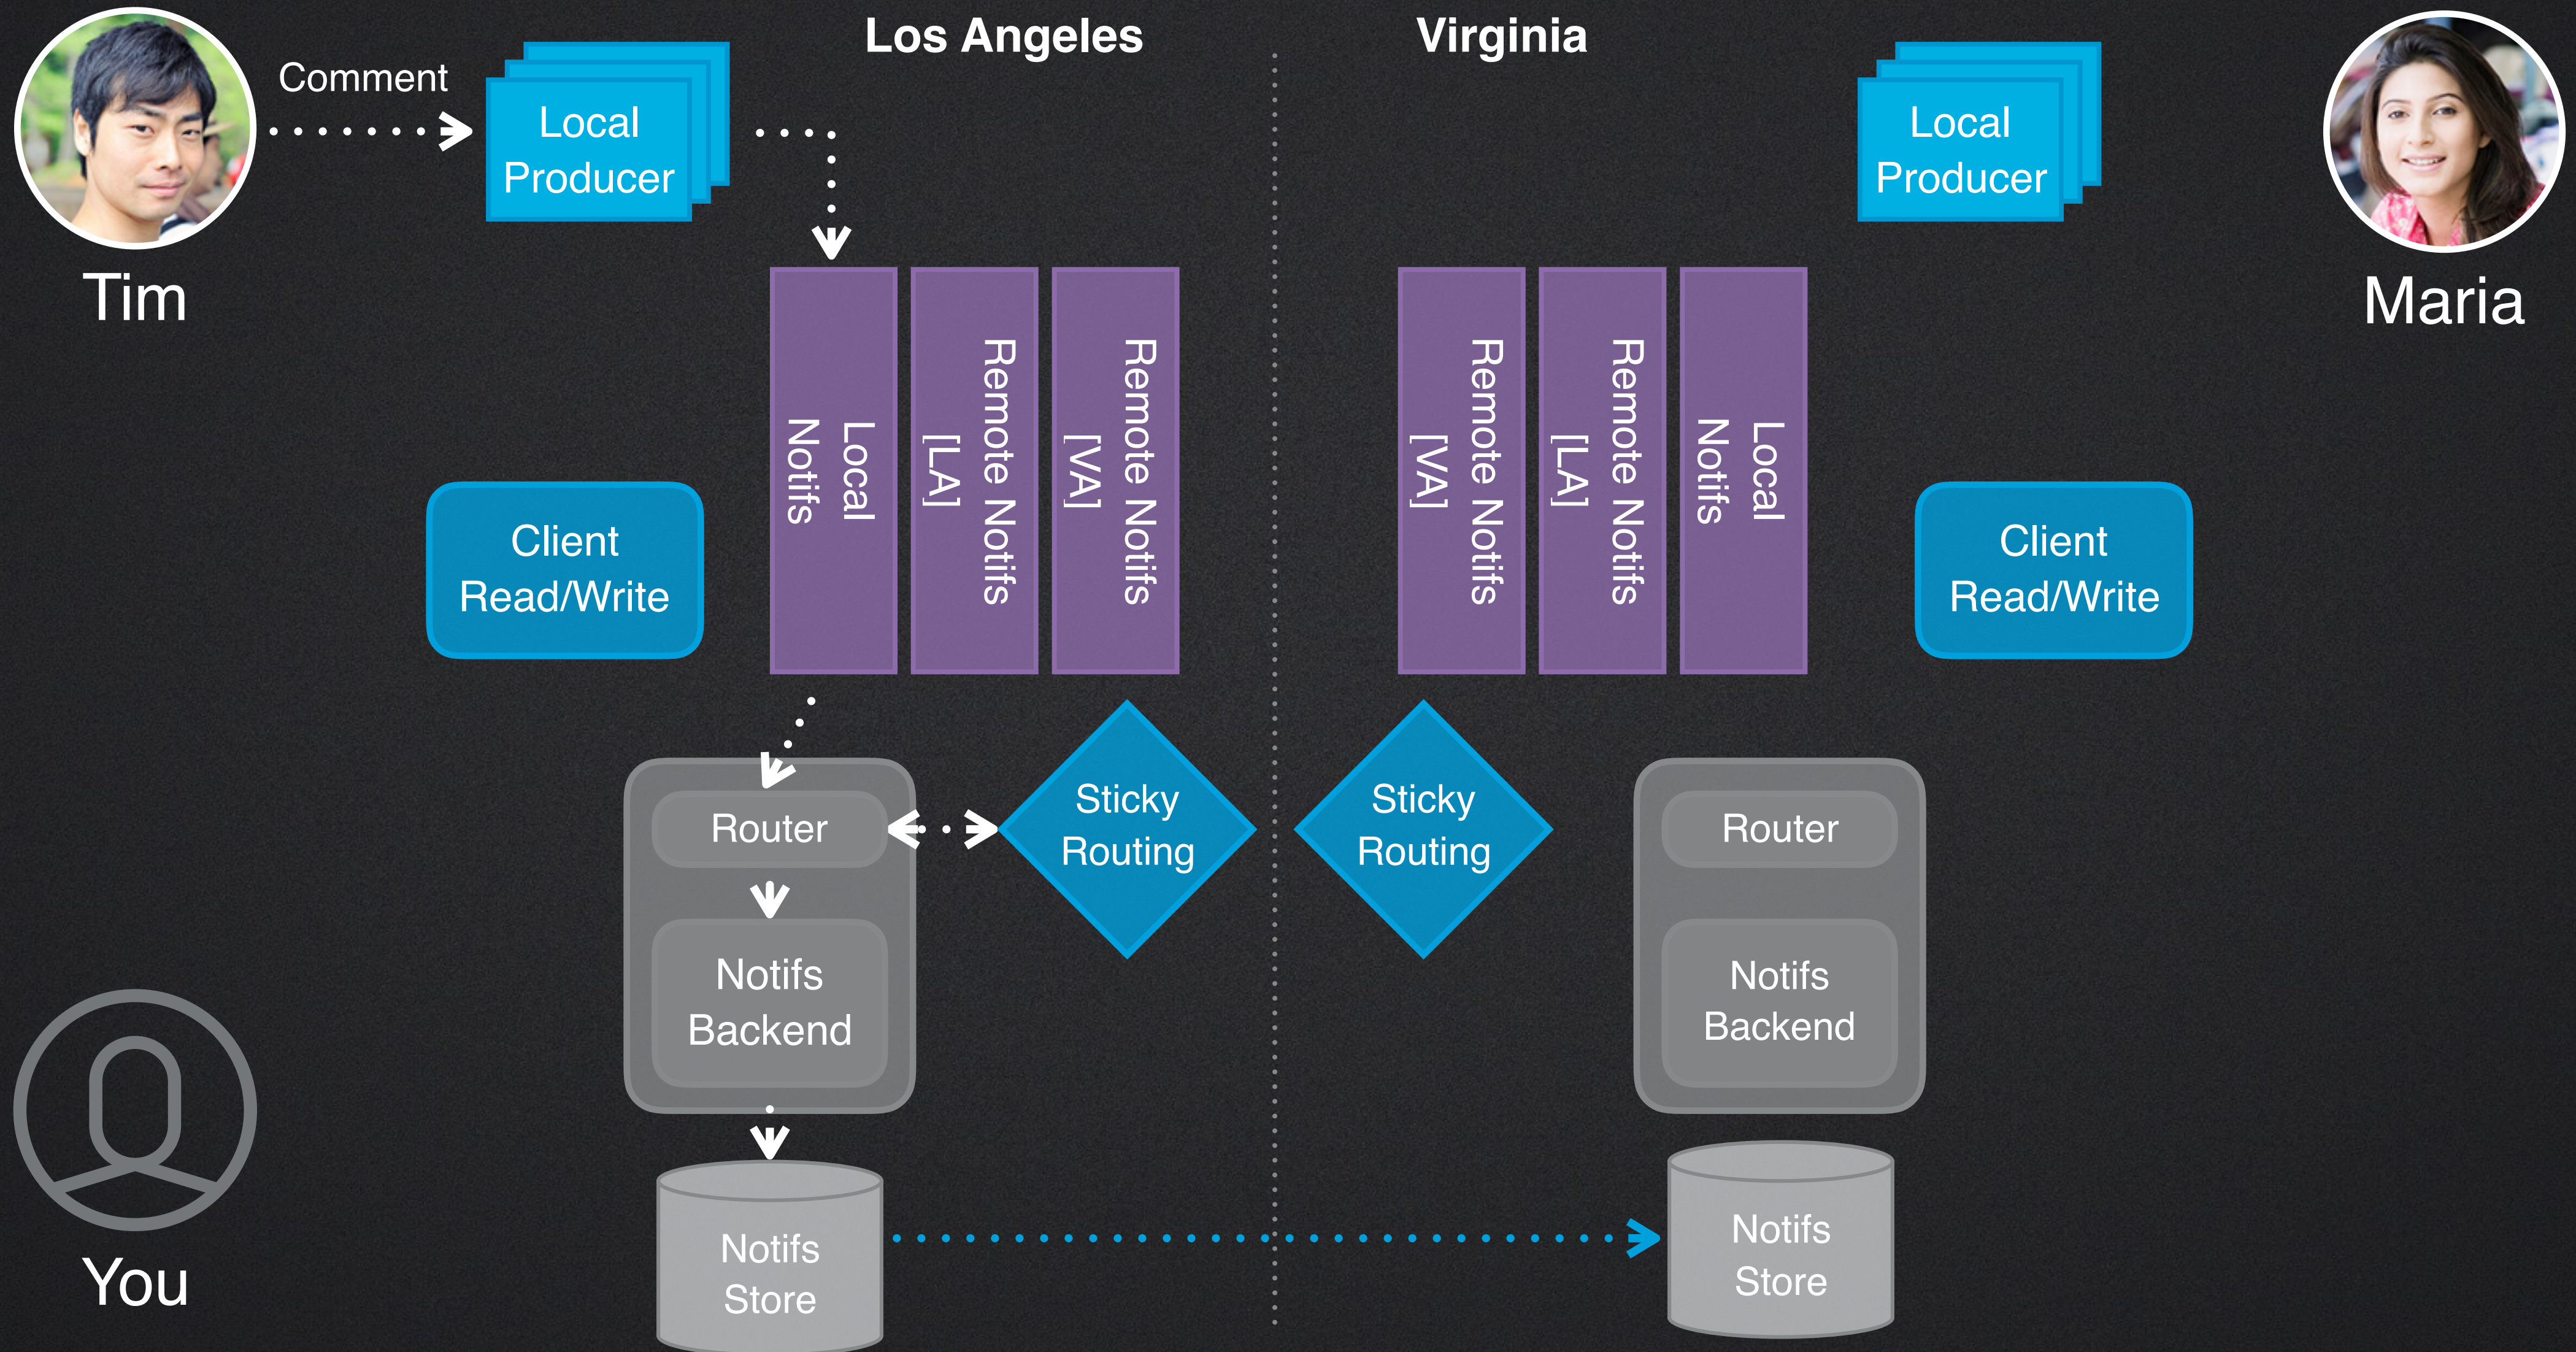


# Edge Case 1: Sticky routing is down



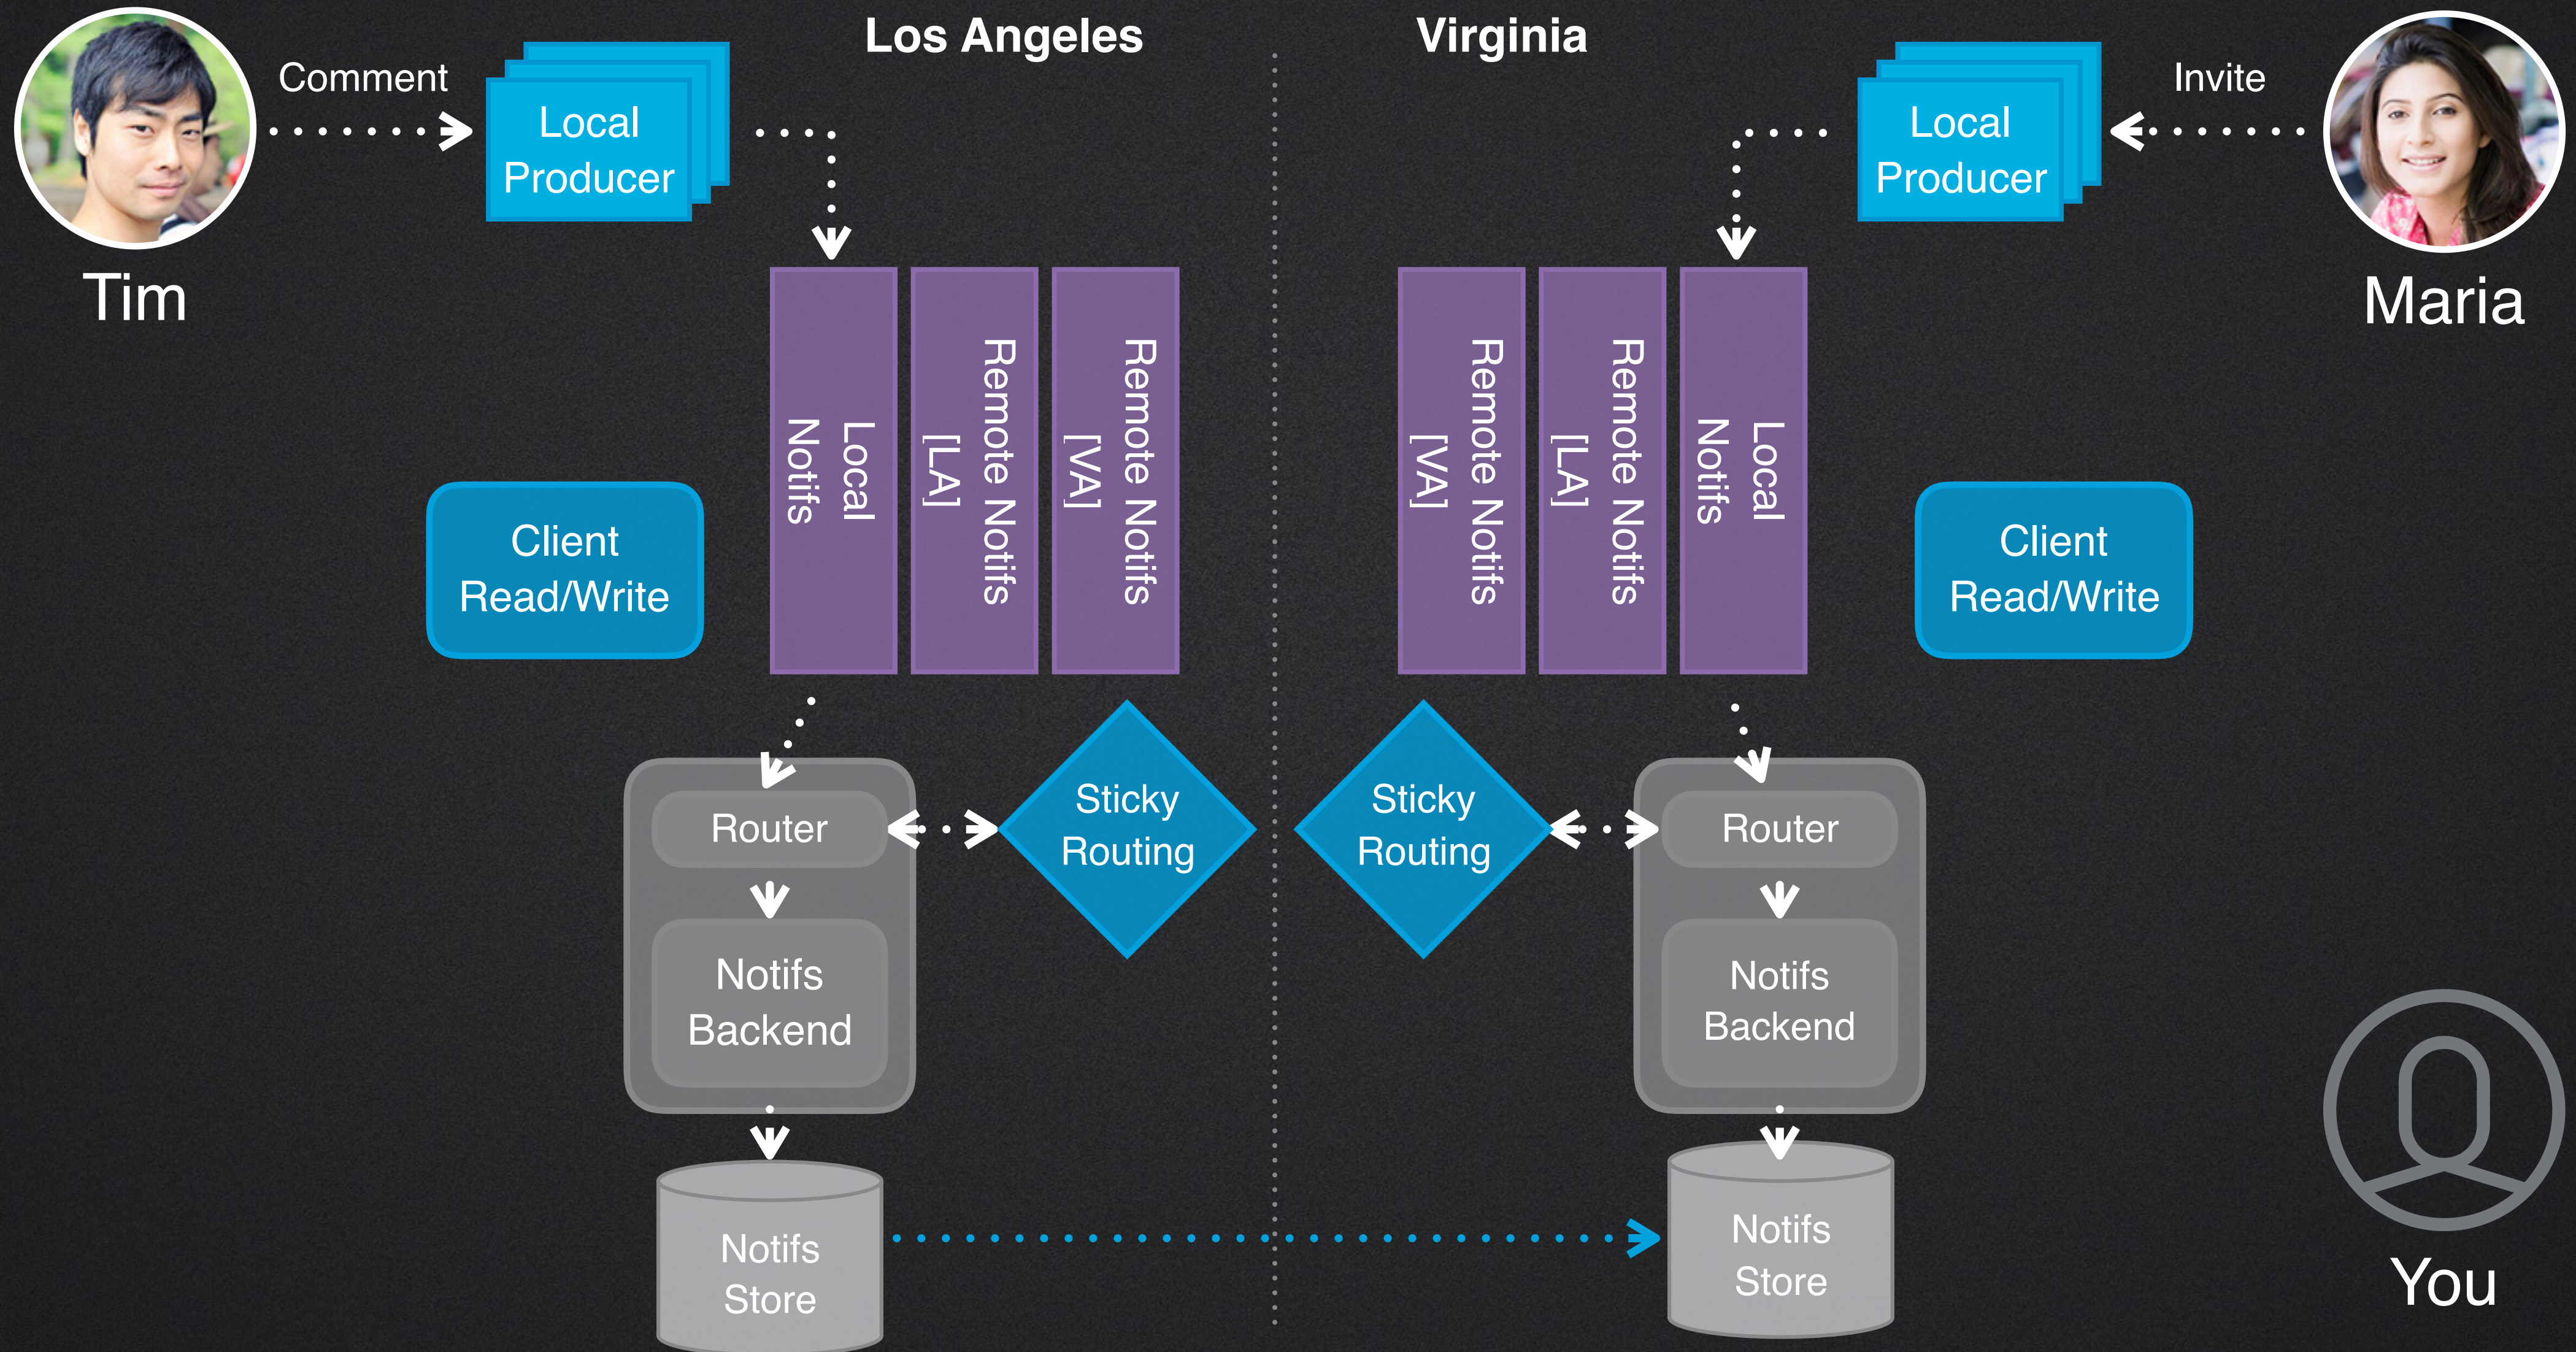


# Edge Case 2: Traffic shift race



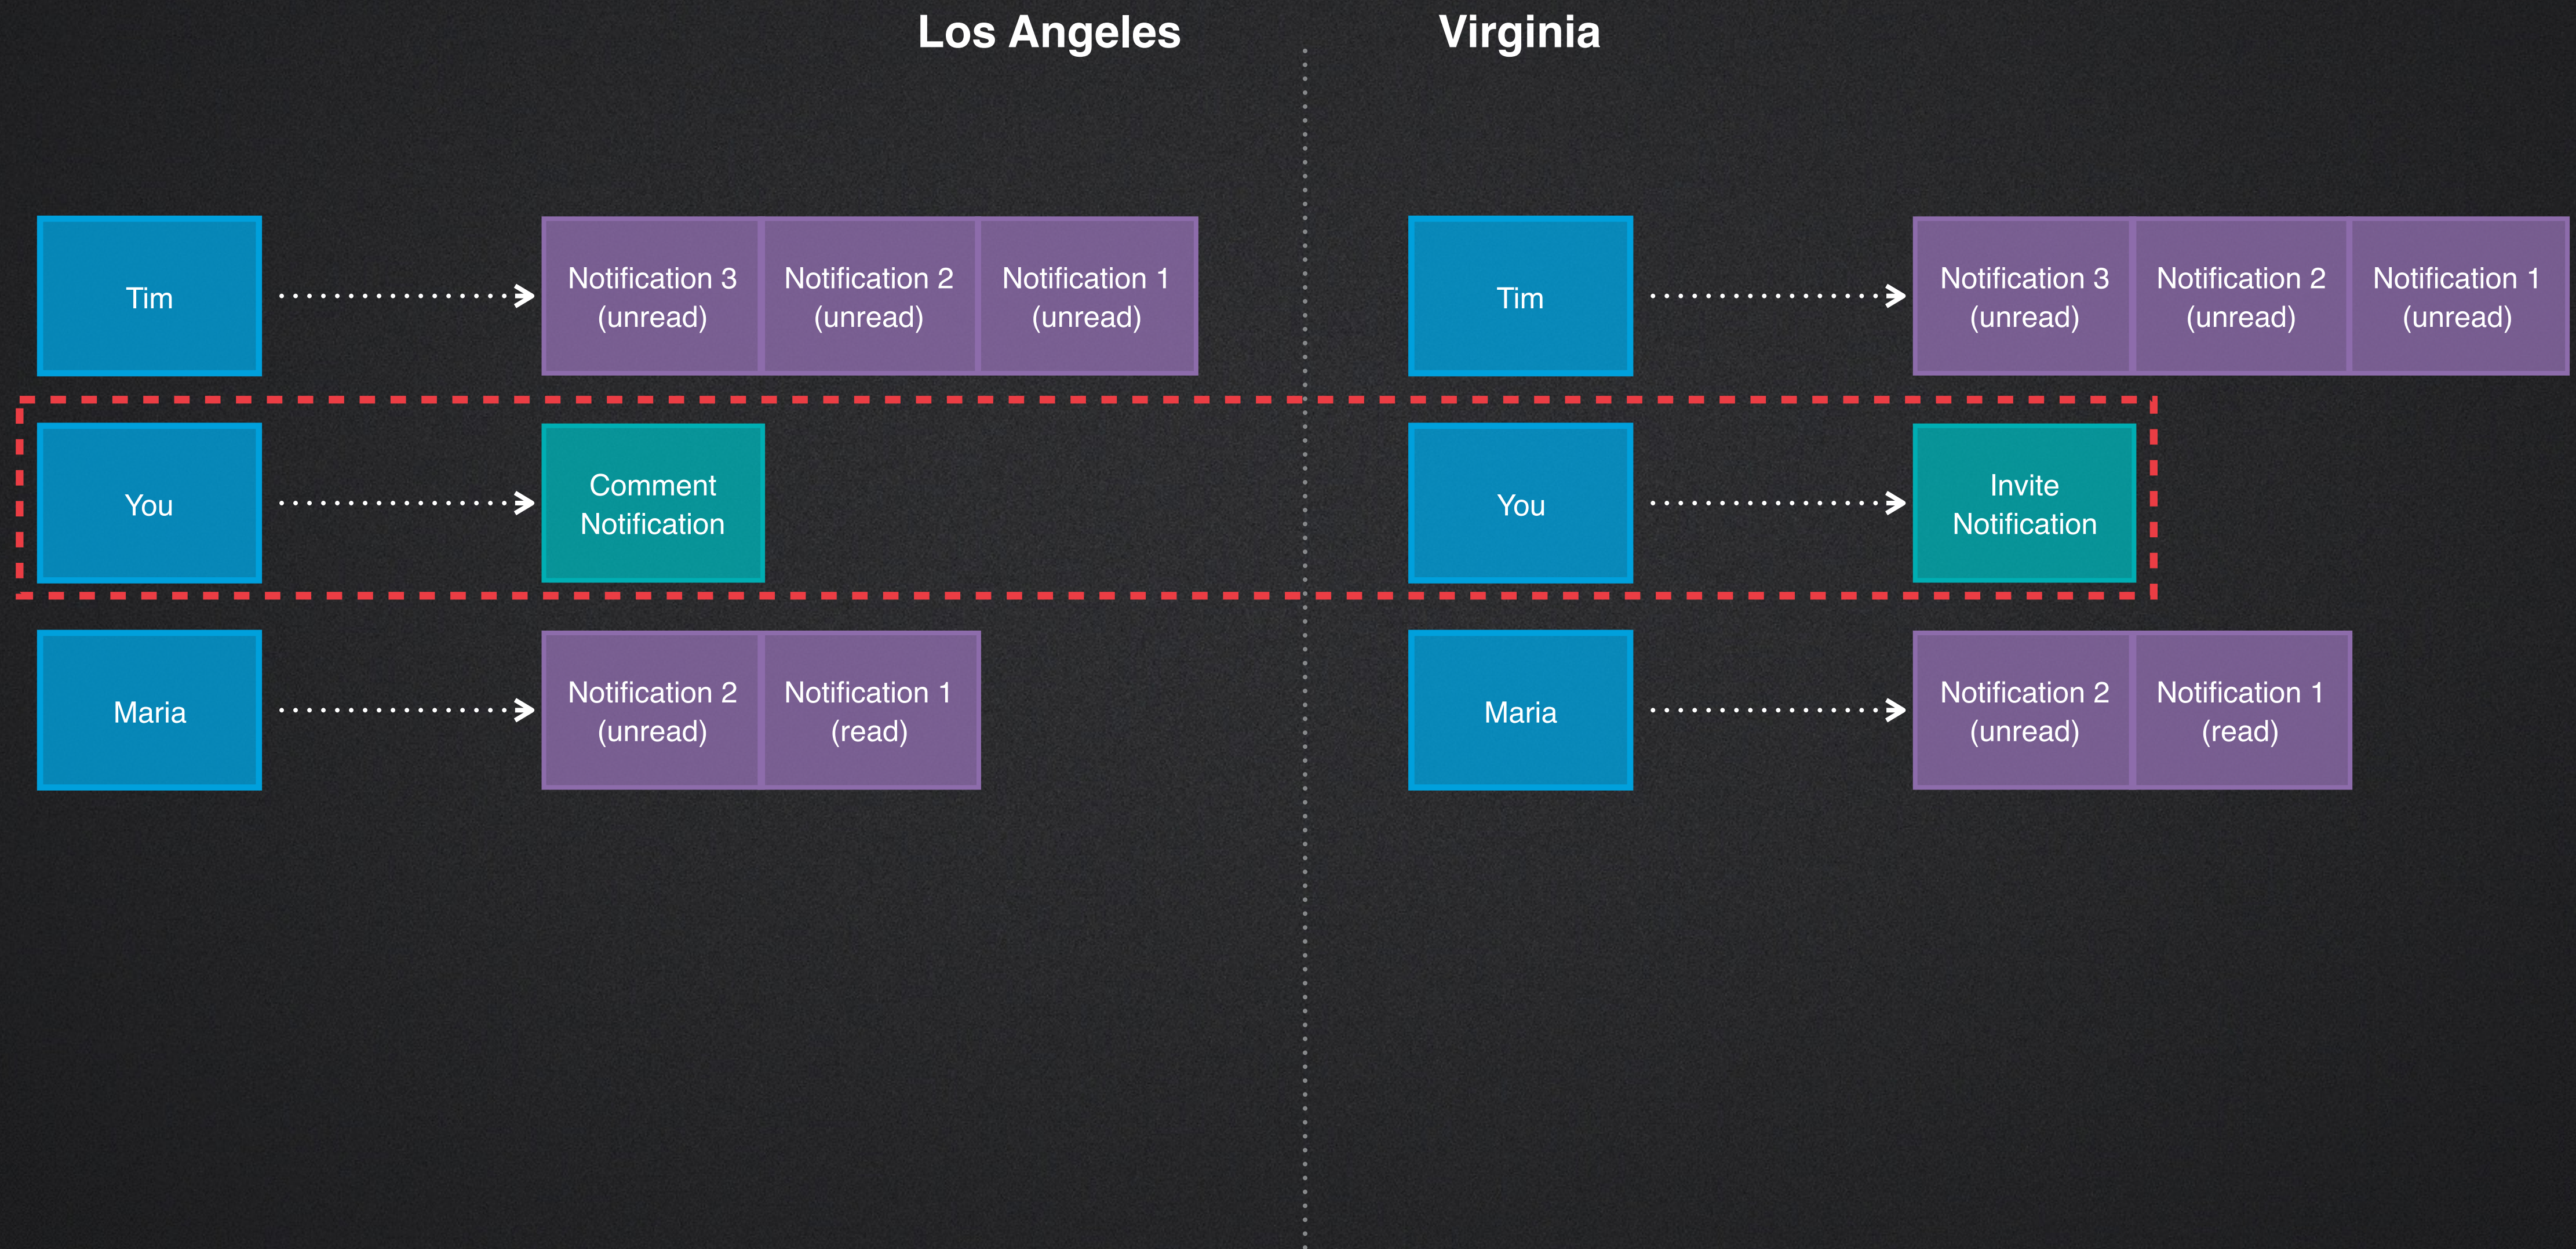


# Edge Case 2: Traffic shift race





# Edge Case 2: Traffic shift race





Only Process Once





2

Personal Data





Your data

is *your* data.



## Complication #1

---

Some of our personal  
databases are very large



95

---

TeraBytes



## Complication #2

---

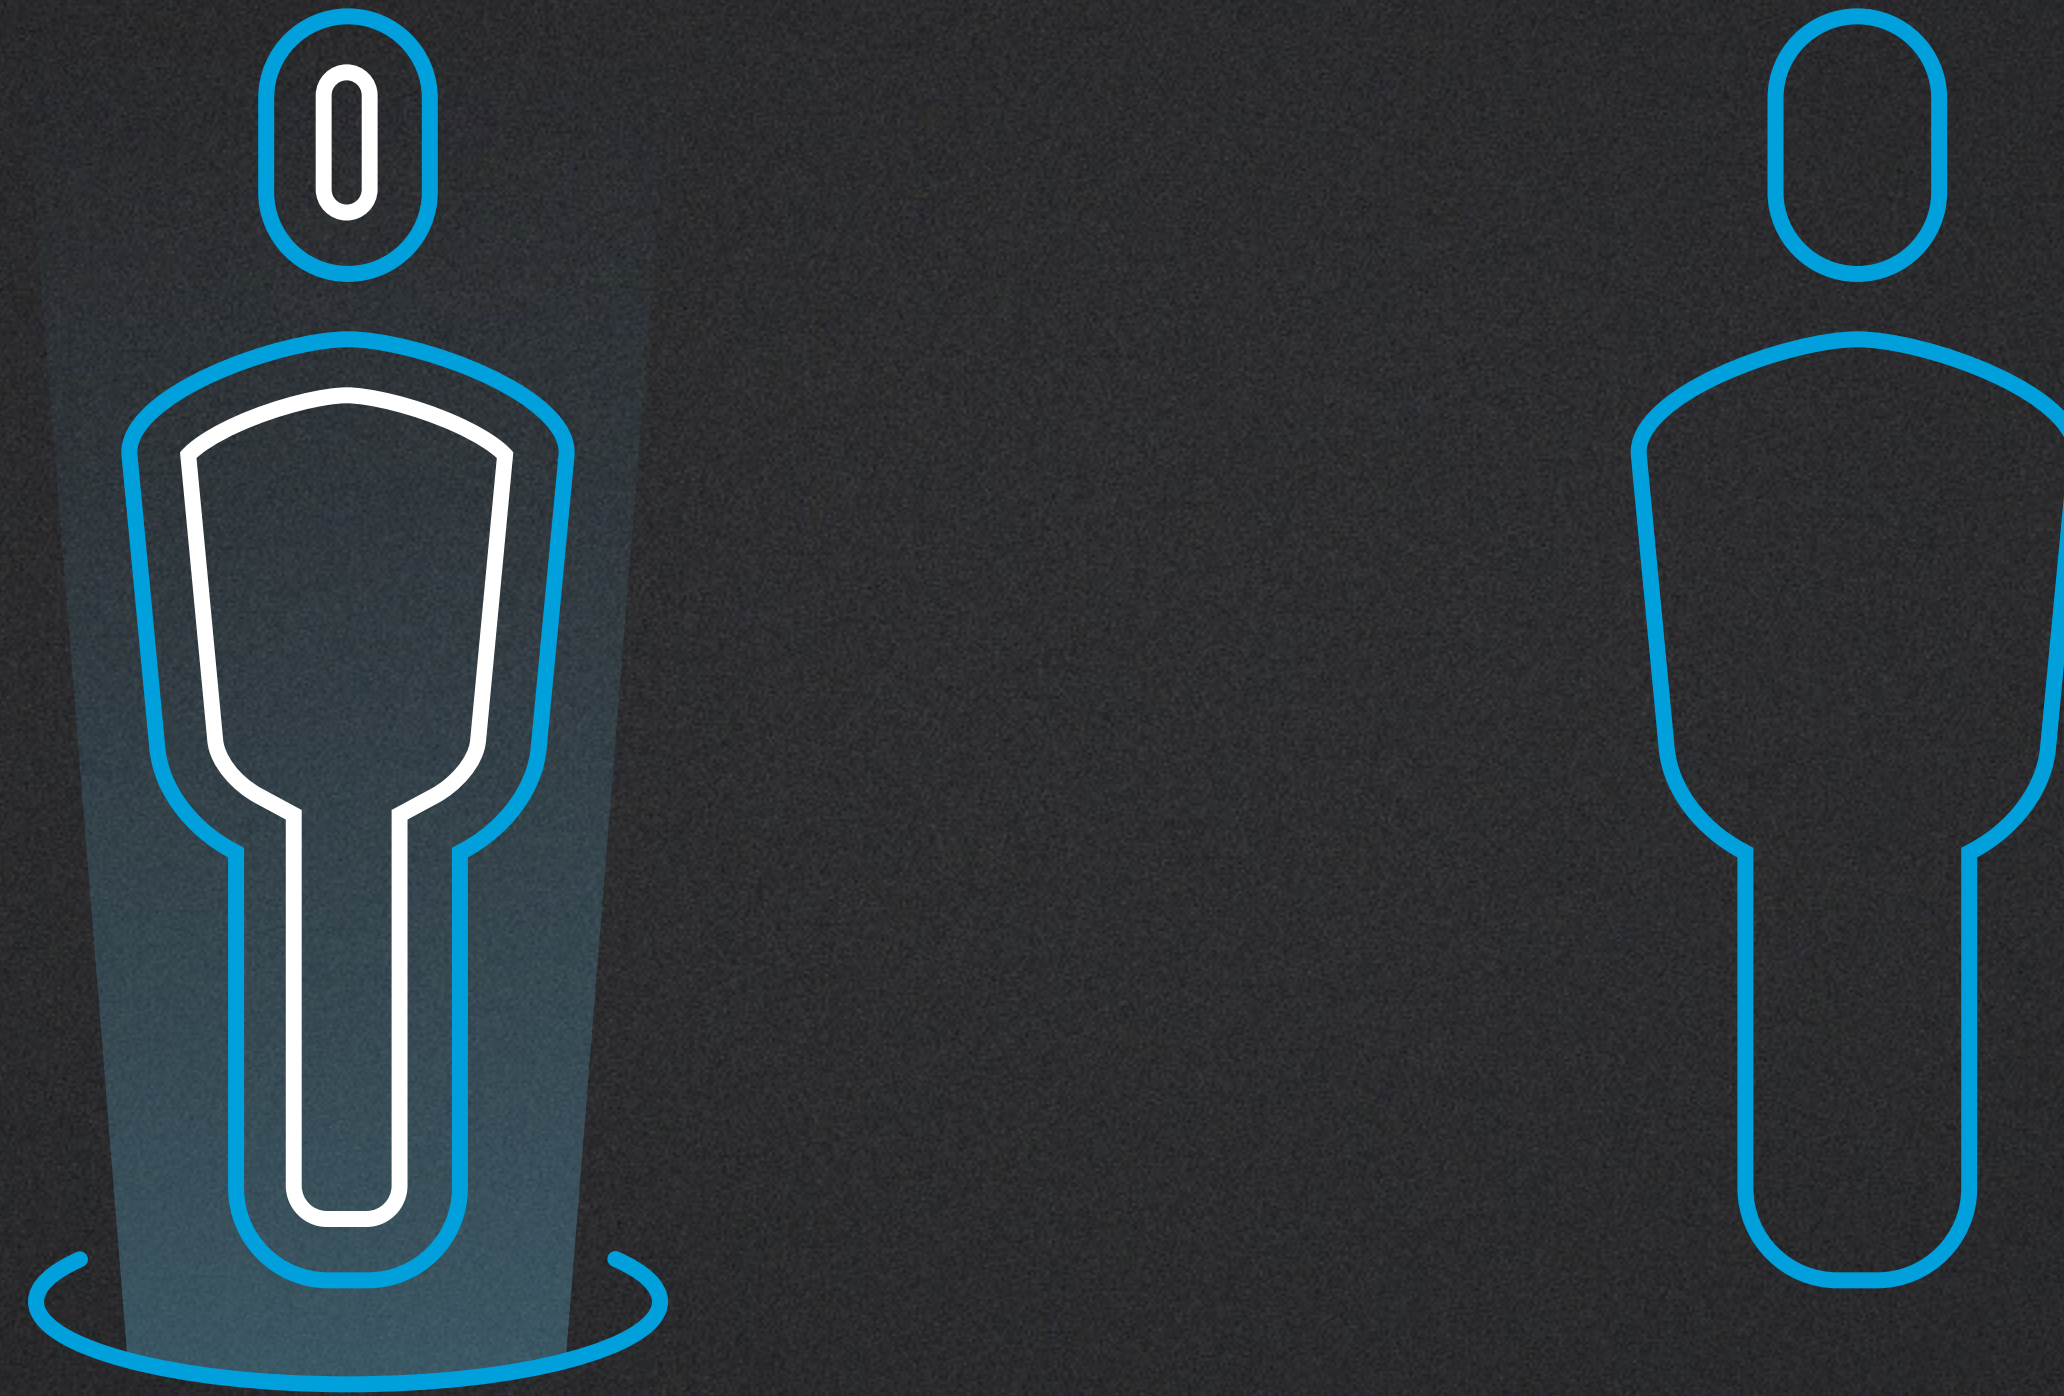
It's illegal to replicate  
personal data outside of  
some countries





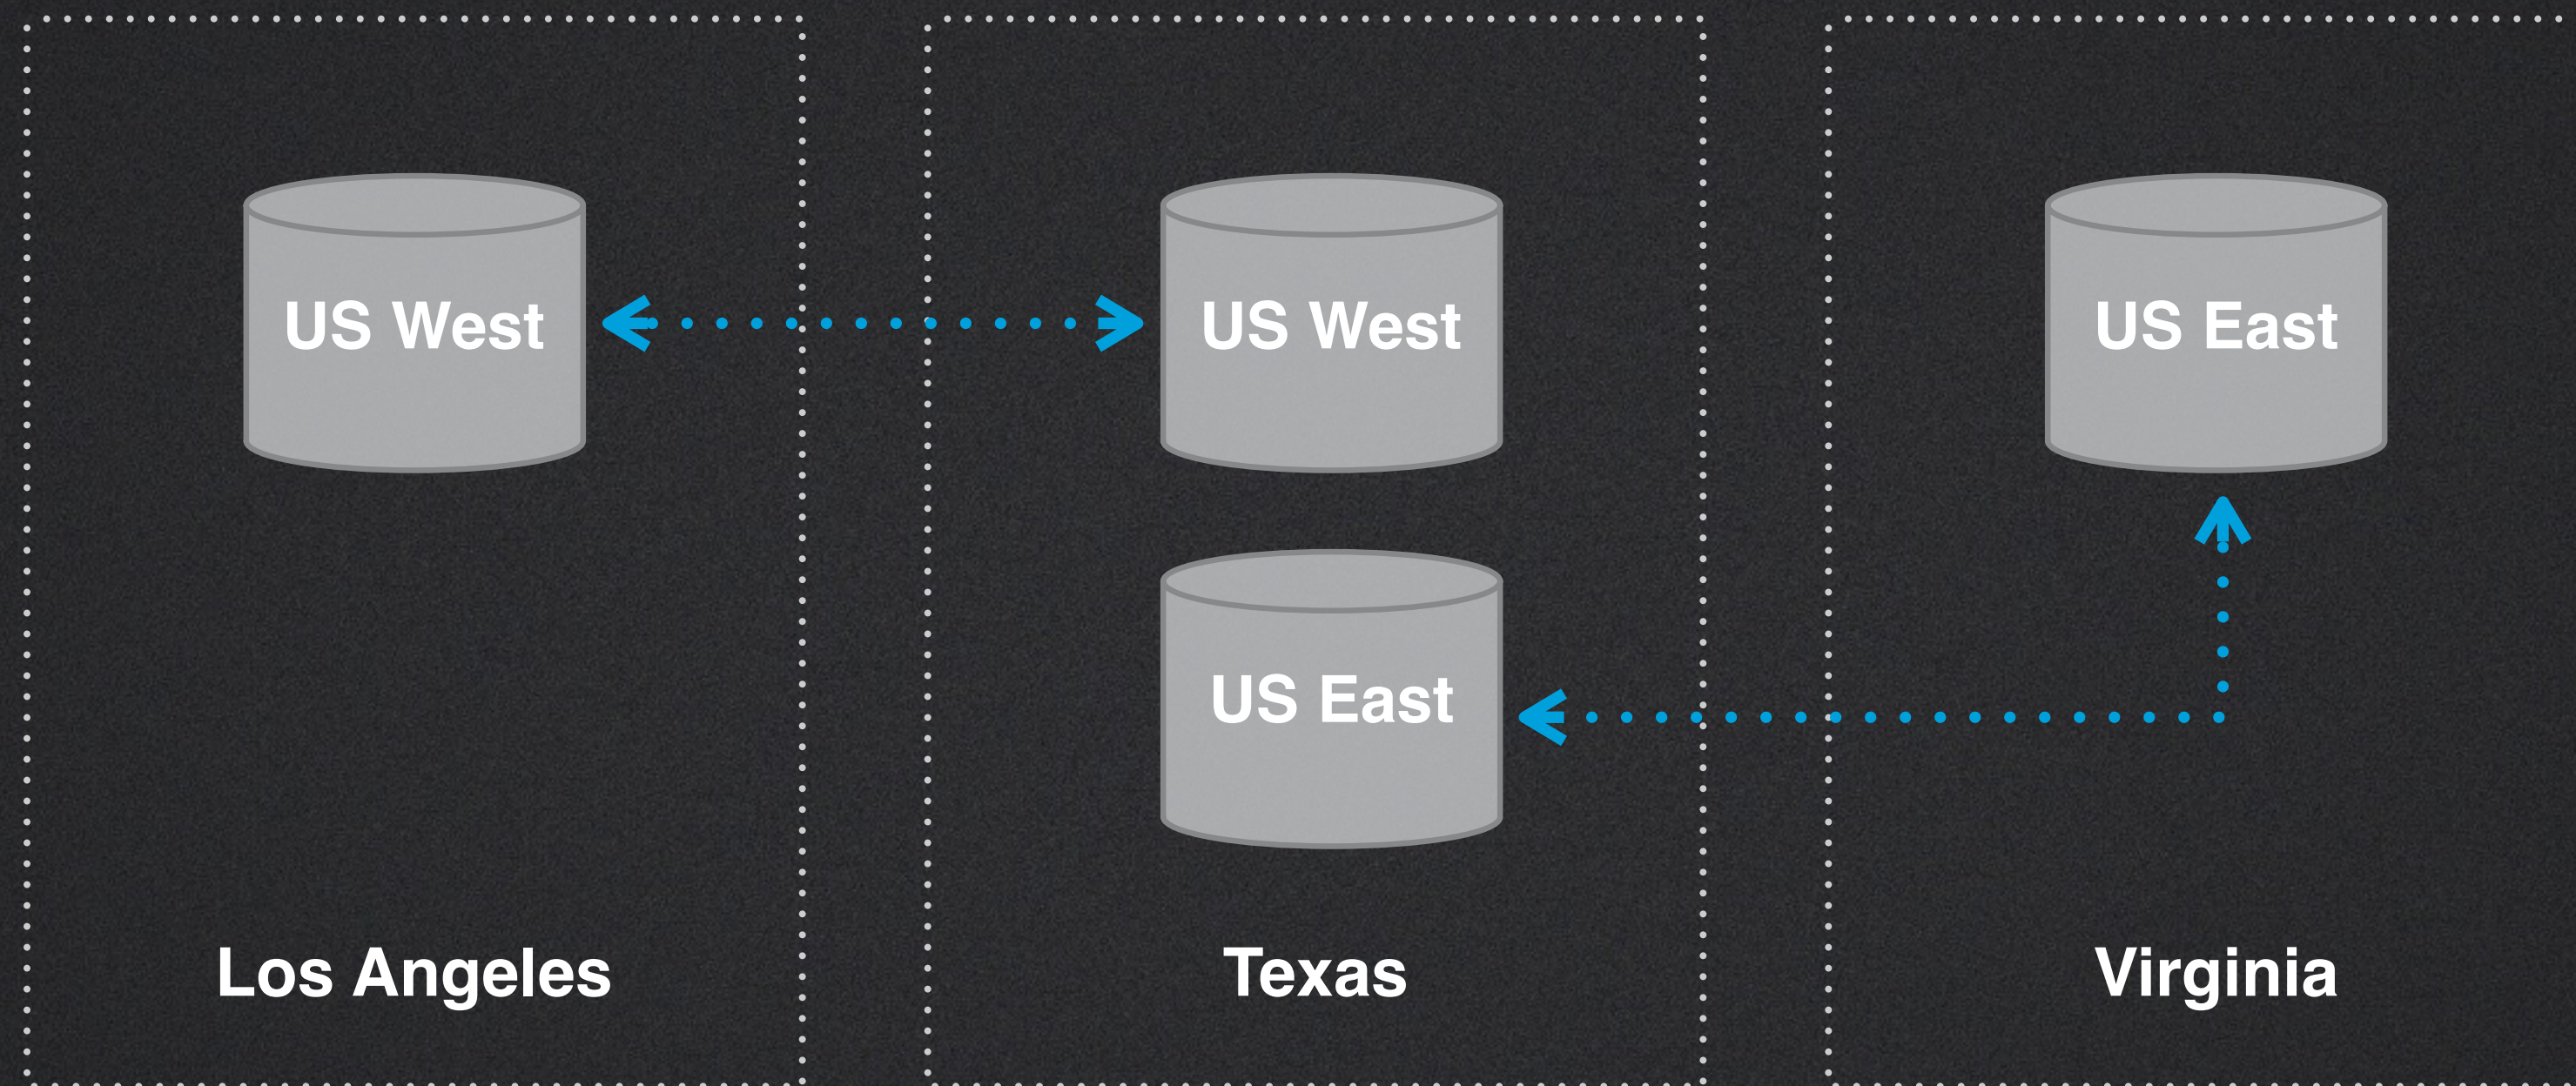


# You only need two copies of Personal Data



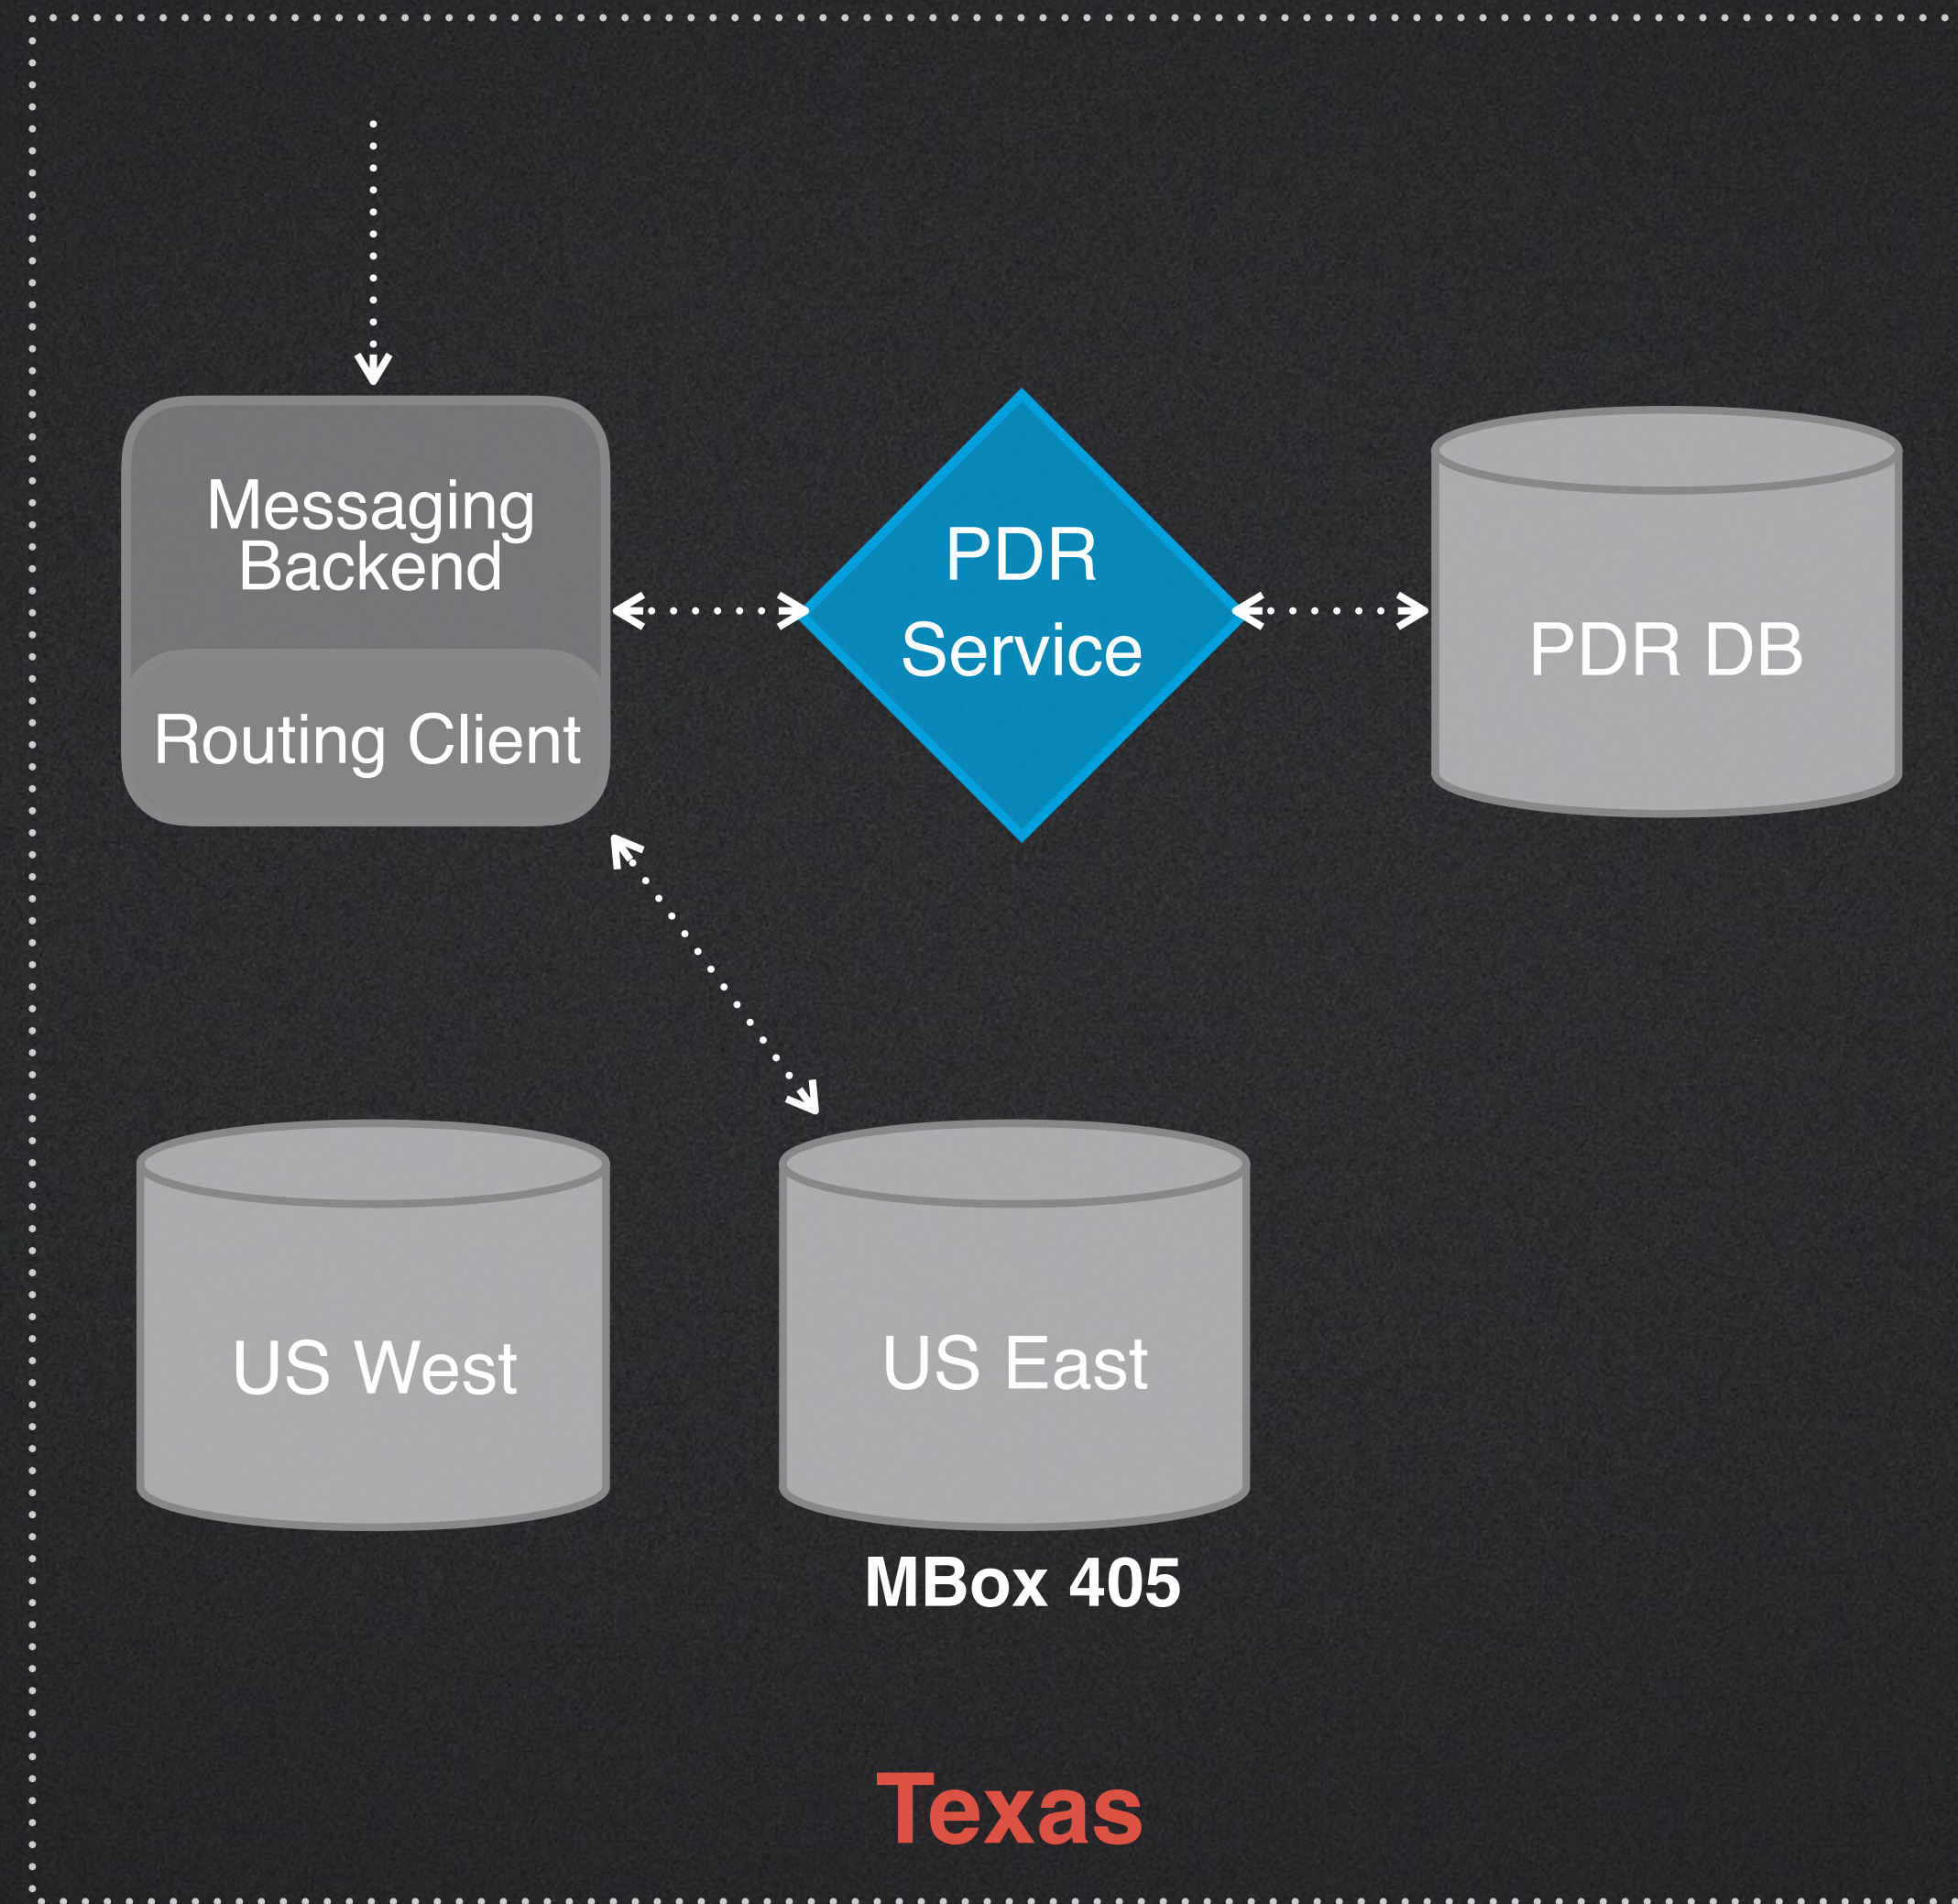


# Solution: Two Clusters



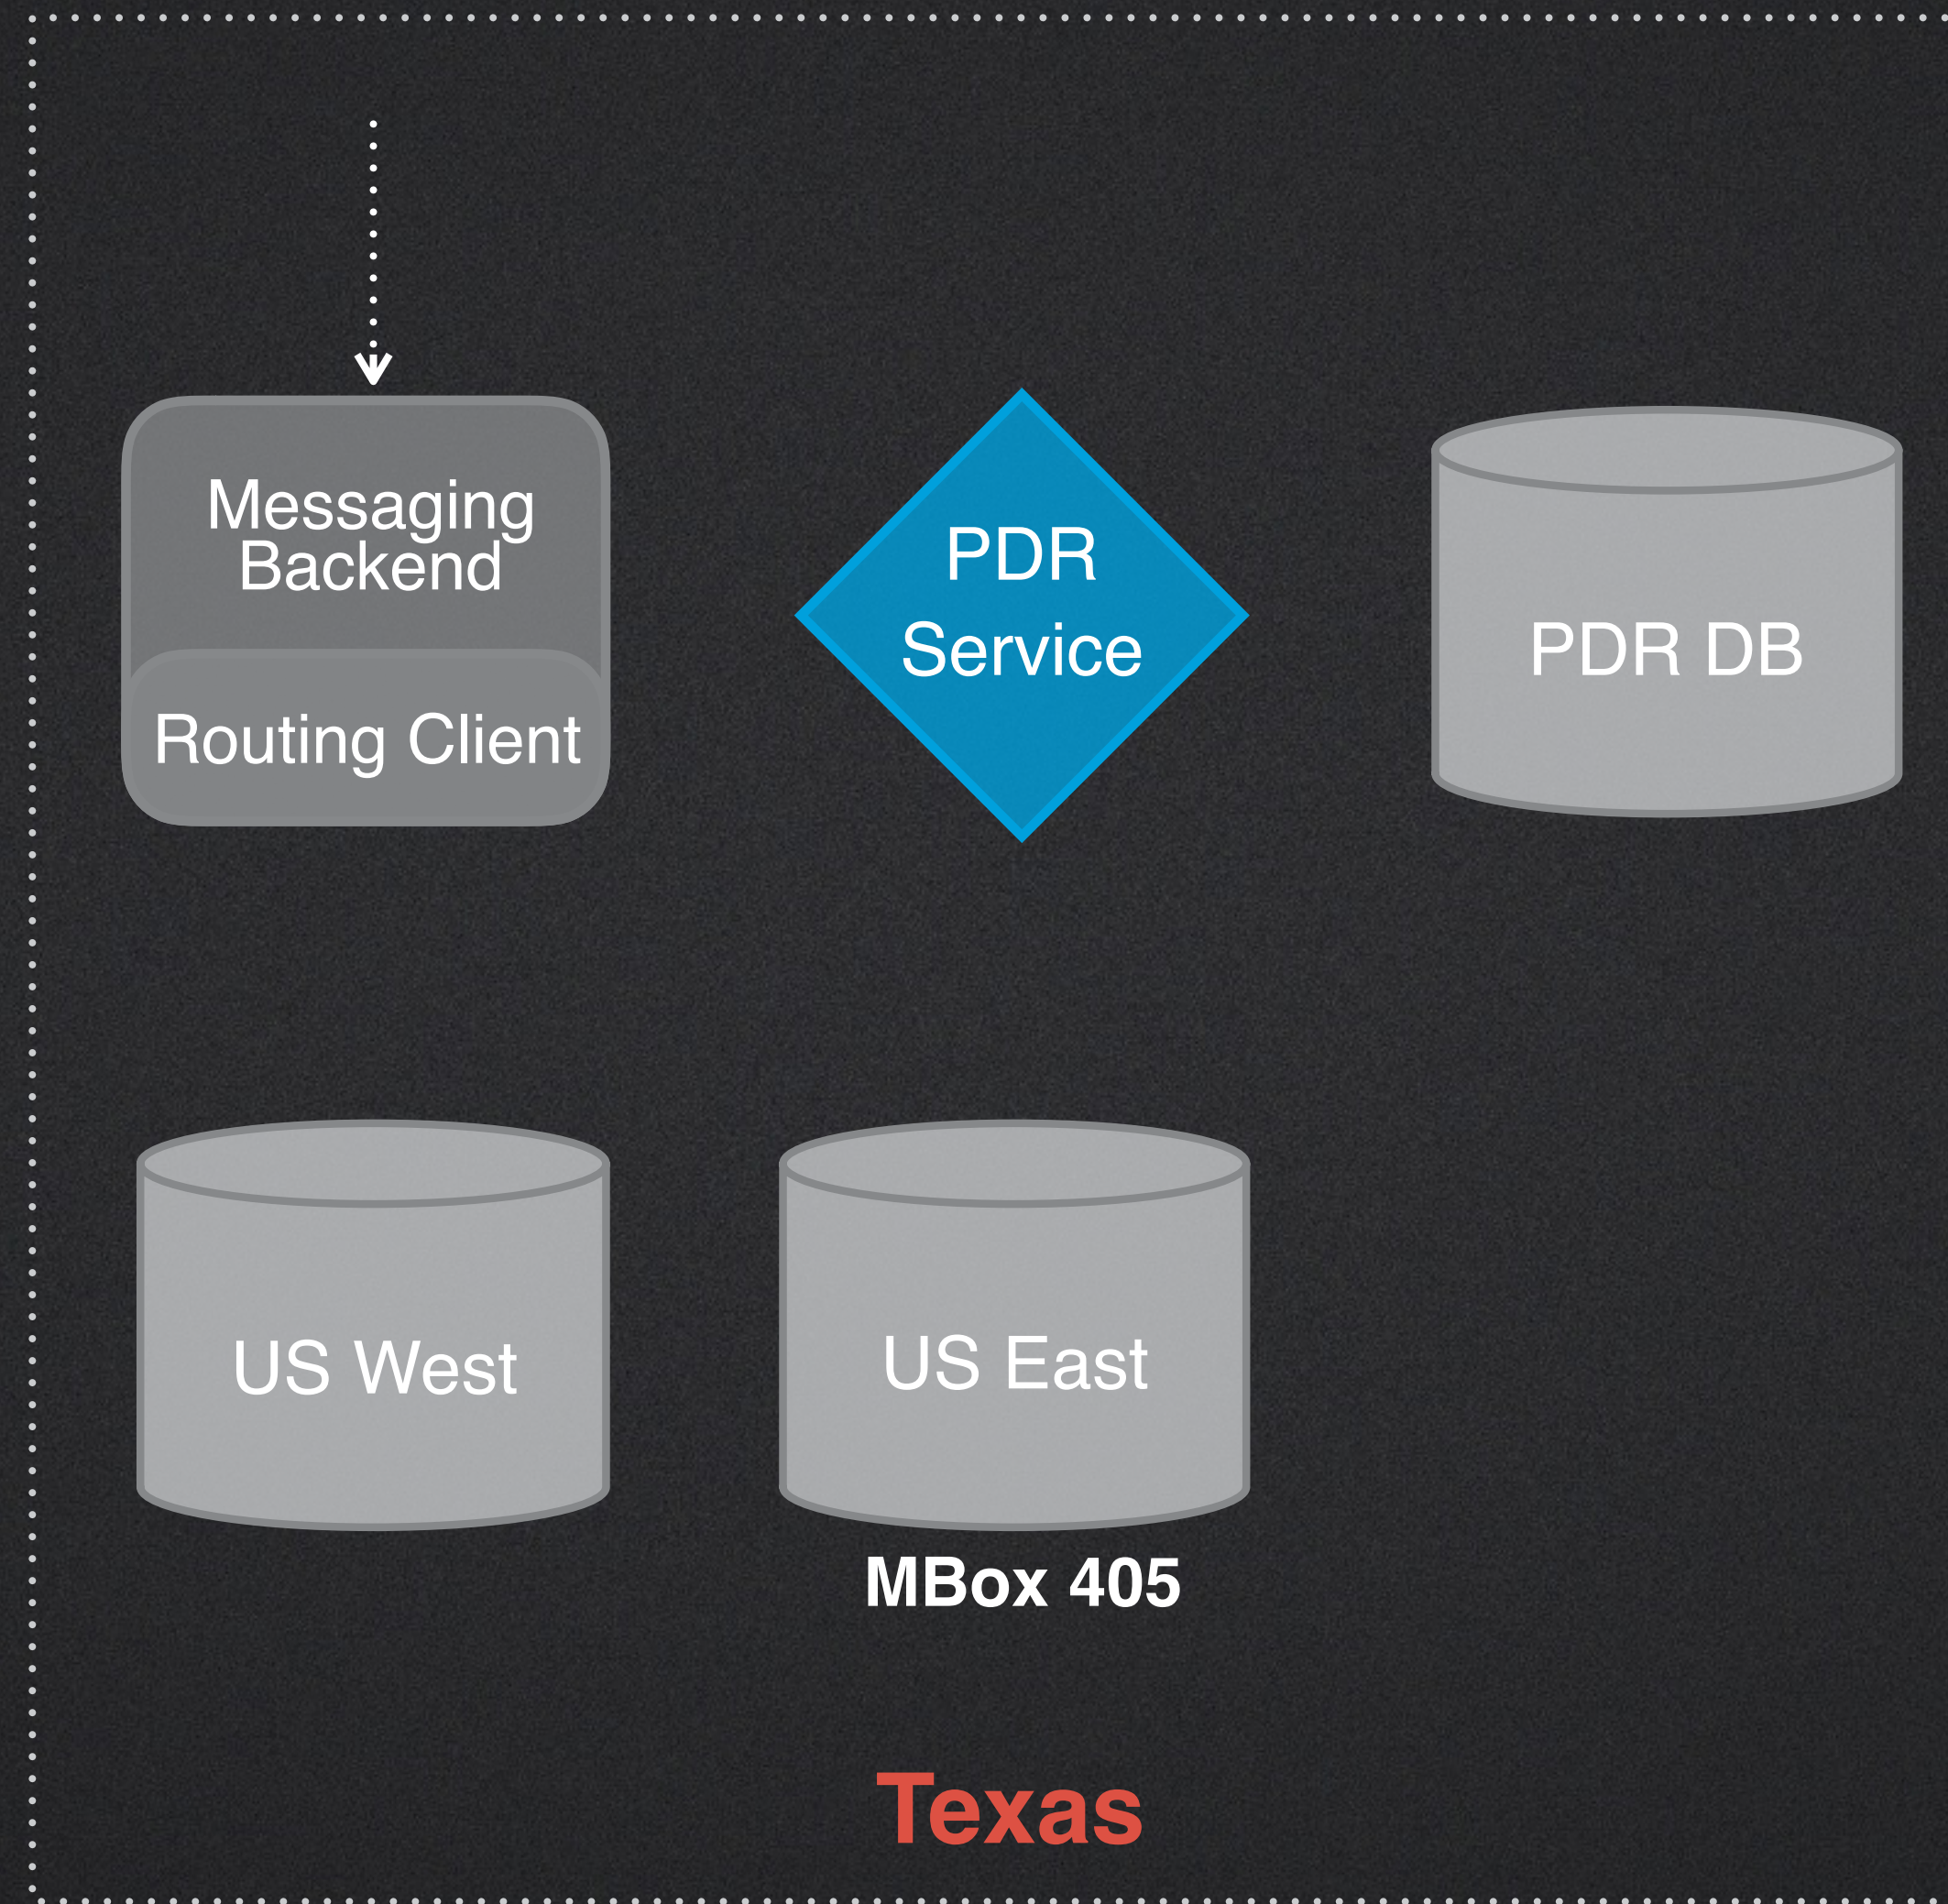


# Solution: Personal Data Routing





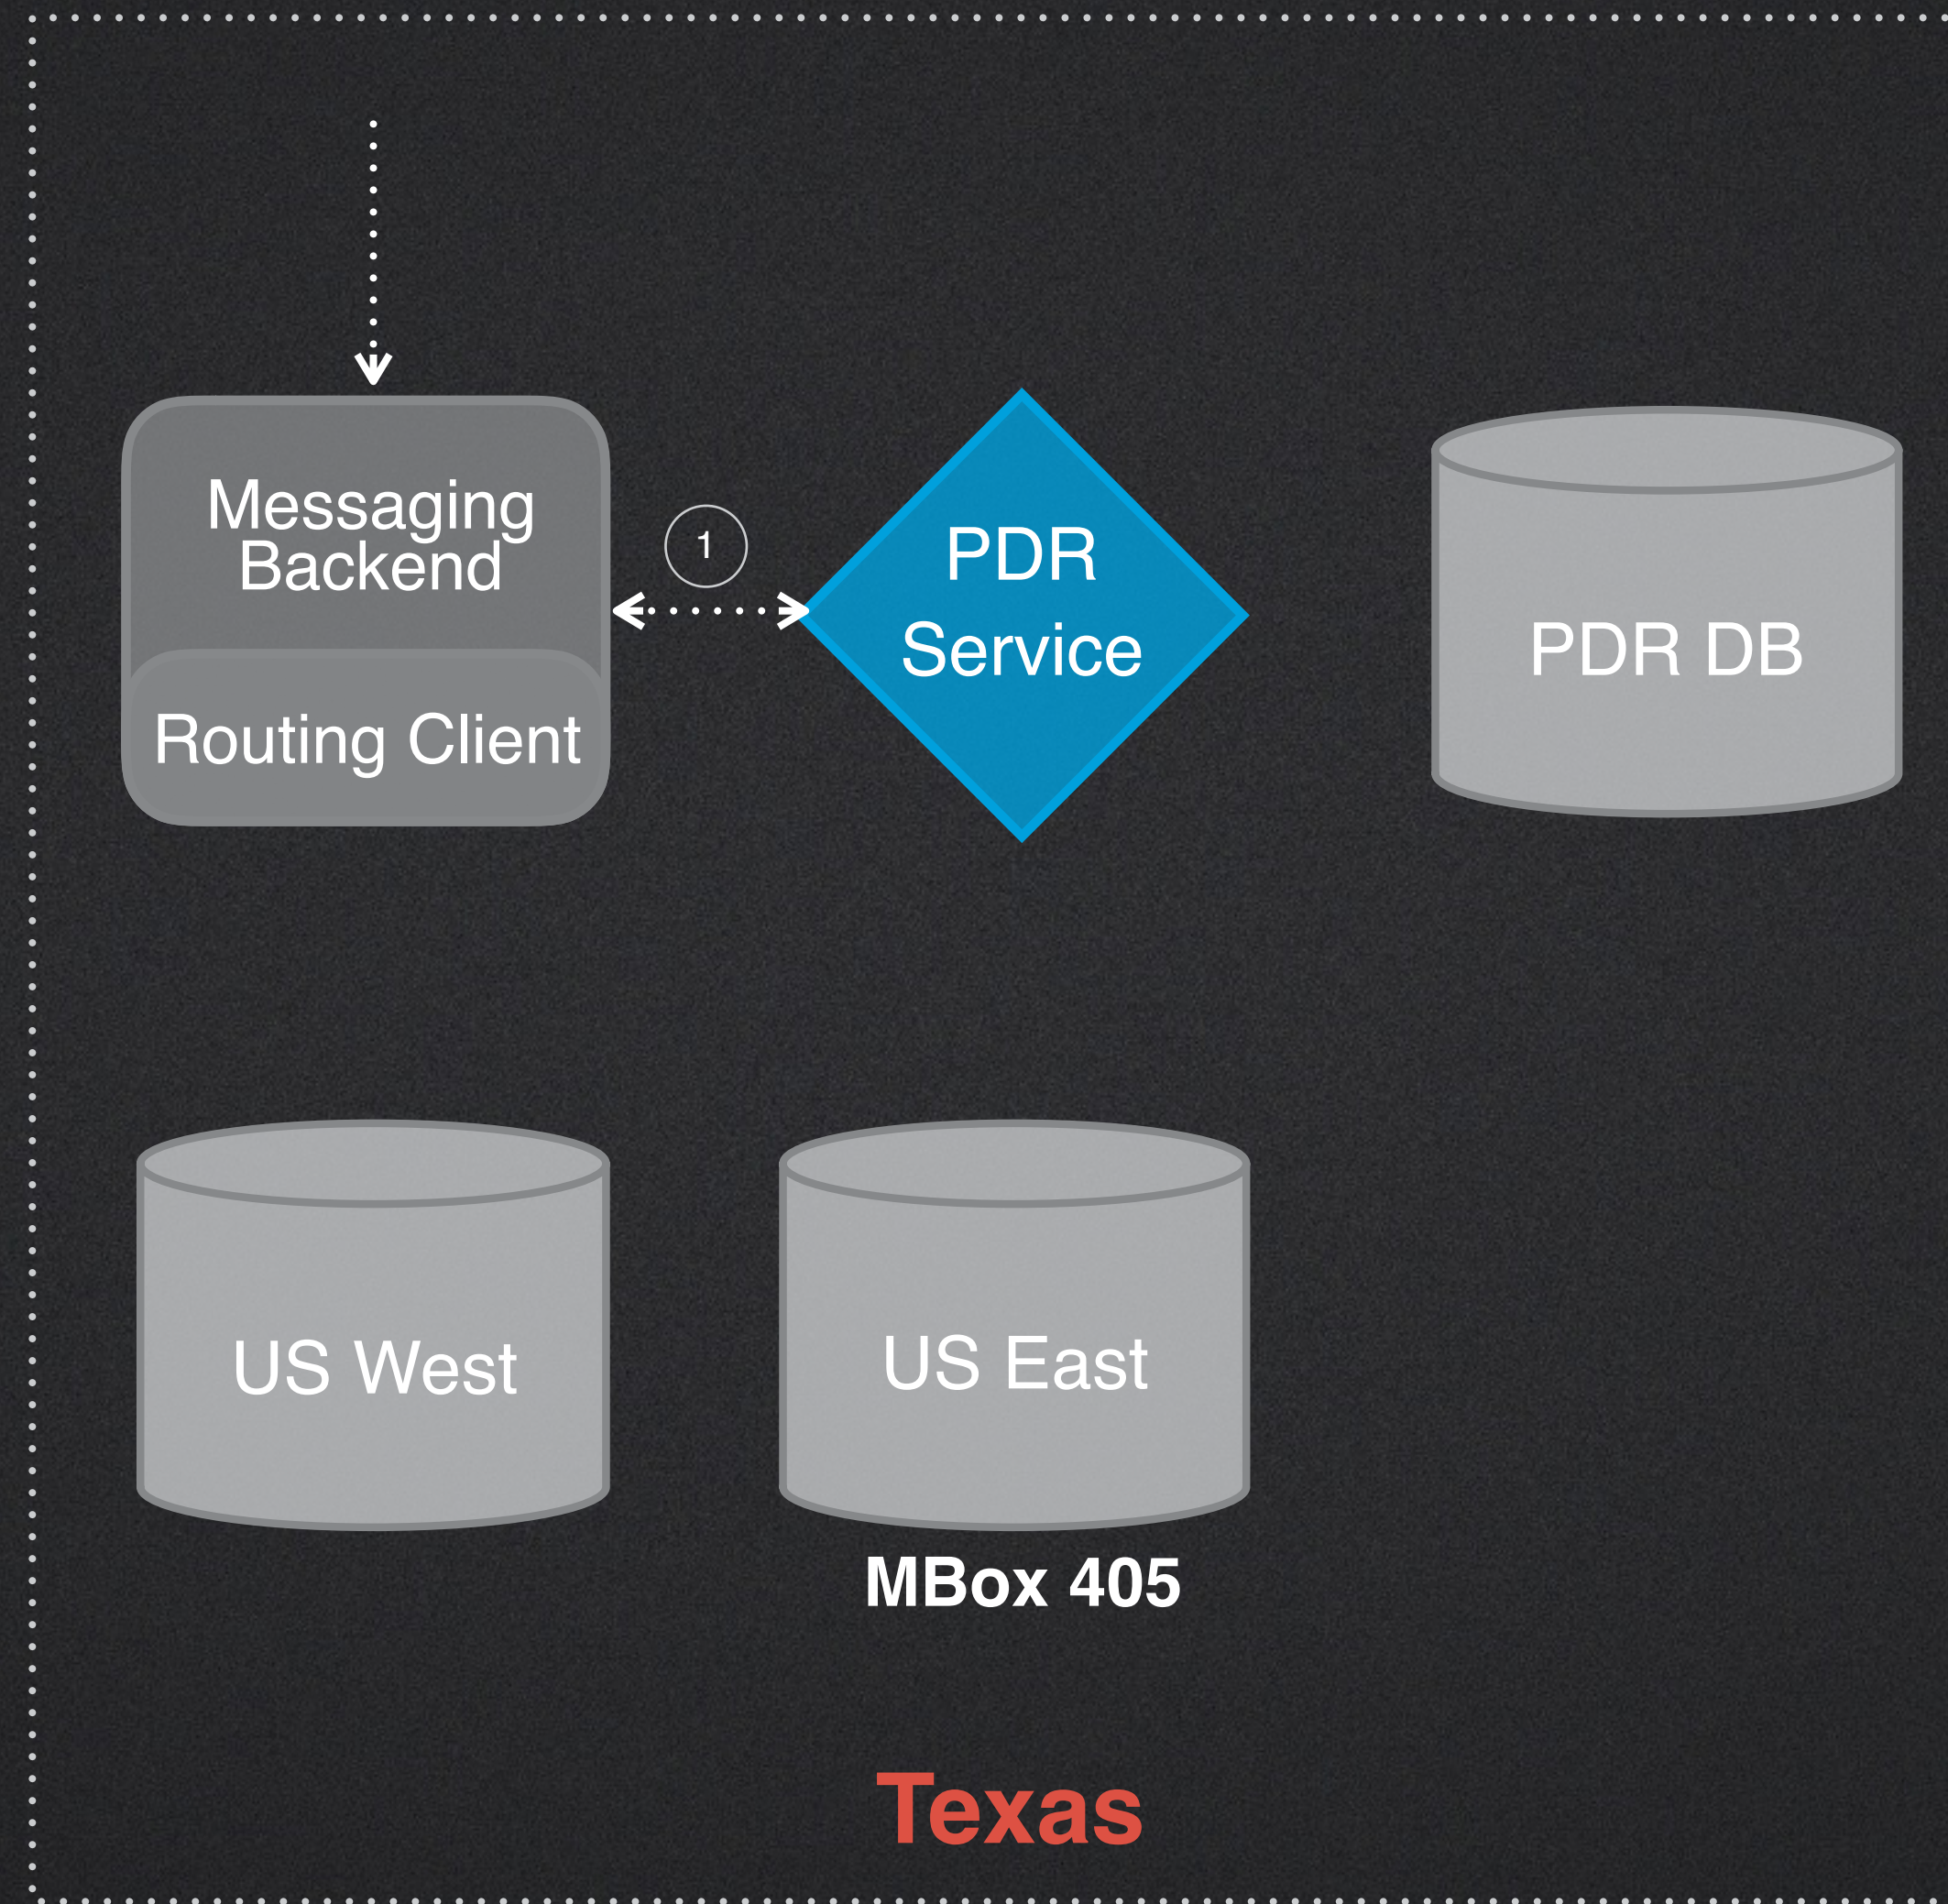
# How PDR works



/MESSAGING/MailboxDB/Mailbox/m-405



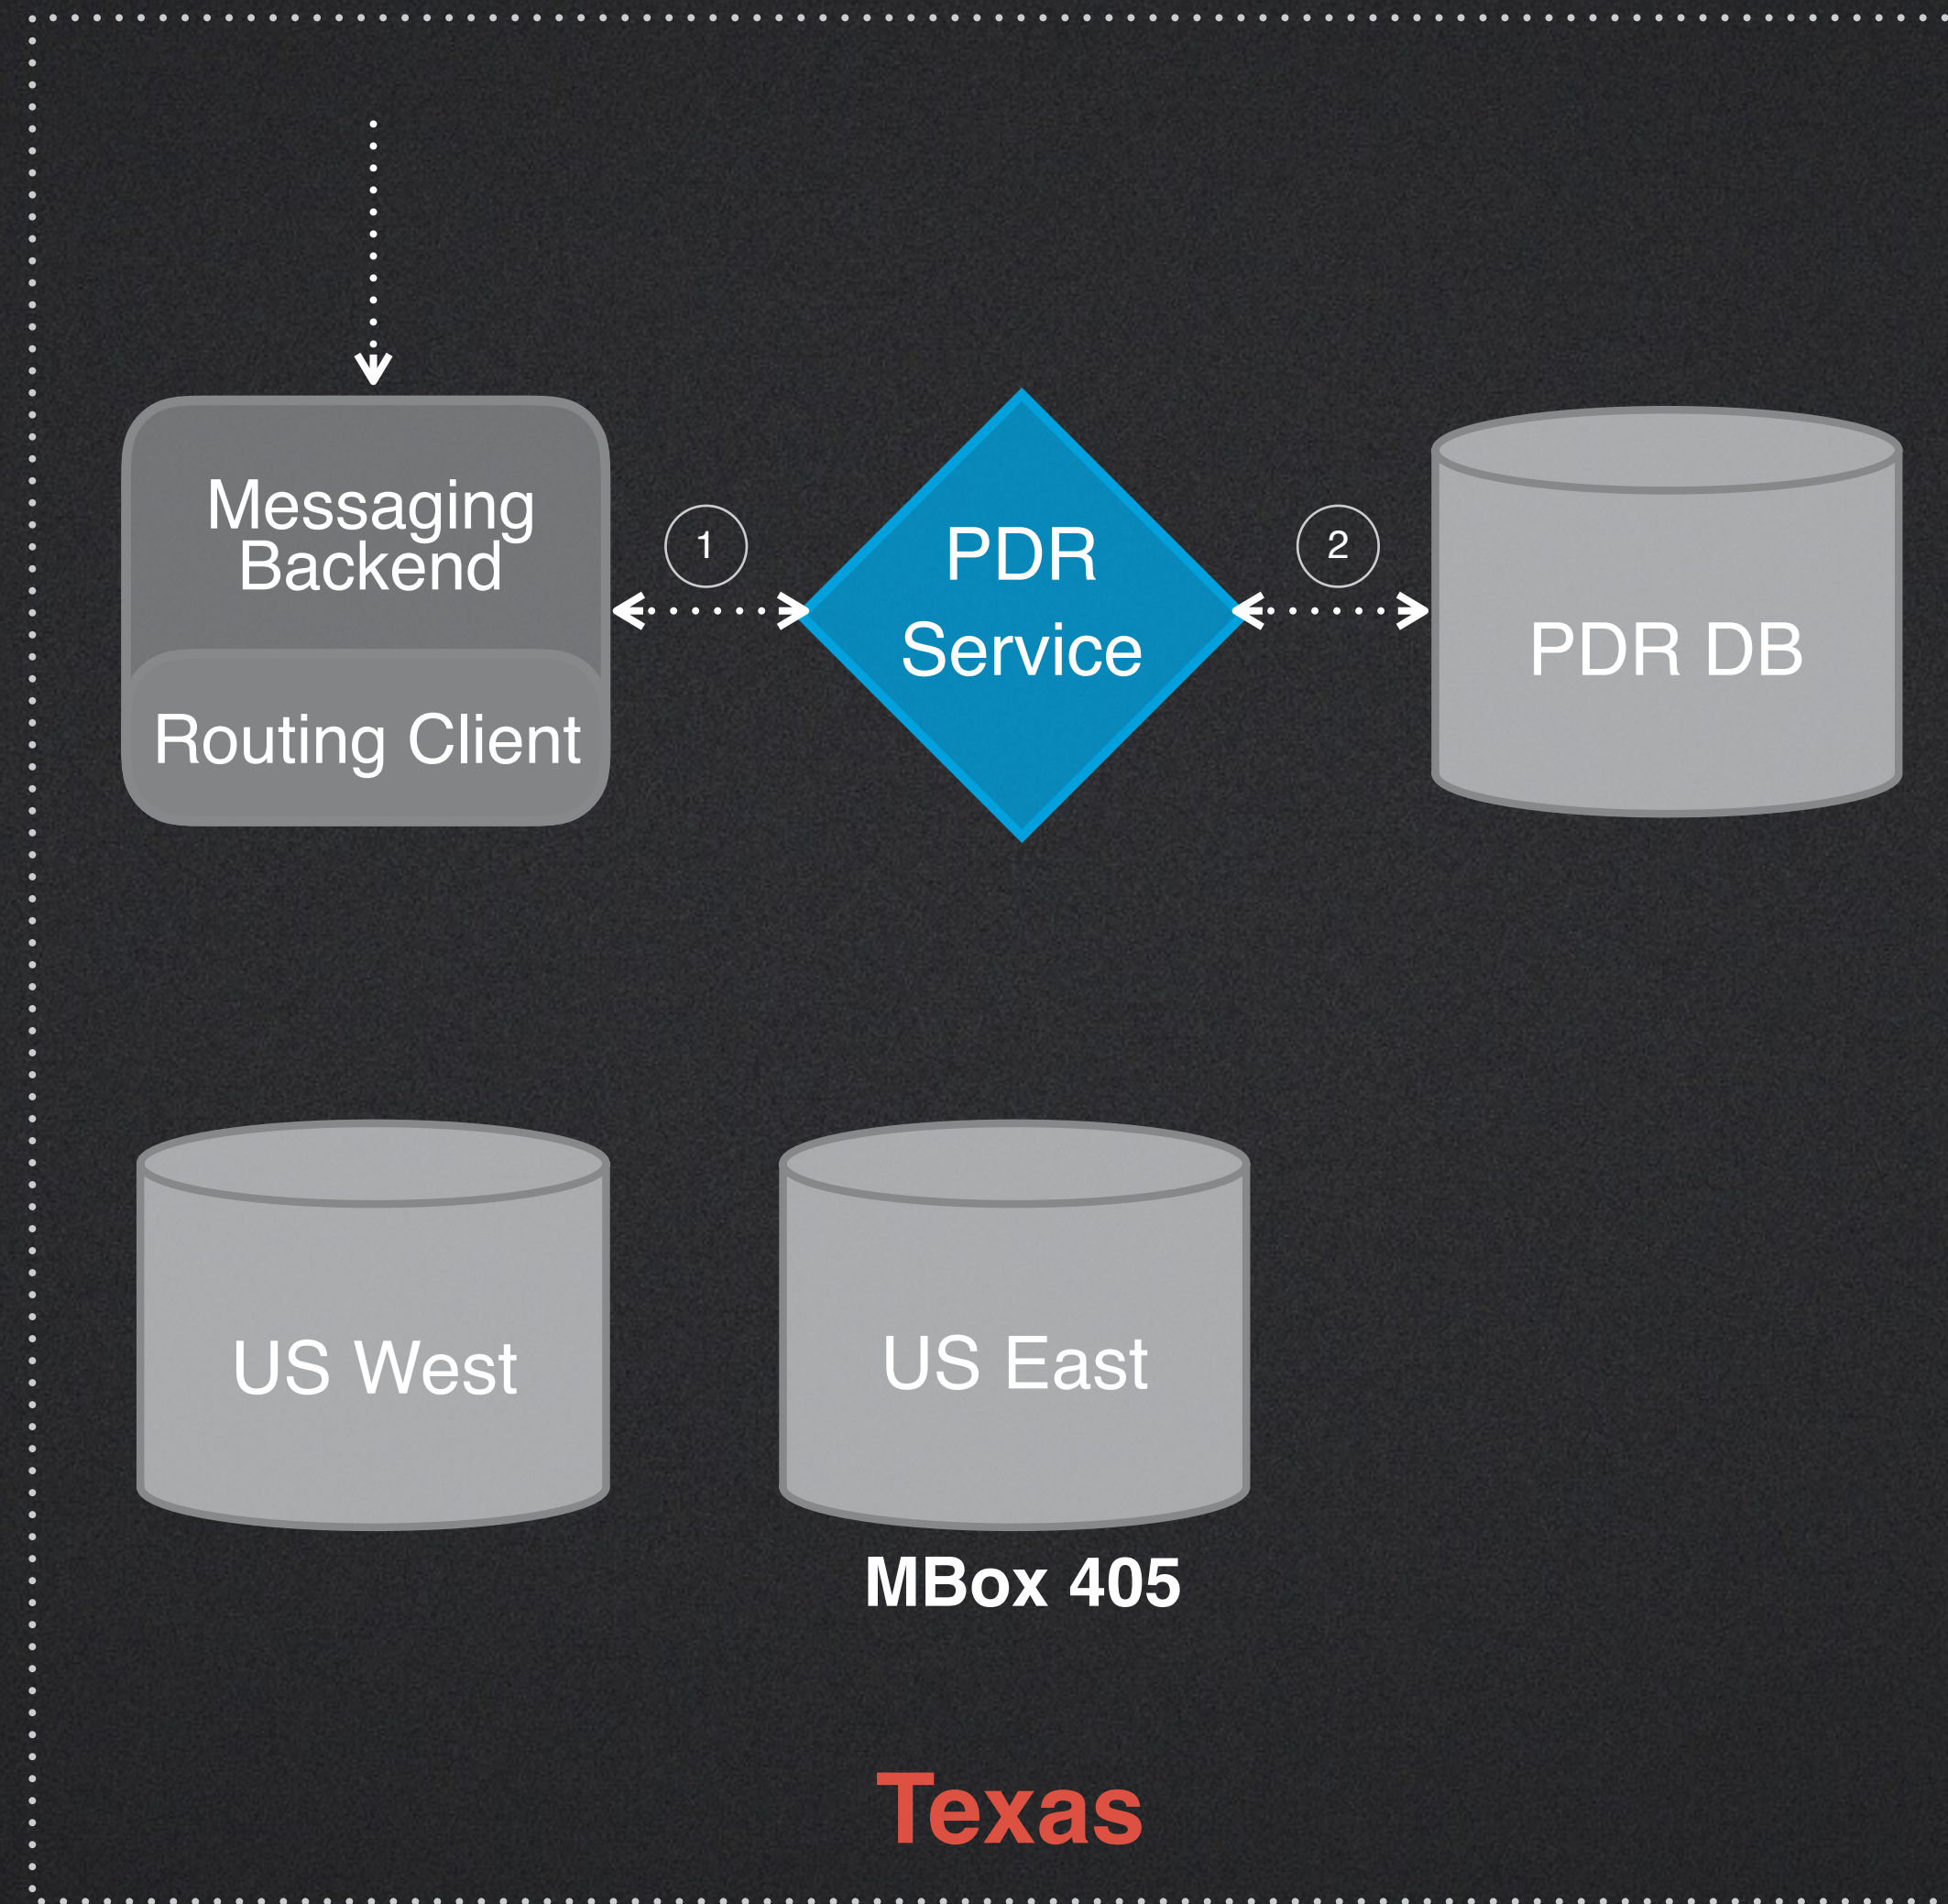
# How PDR works



/MESSAGING/MailboxDB/Mailbox/m-405



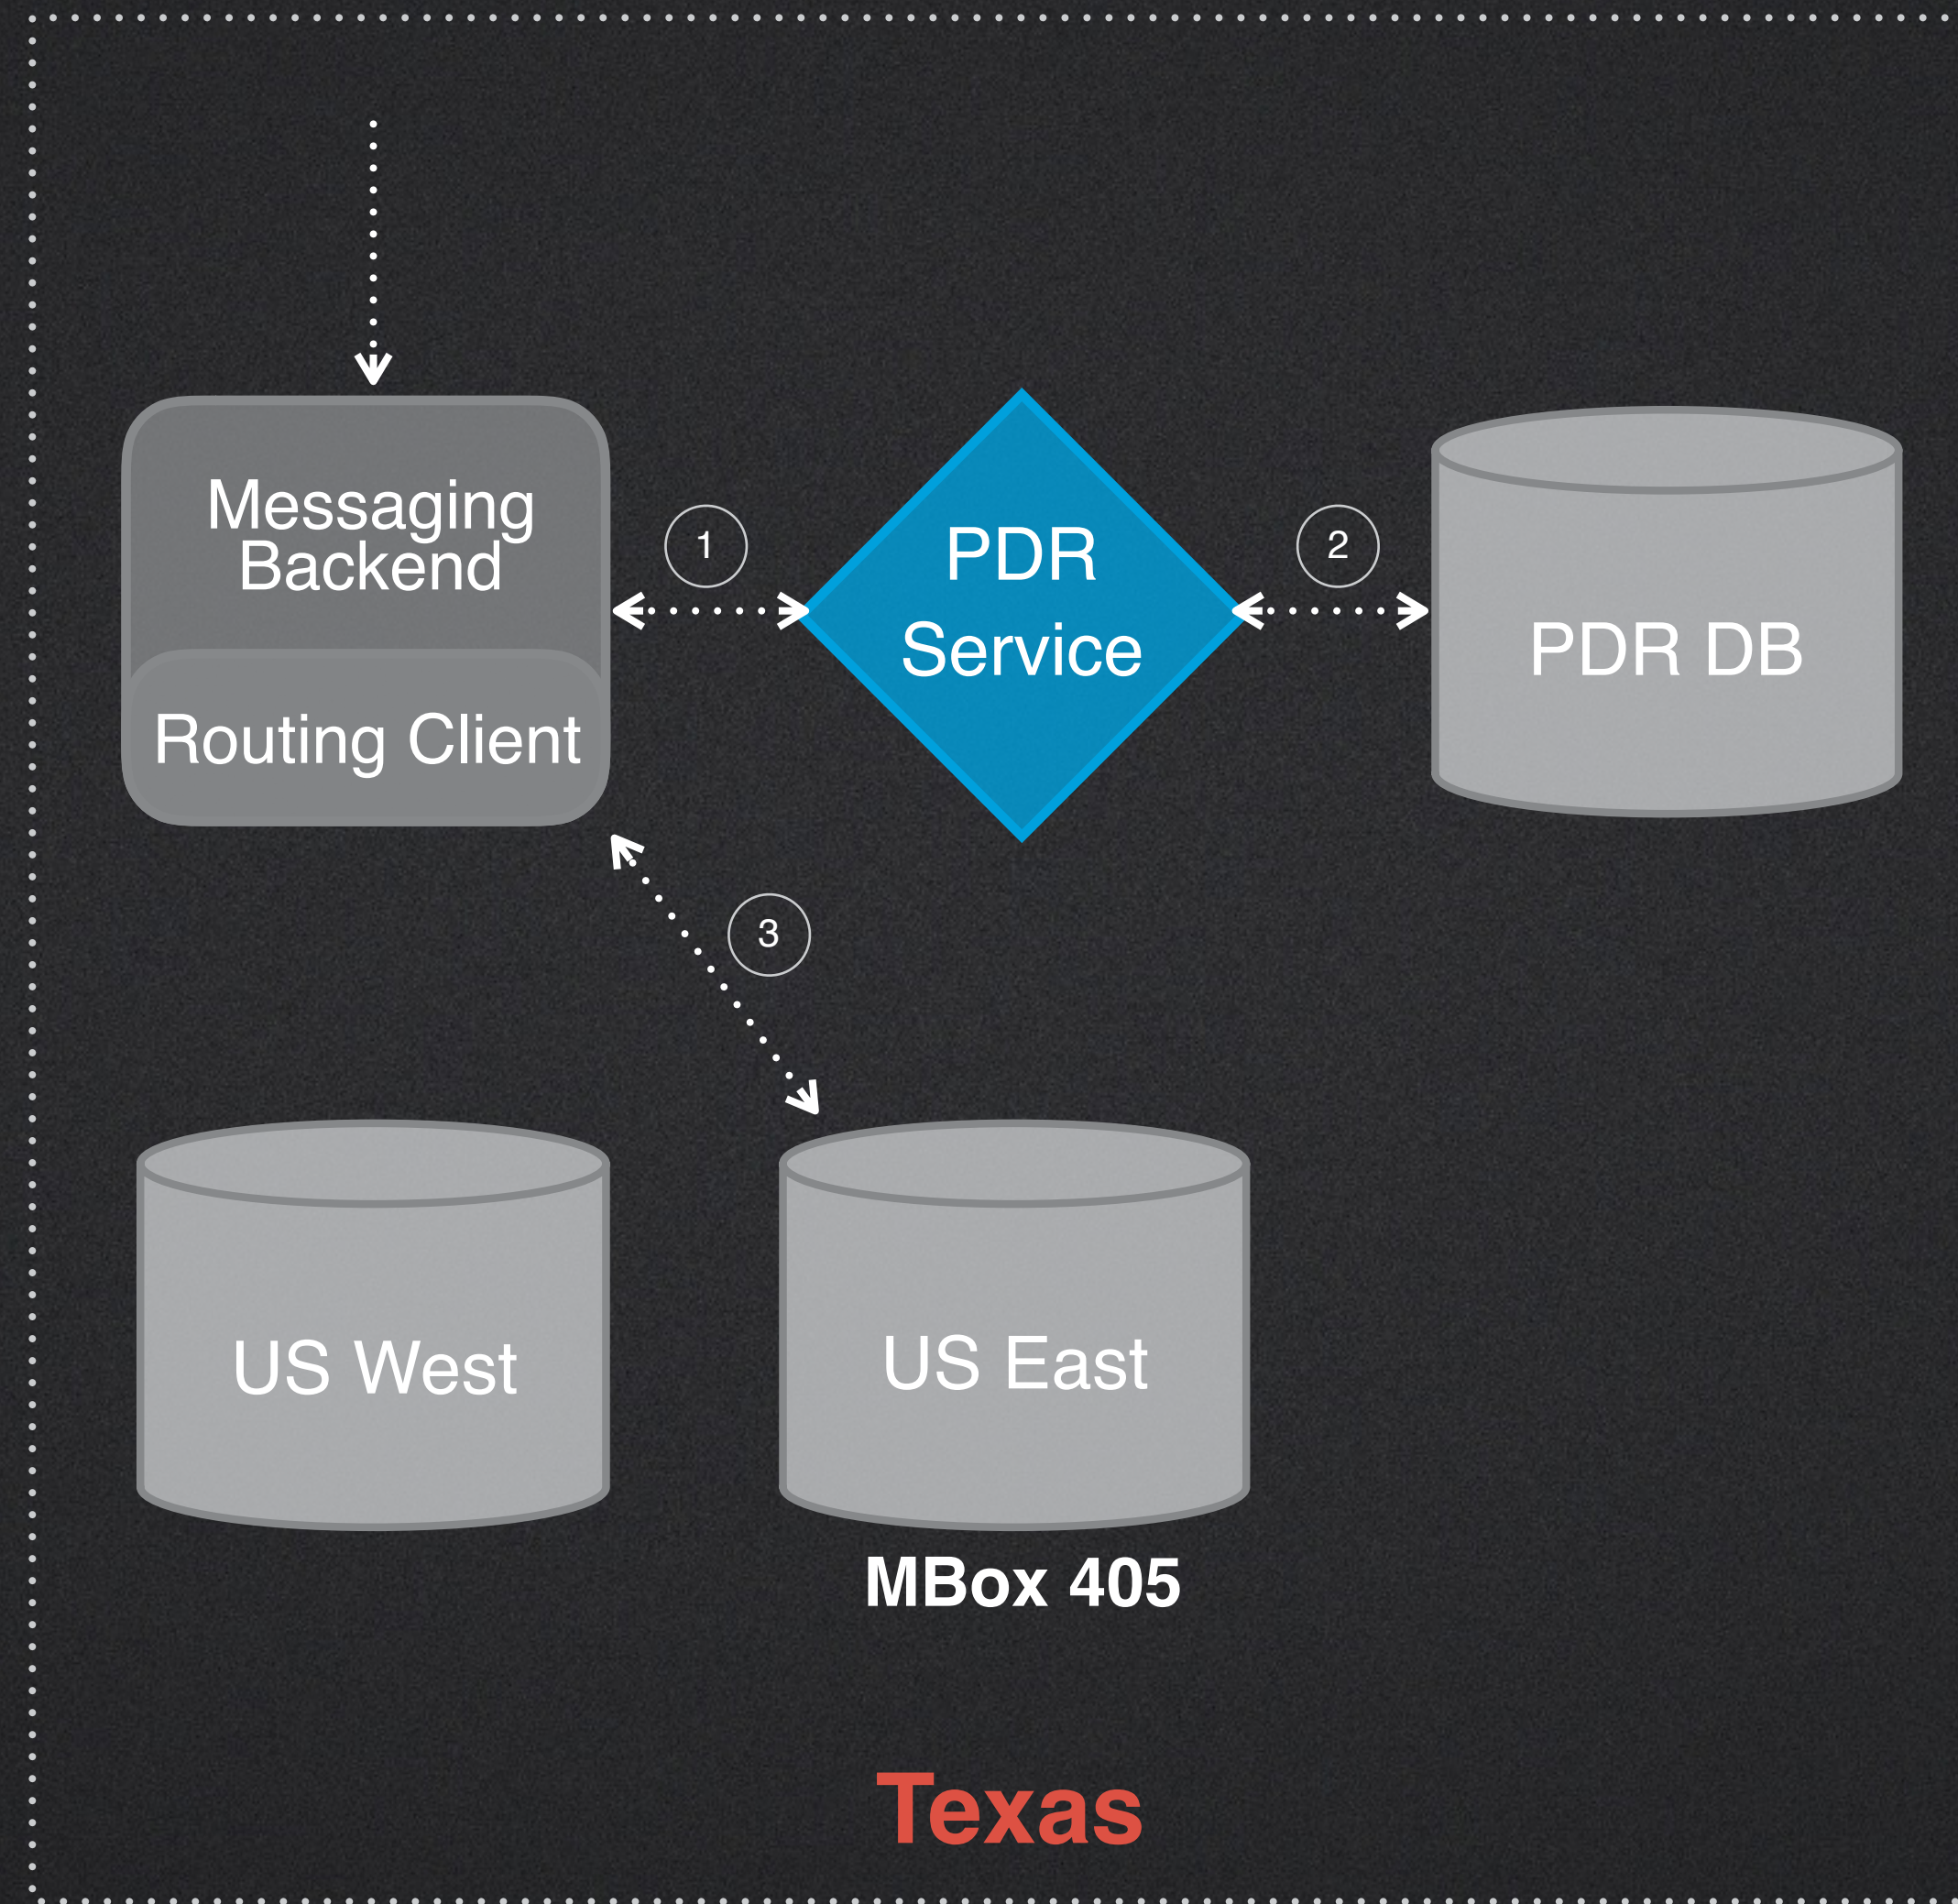
# How PDR works



/MESSAGING/MailboxDB/Mailbox/m-405



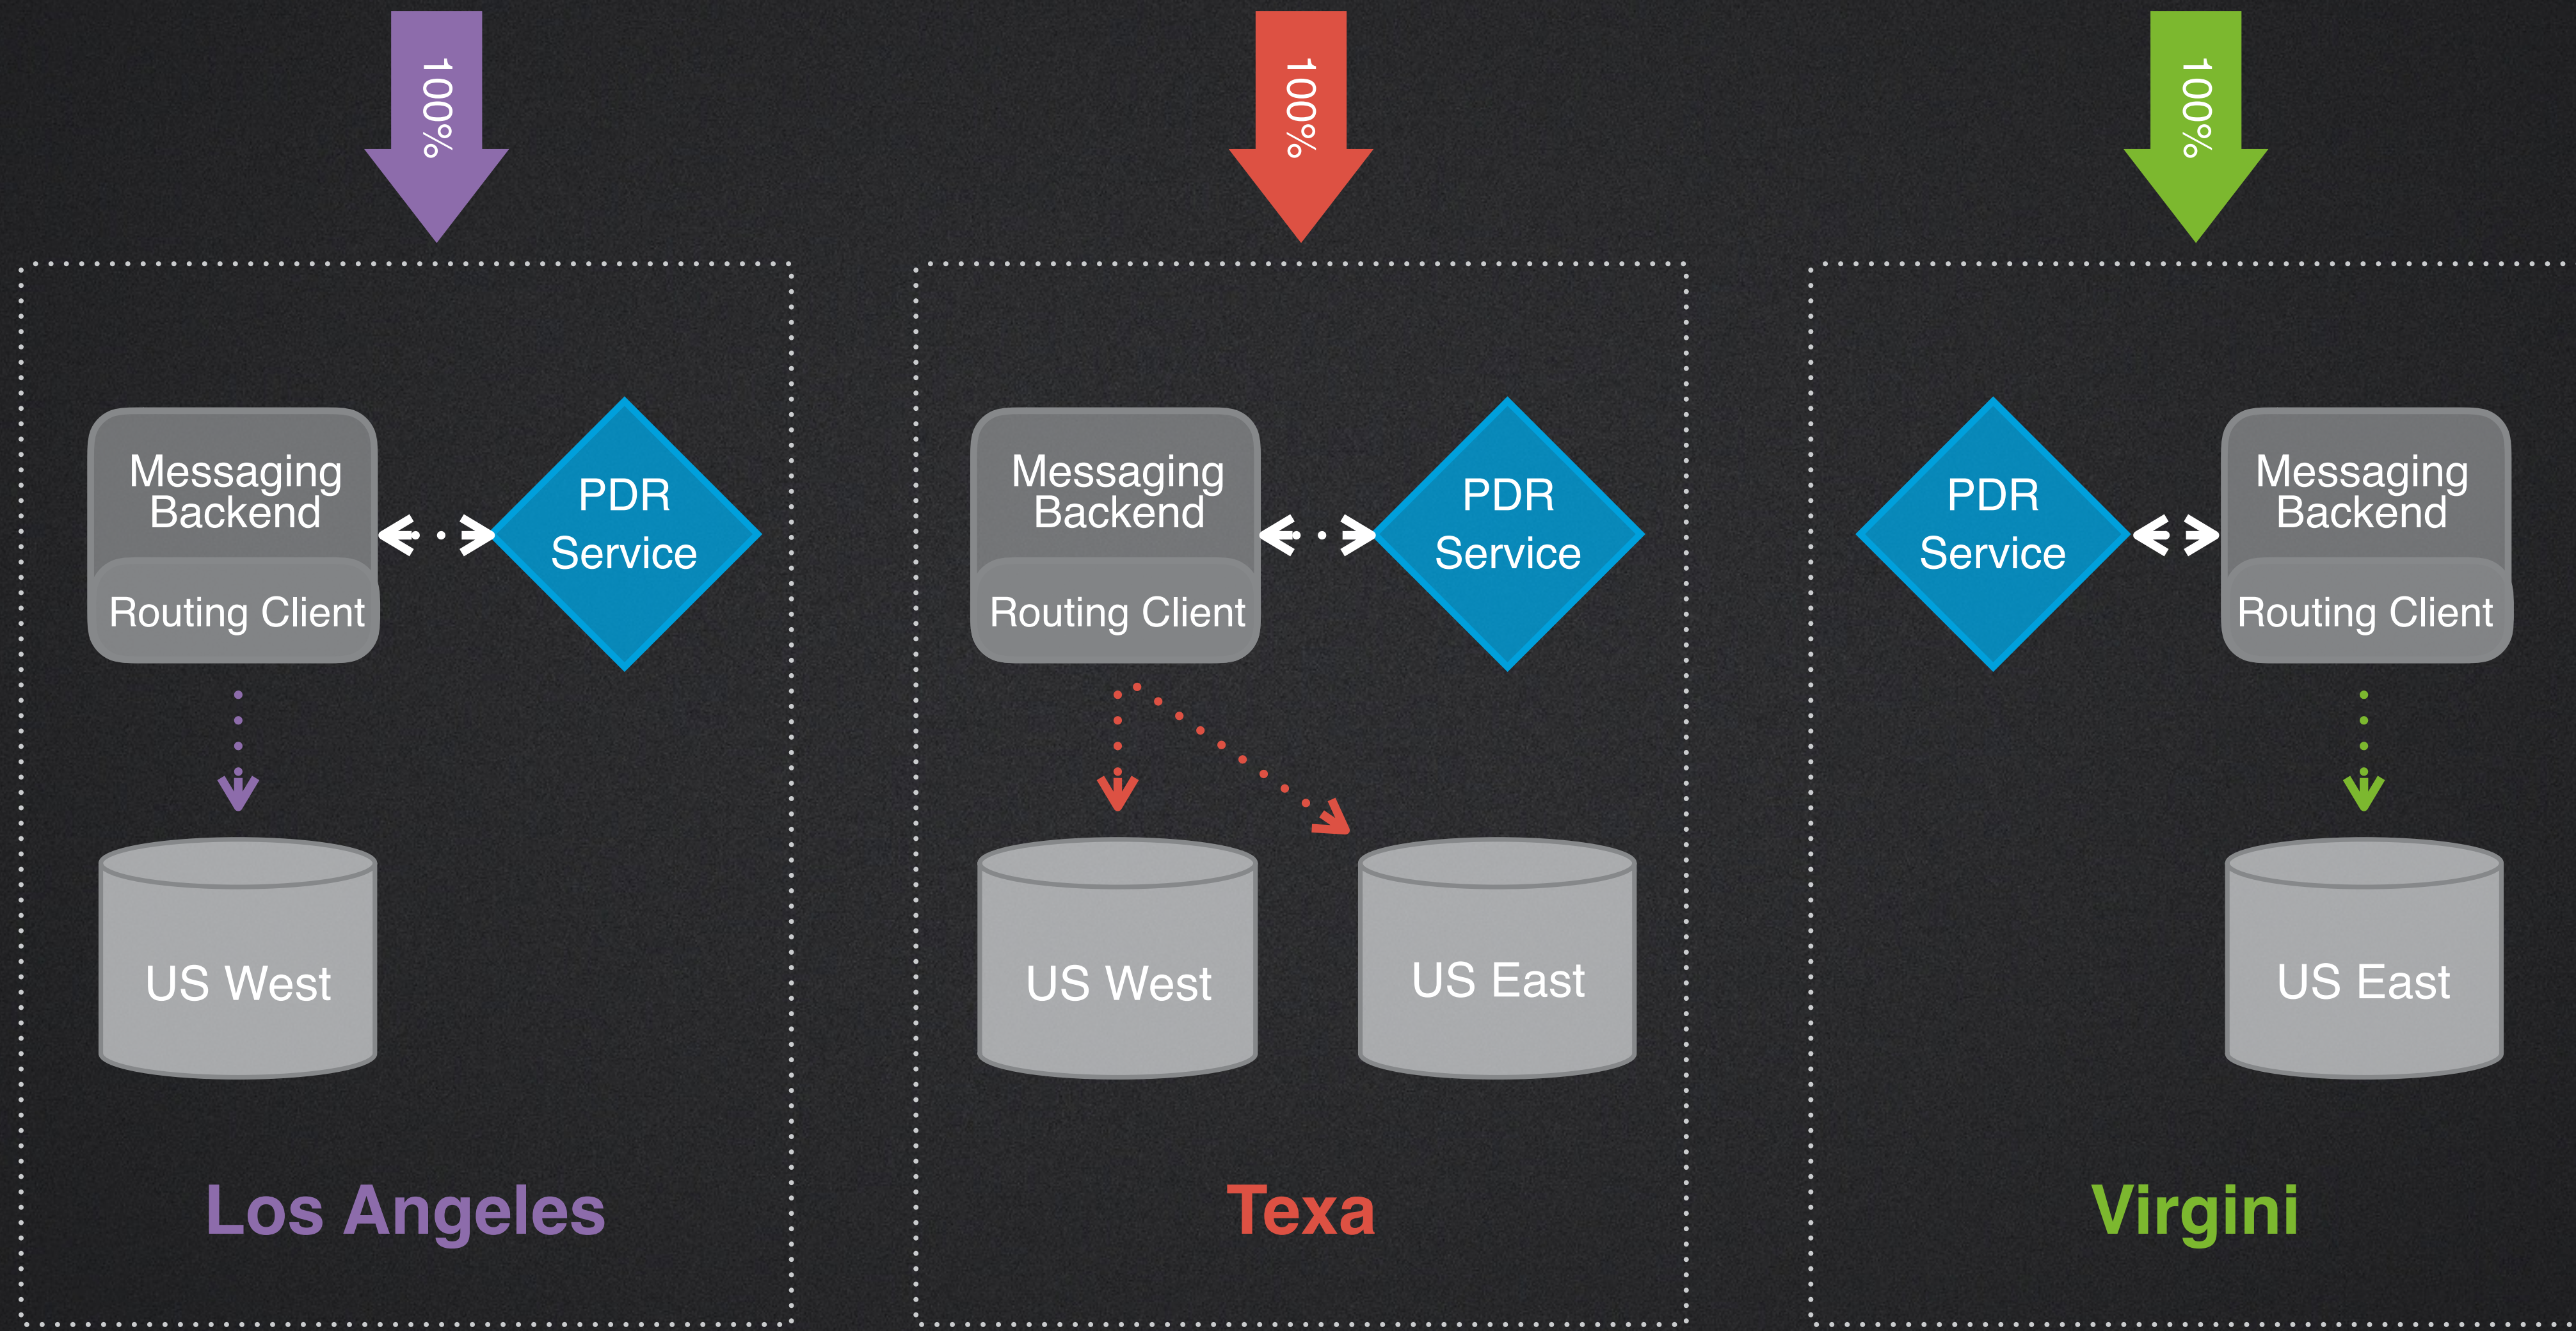
# How PDR works



/MESSAGING-US\_EAST/MailboxDB/Mailbox/m-405

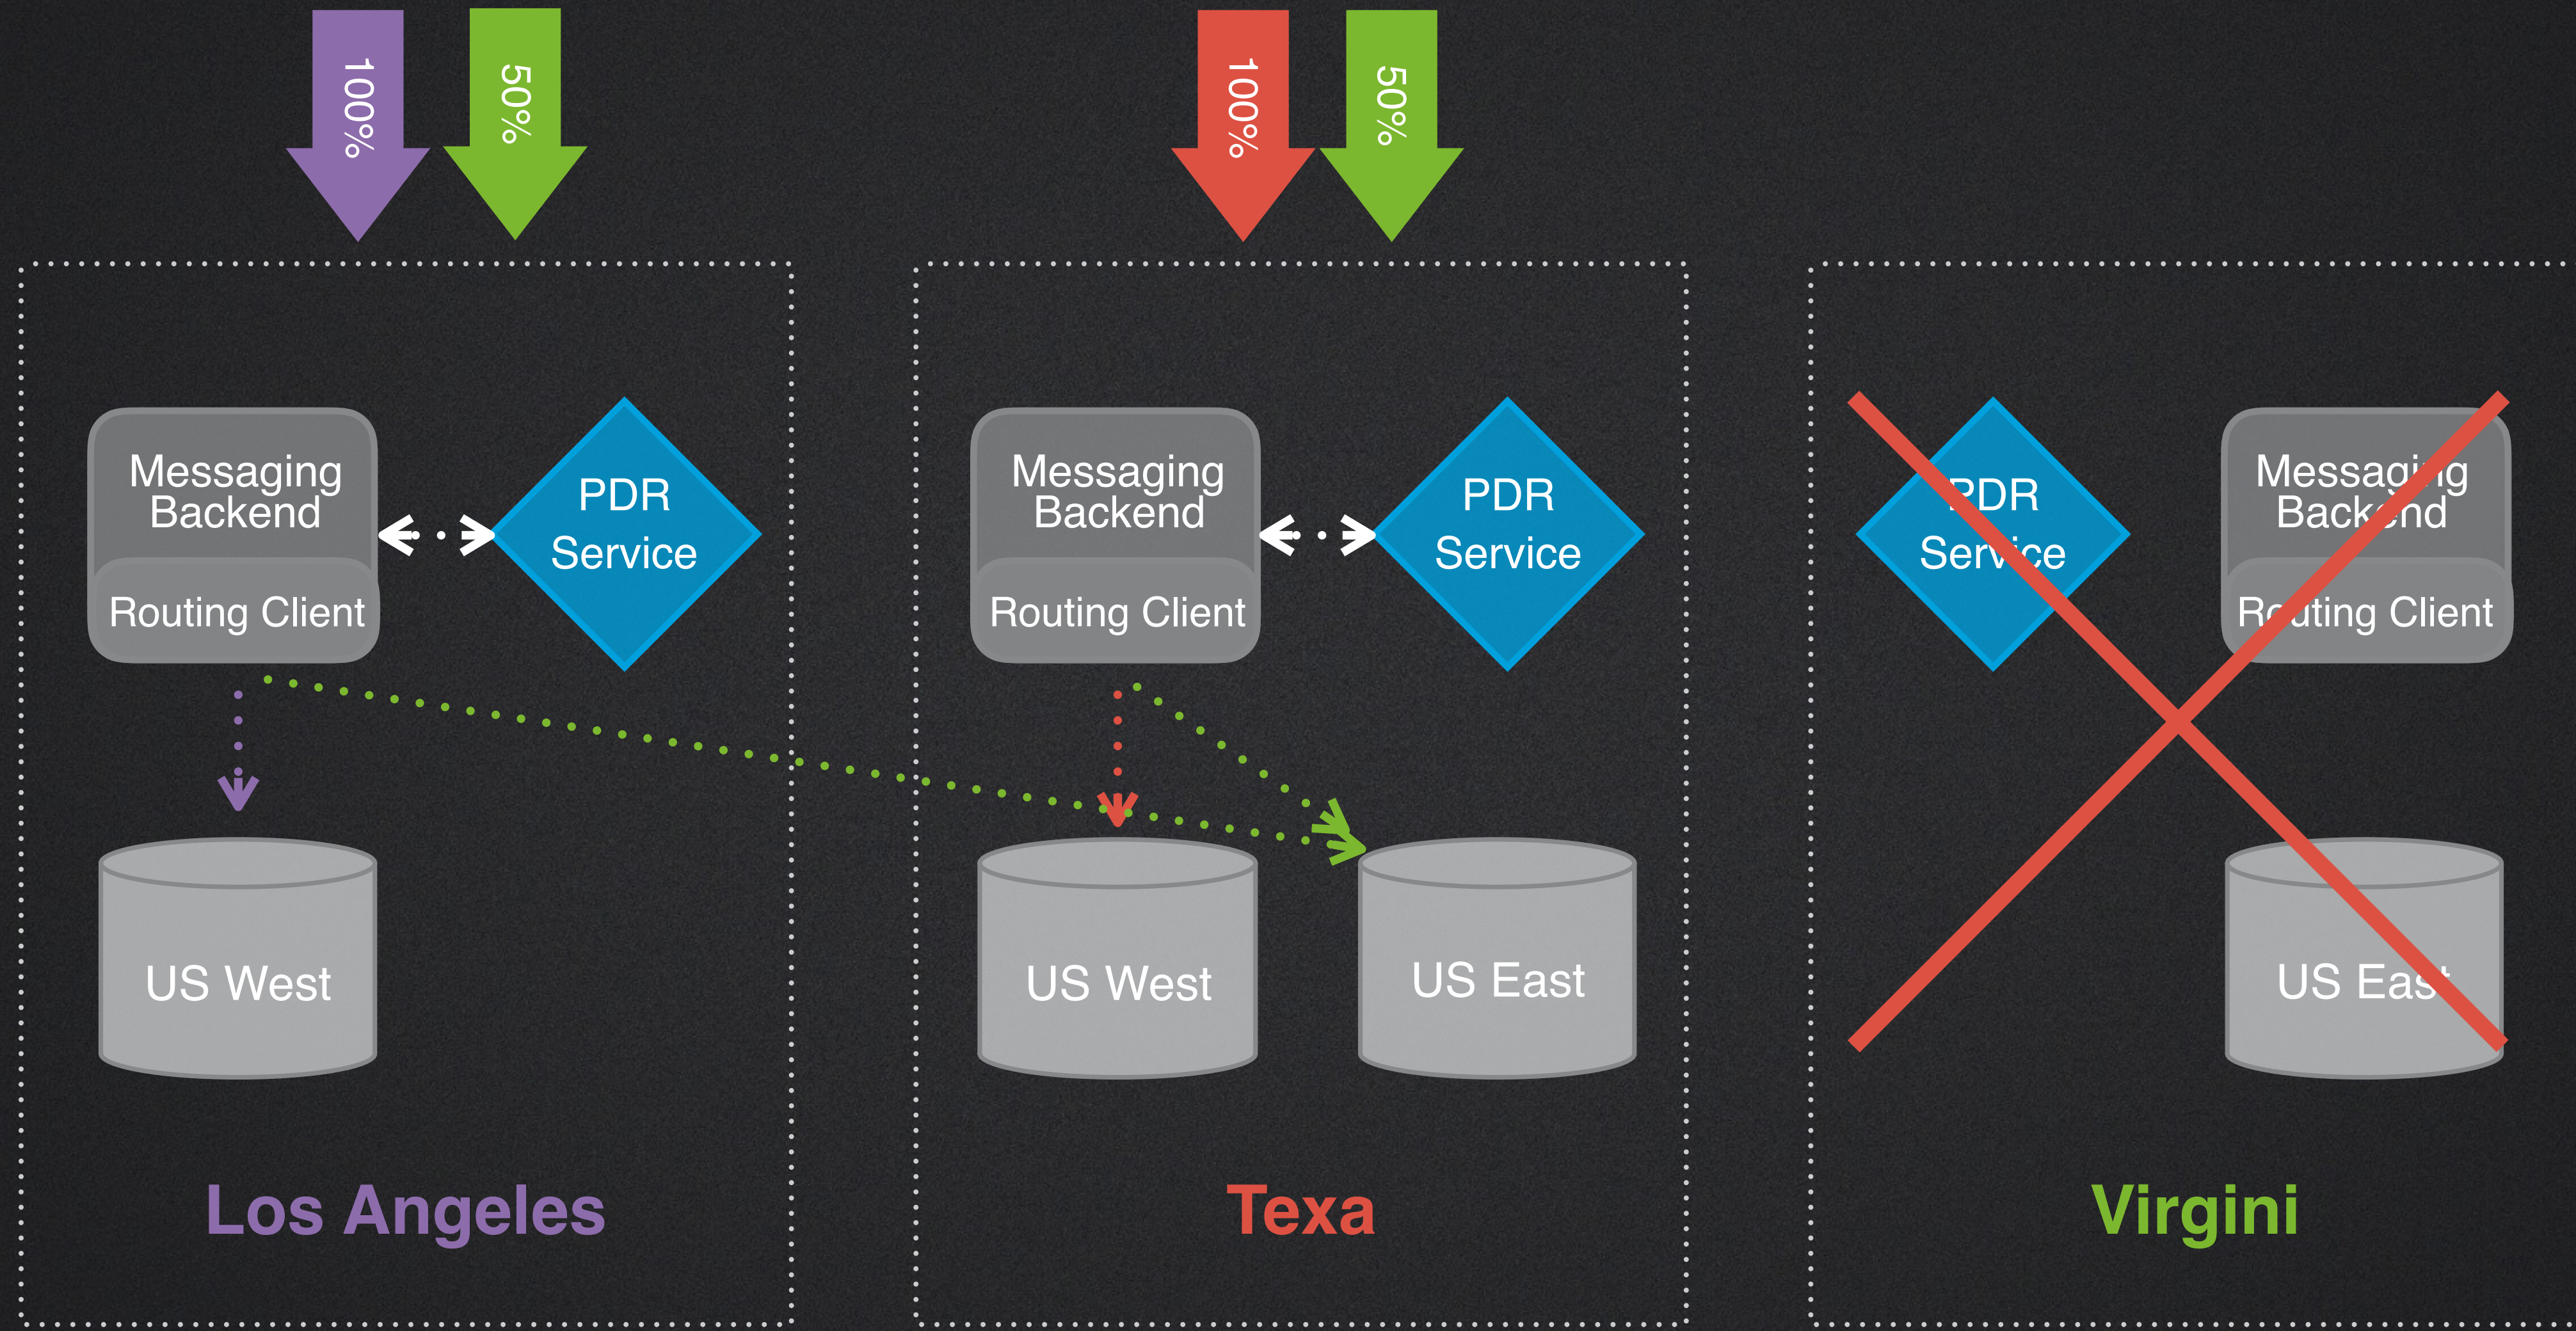


# What happens when we shift traffic





# What happens when we shift traffic

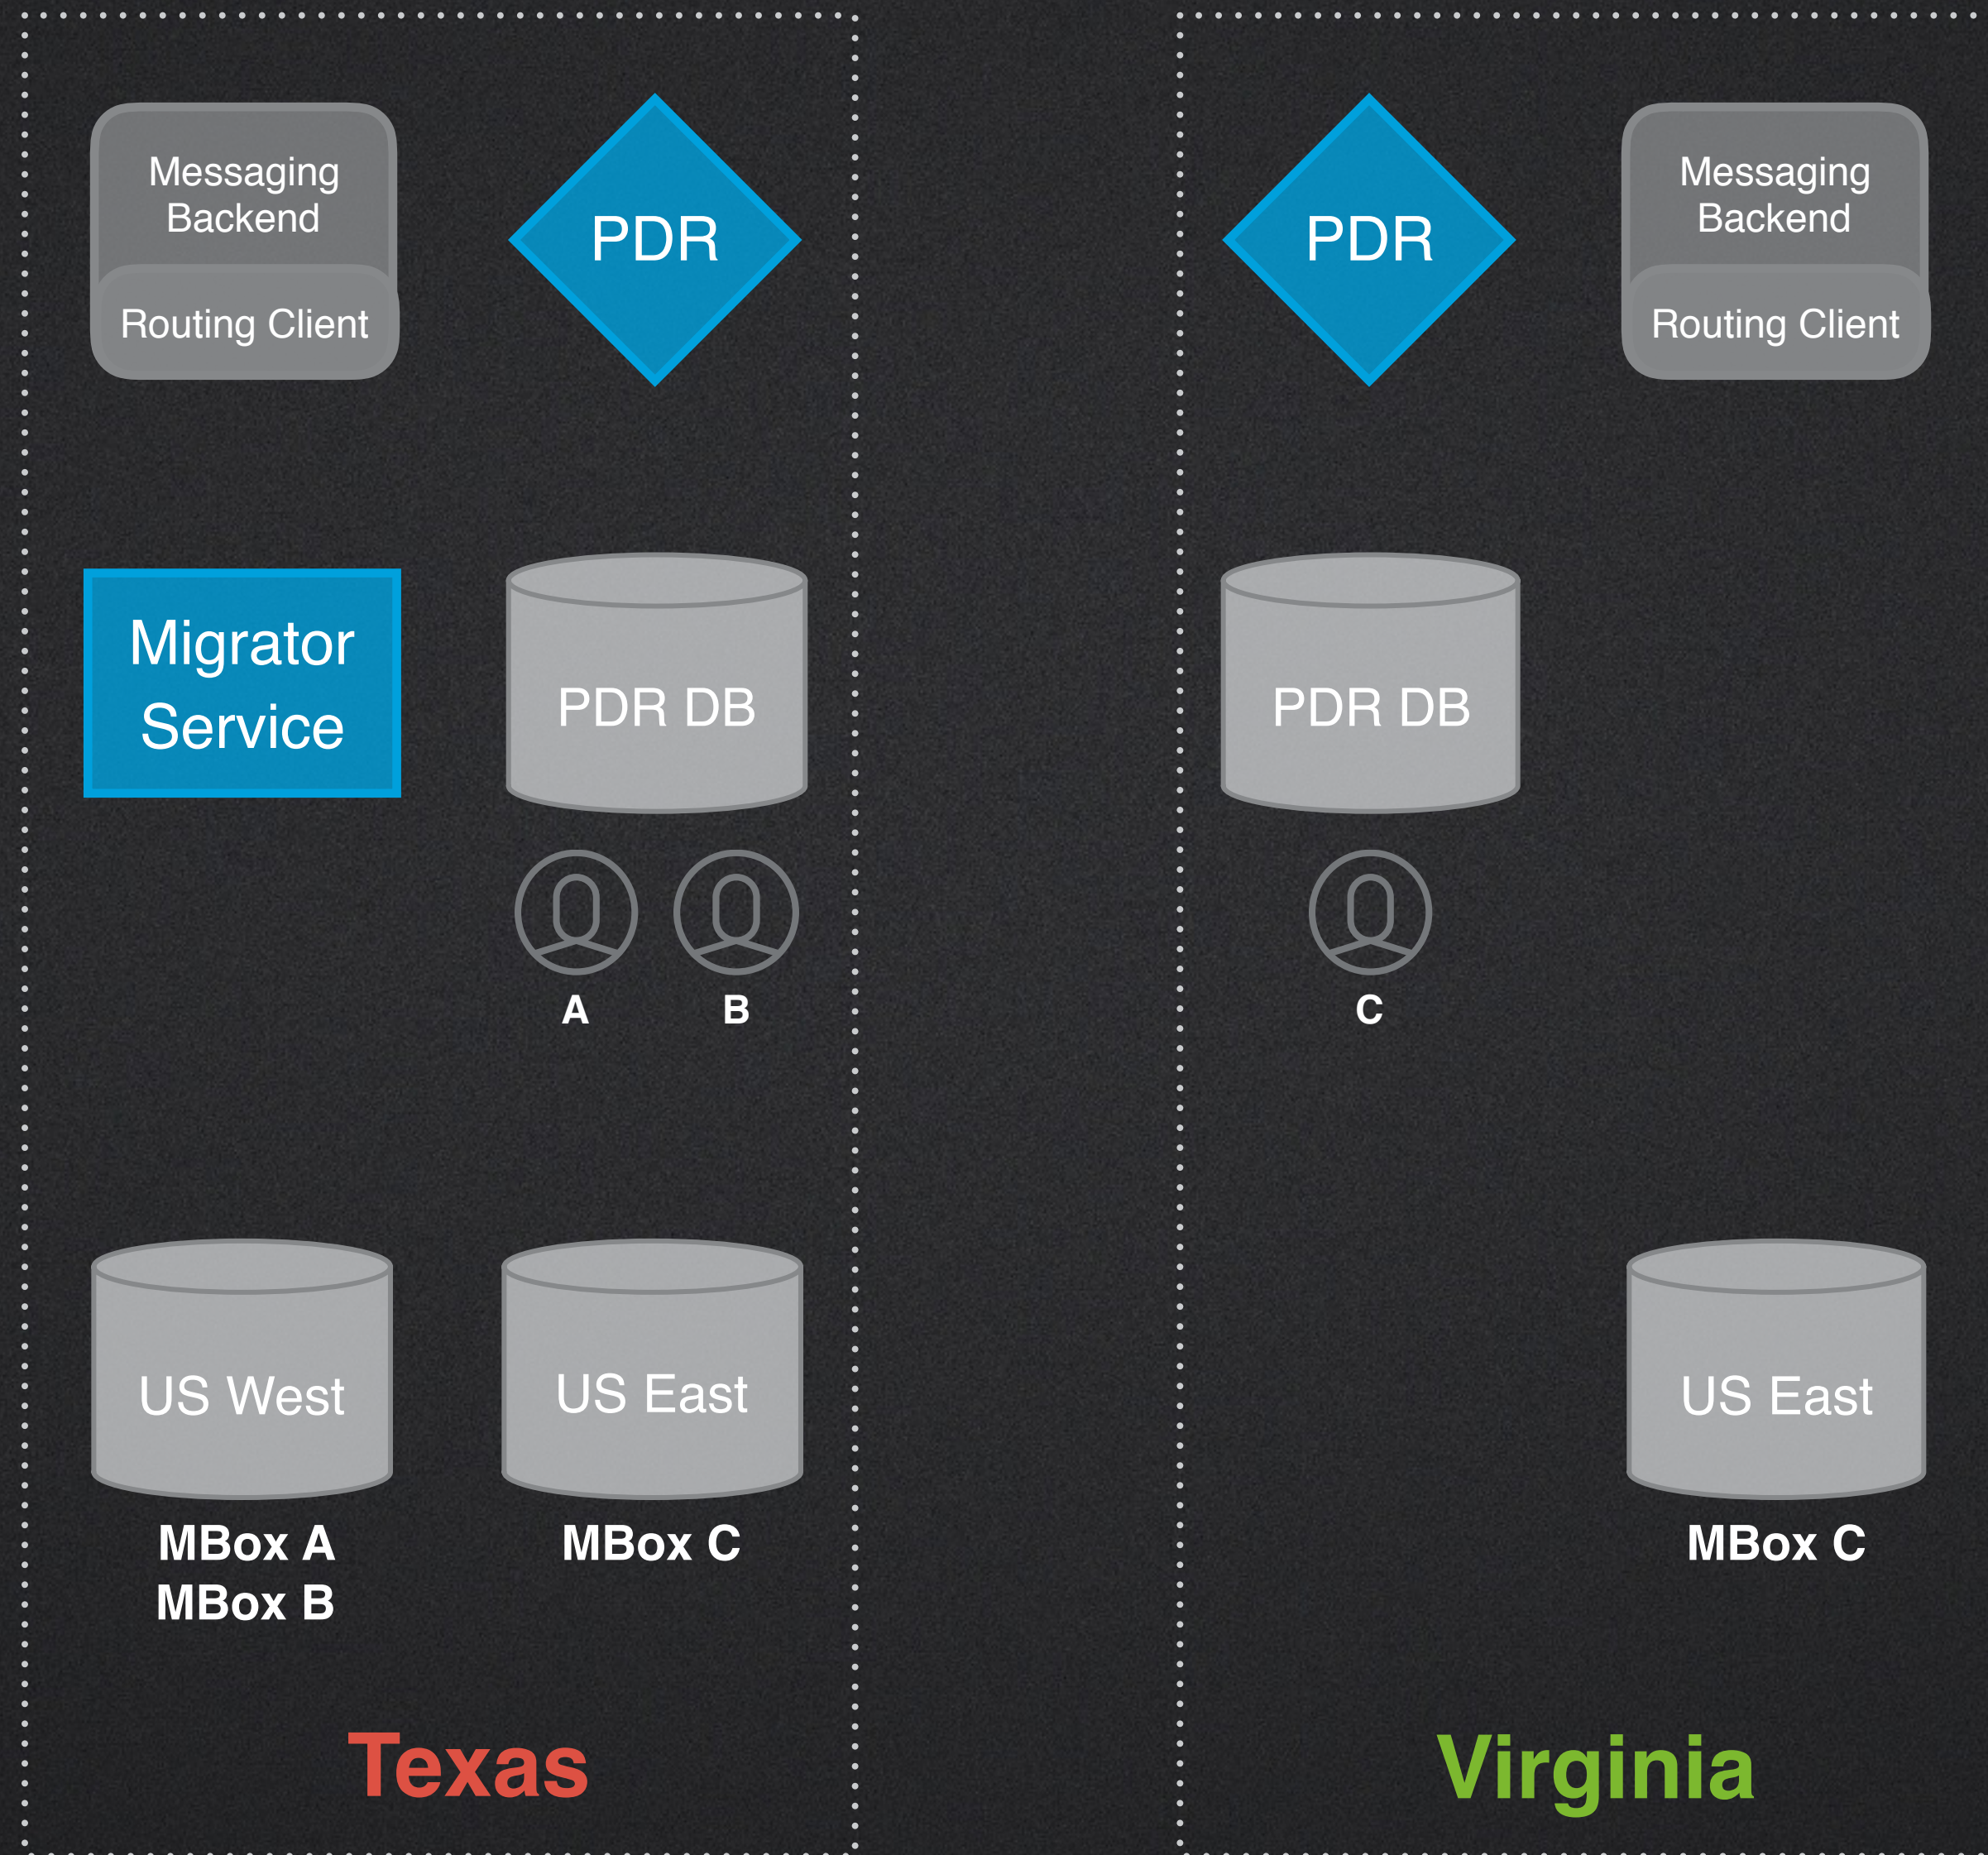






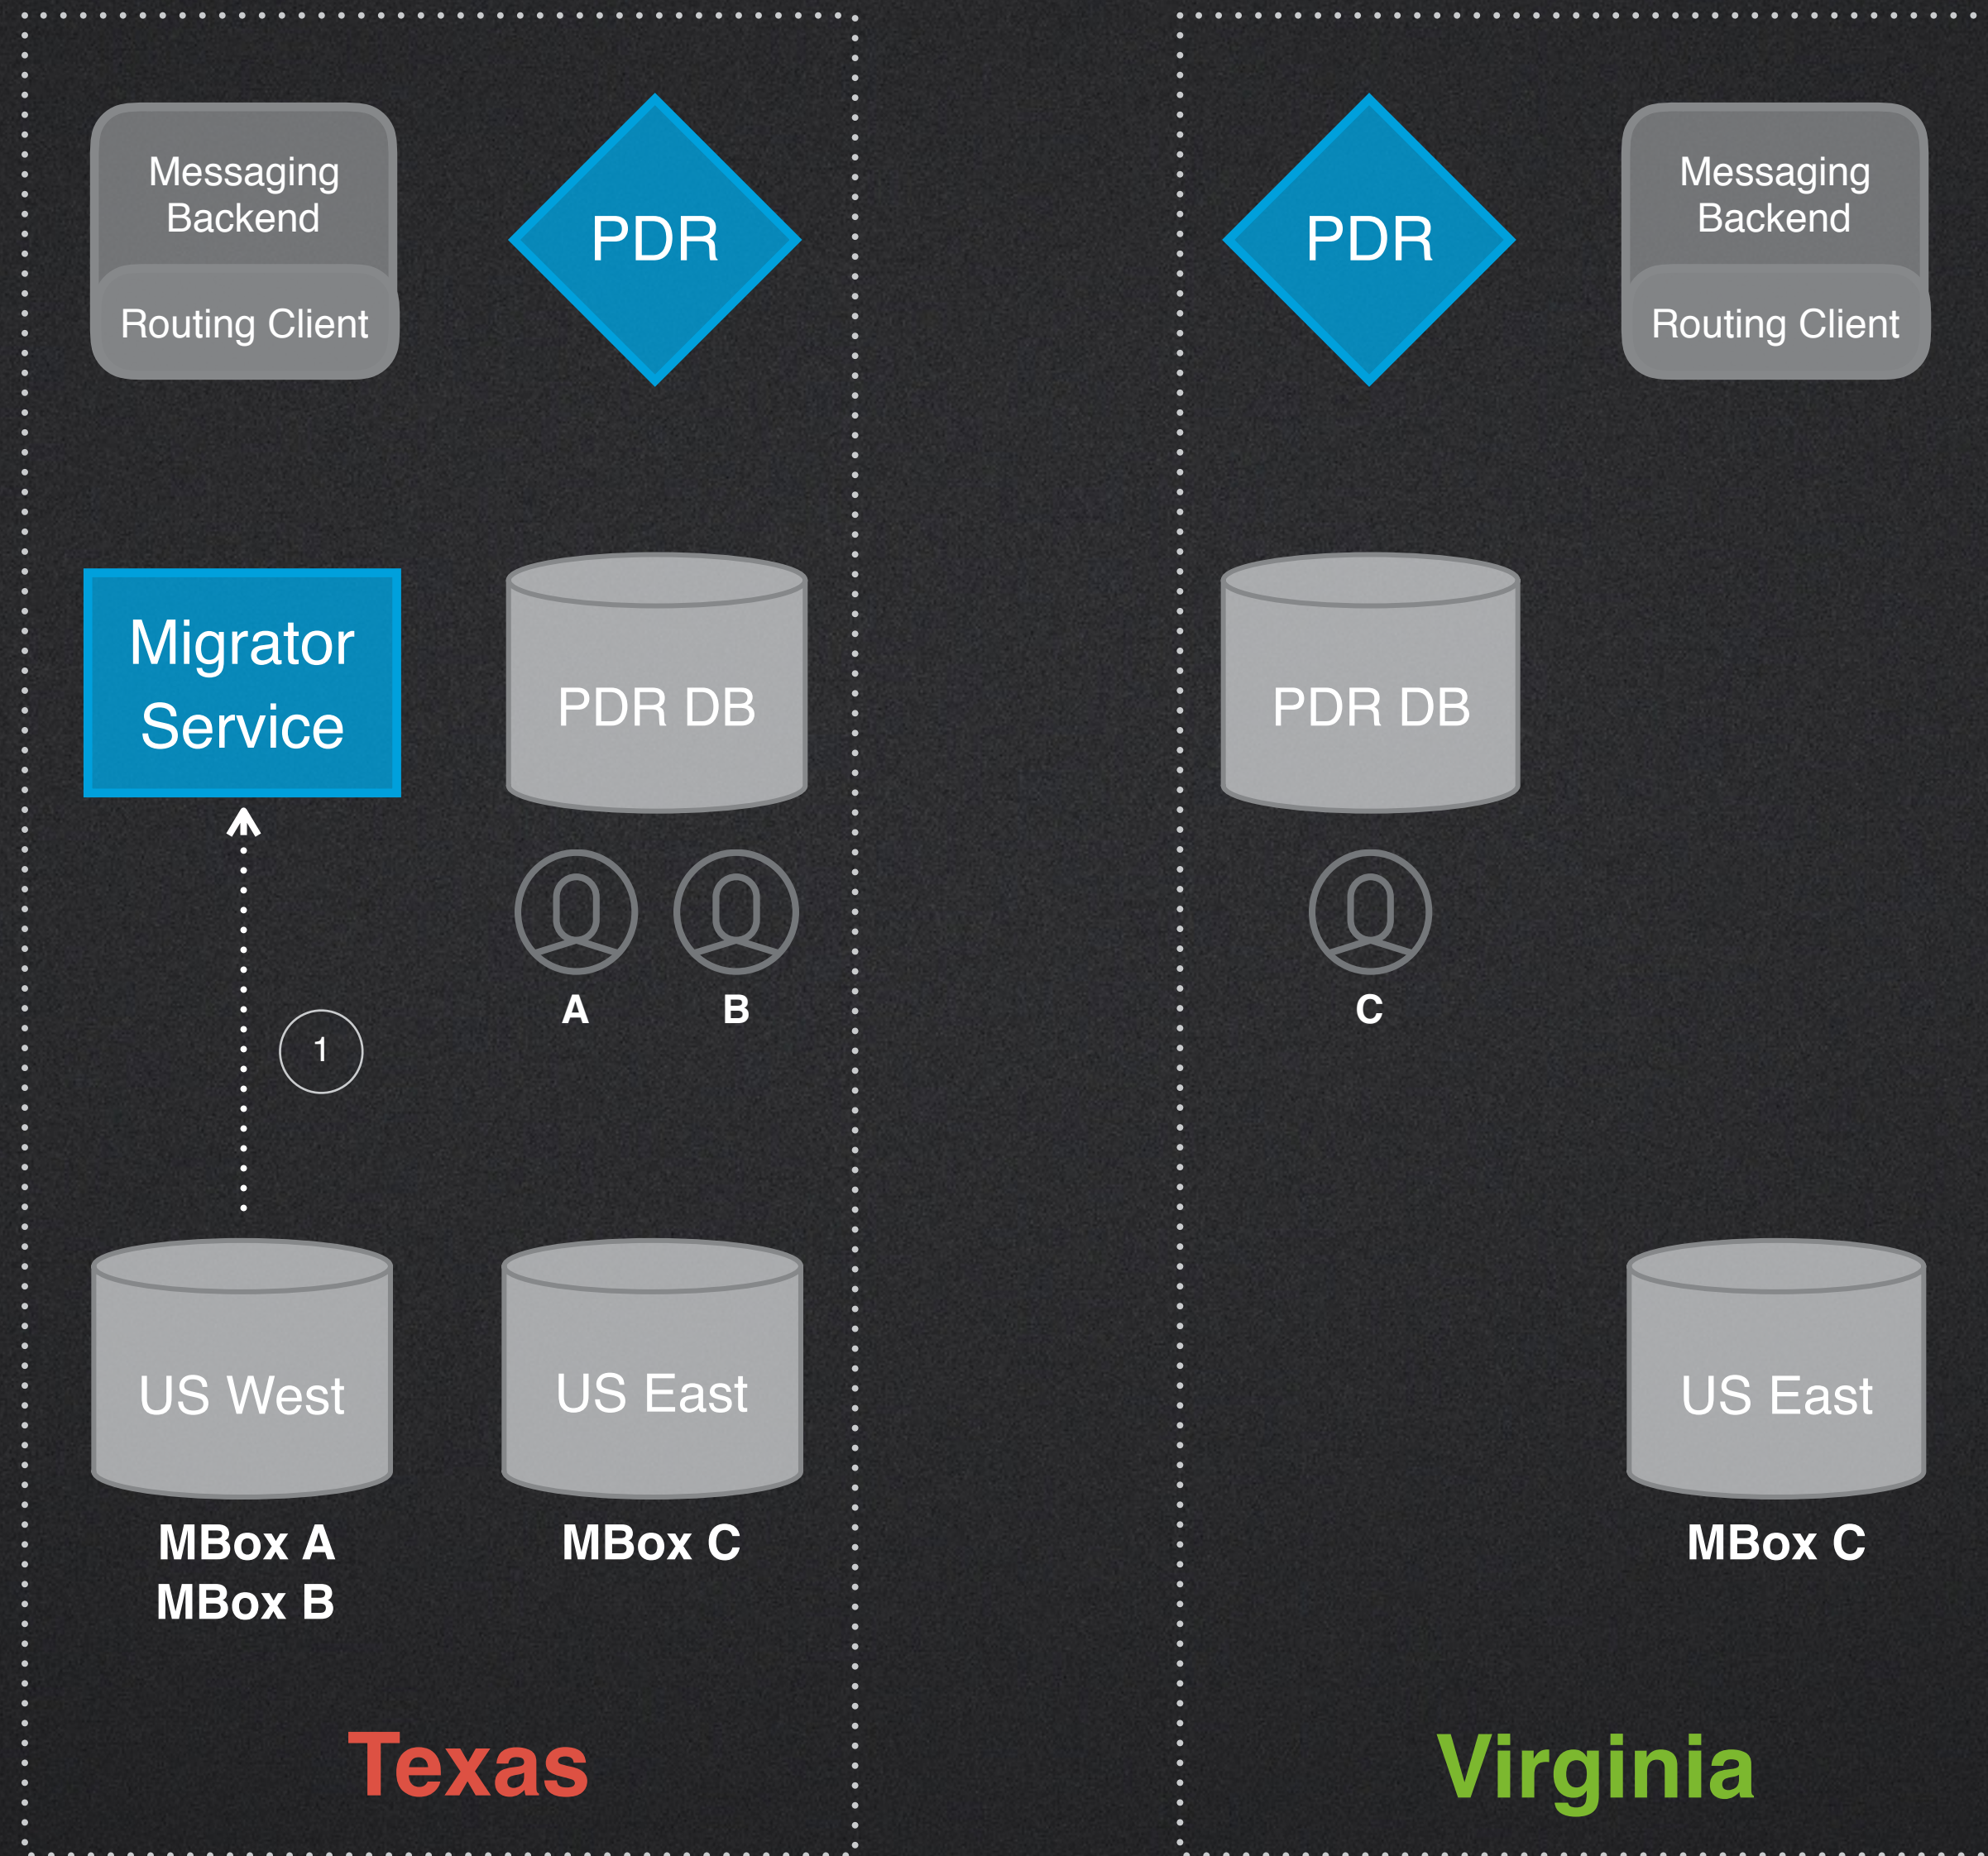


# Mailbox Migration



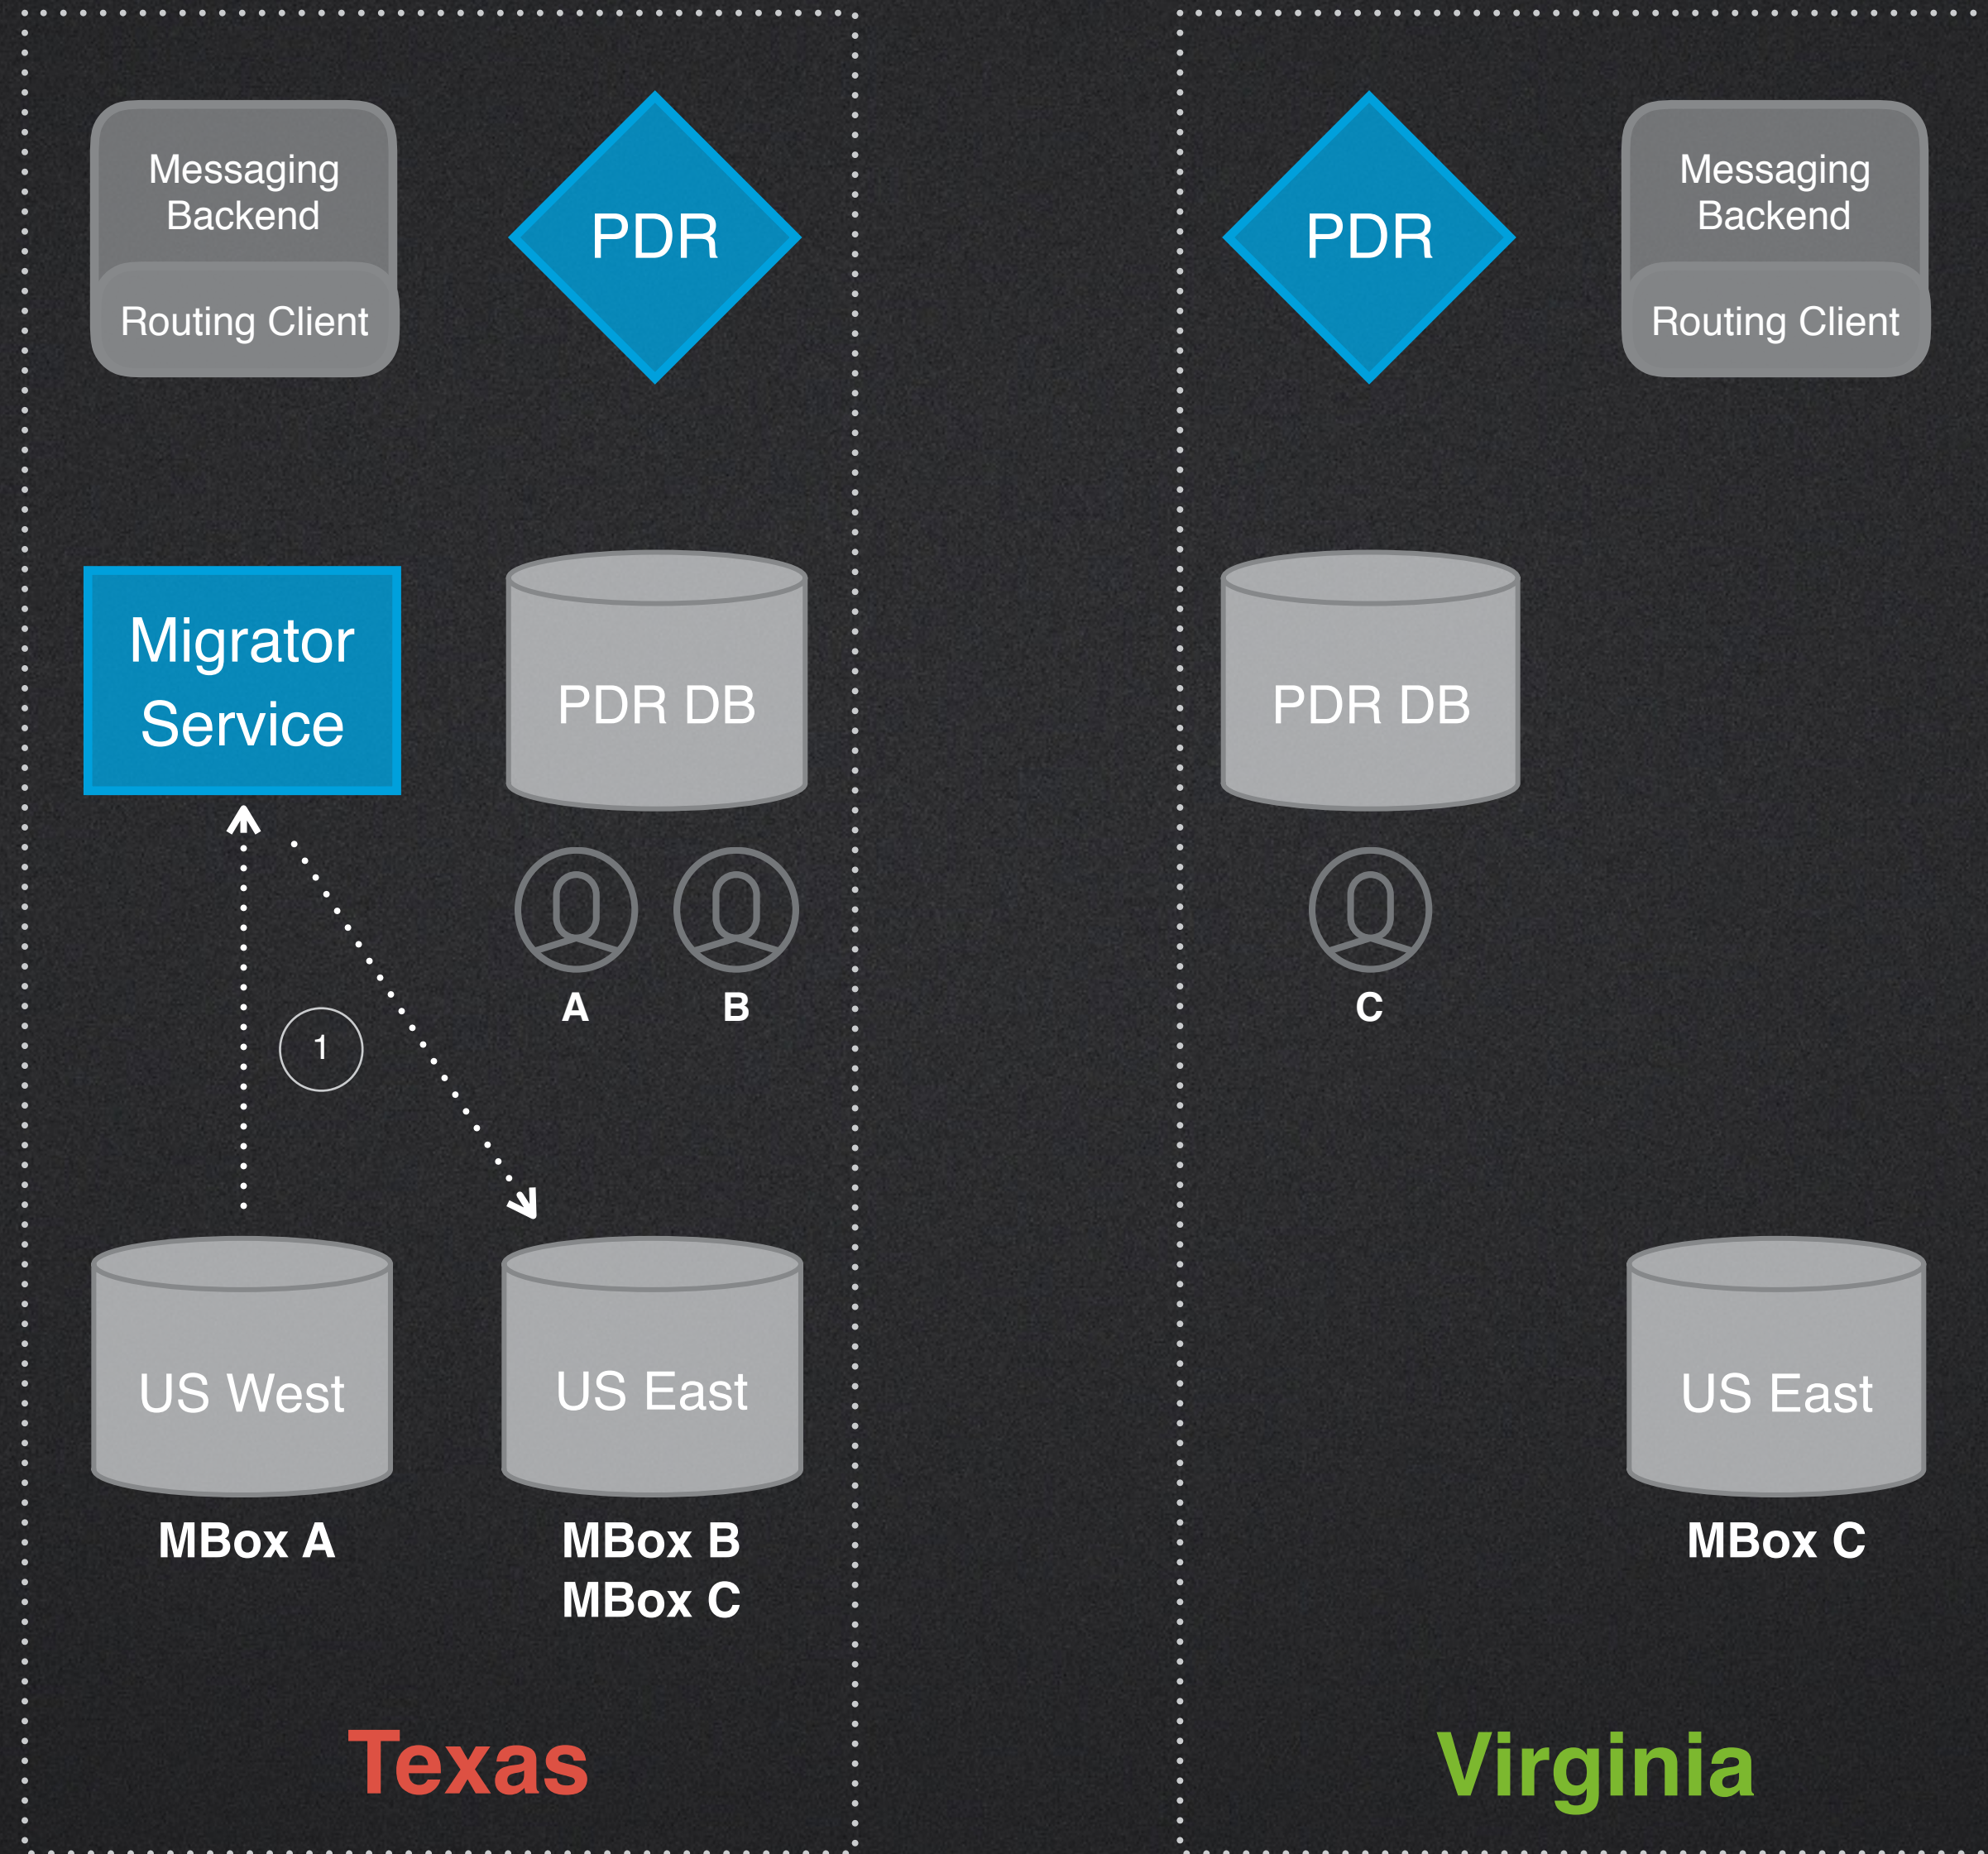


# Mailbox Migration



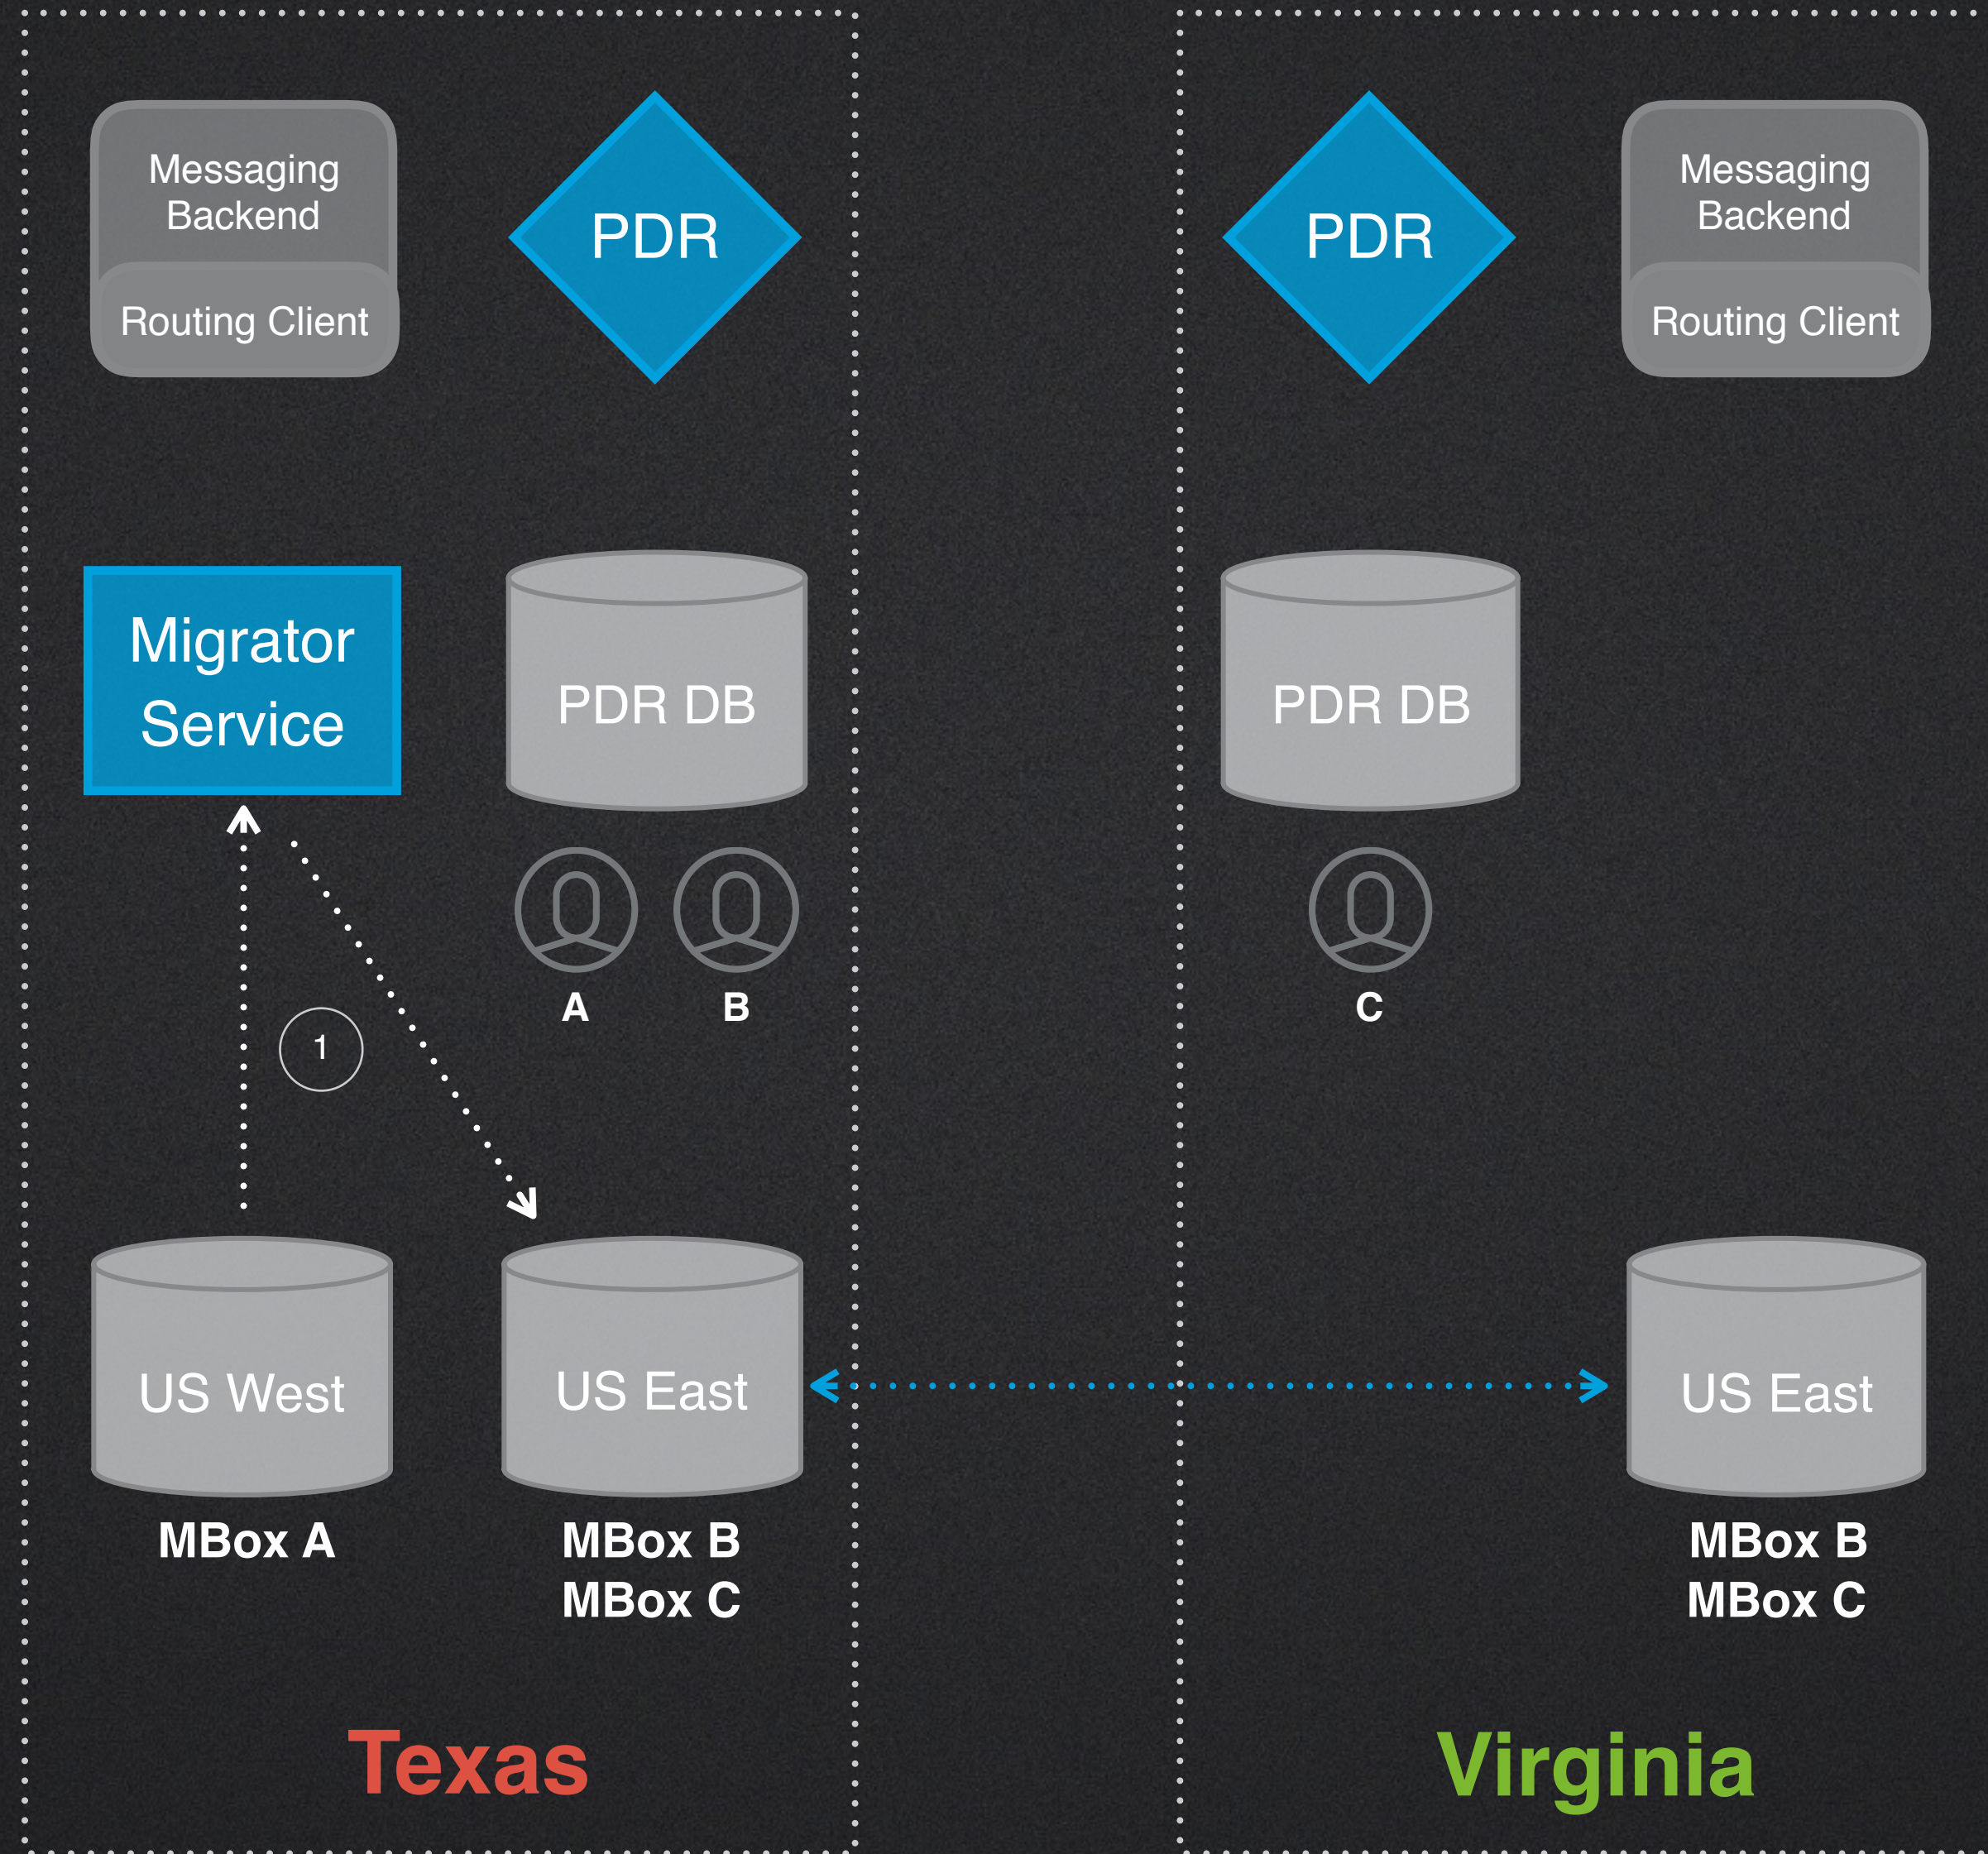


# Mailbox Migration



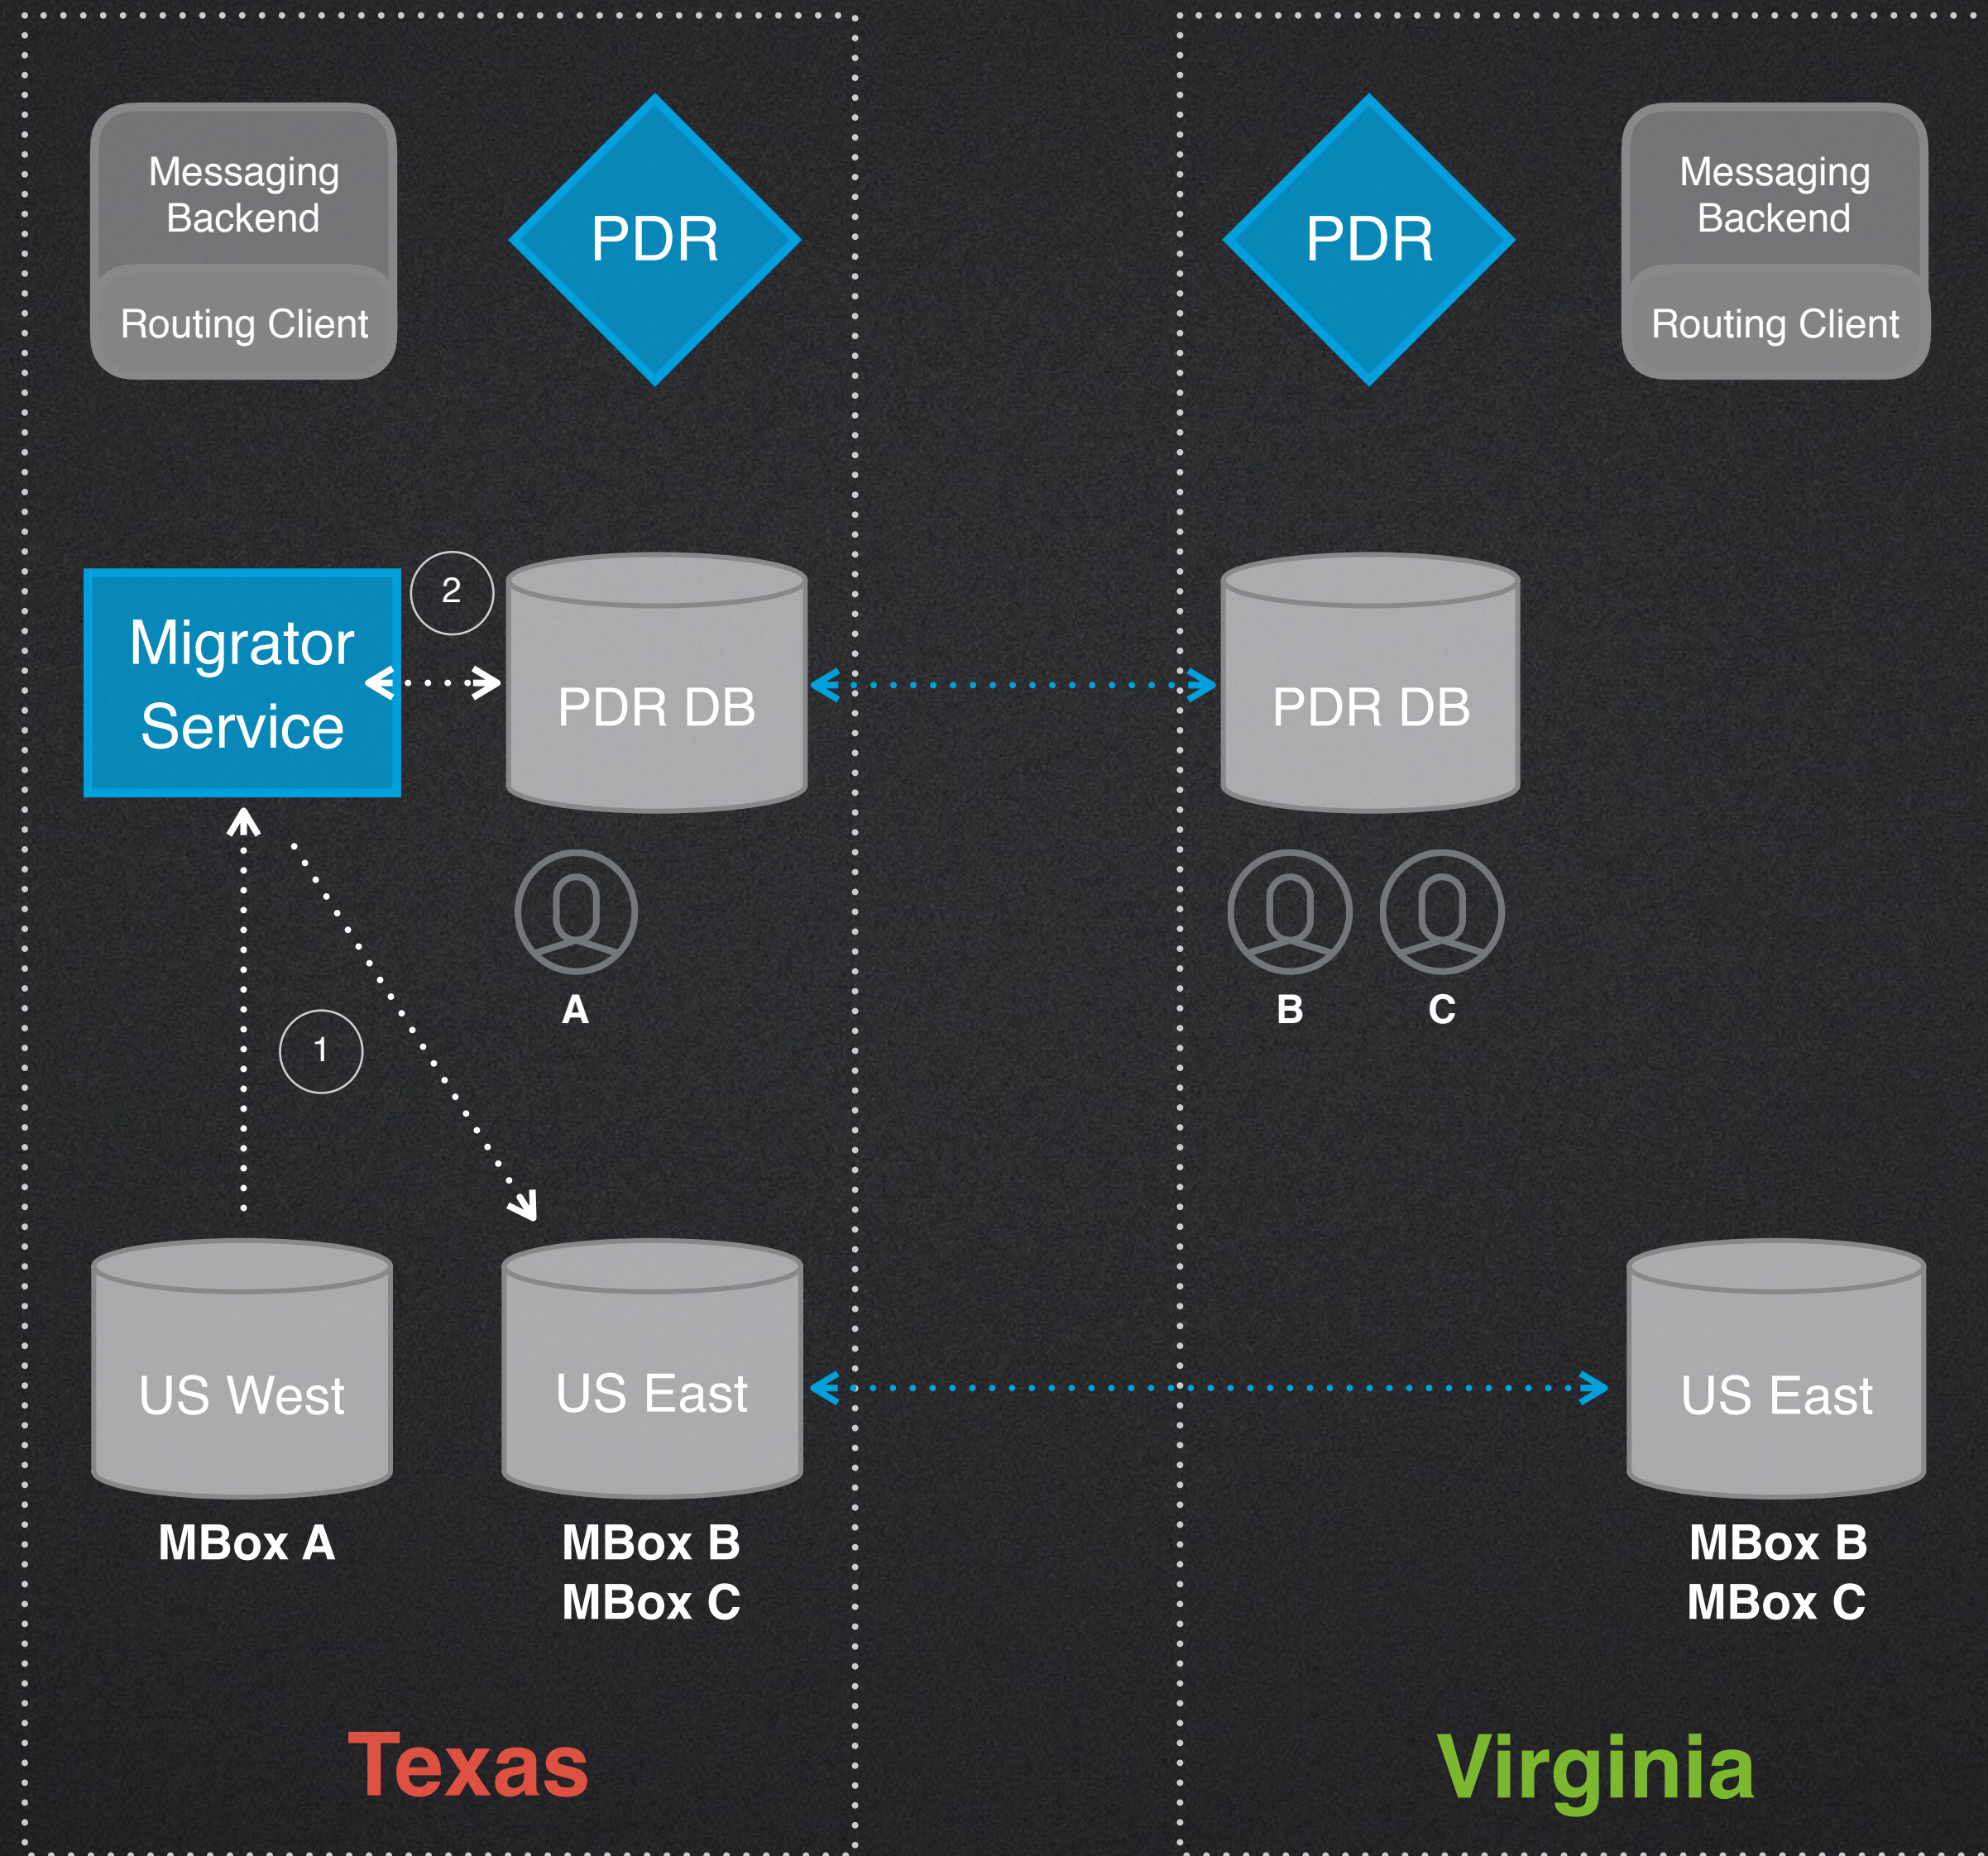


# Mailbox Migration



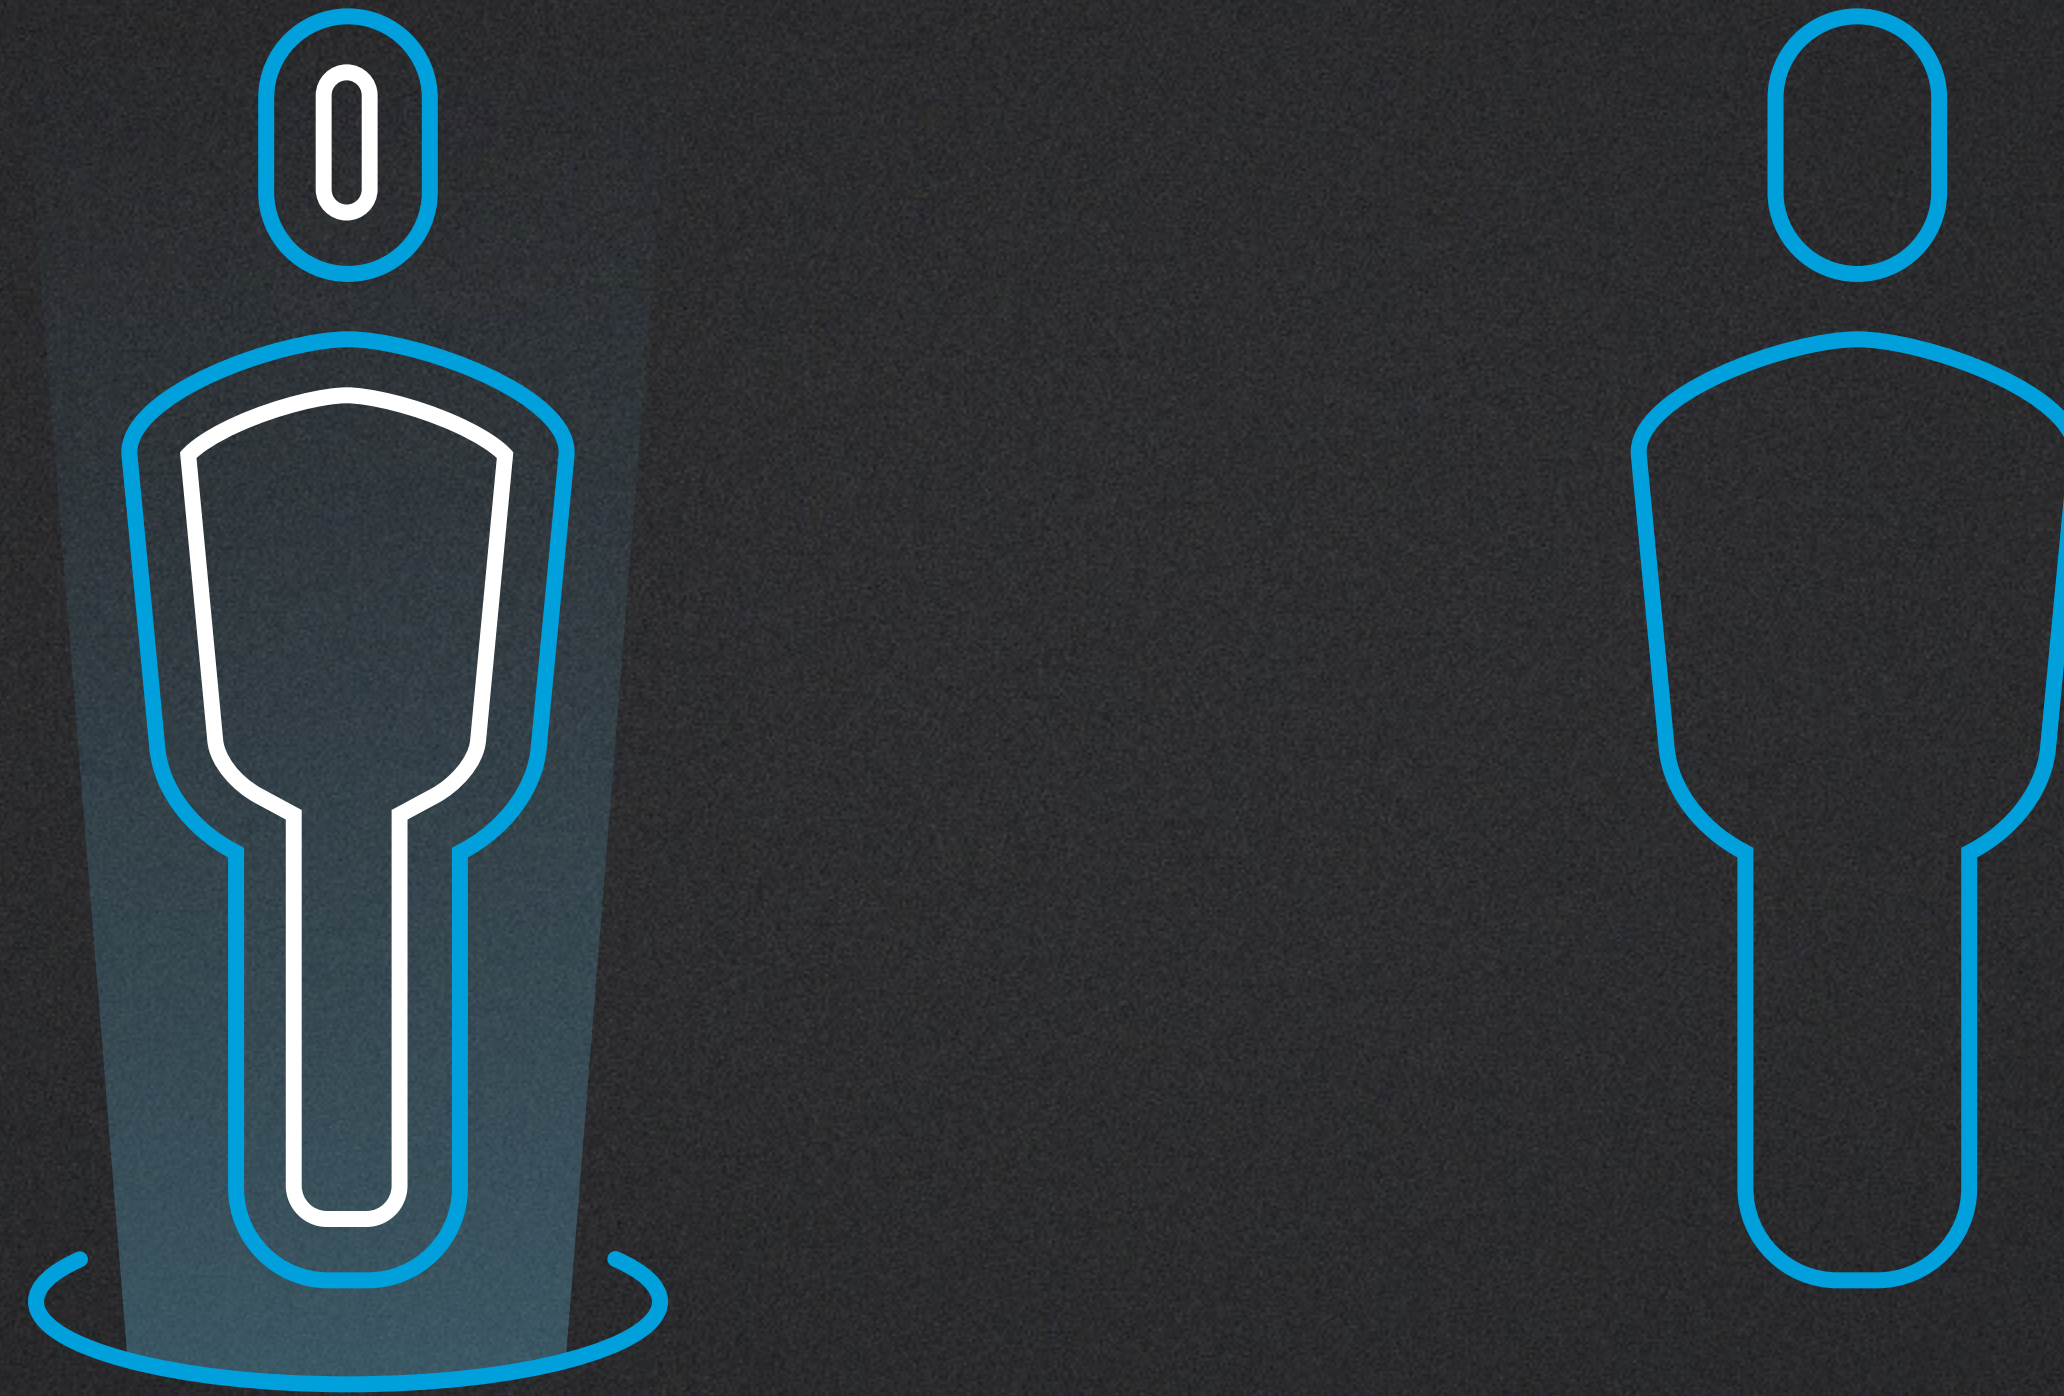


# Mailbox Migration





You only need two copies of Personal Data





3

# Cache Invalidation & Replication Conflicts



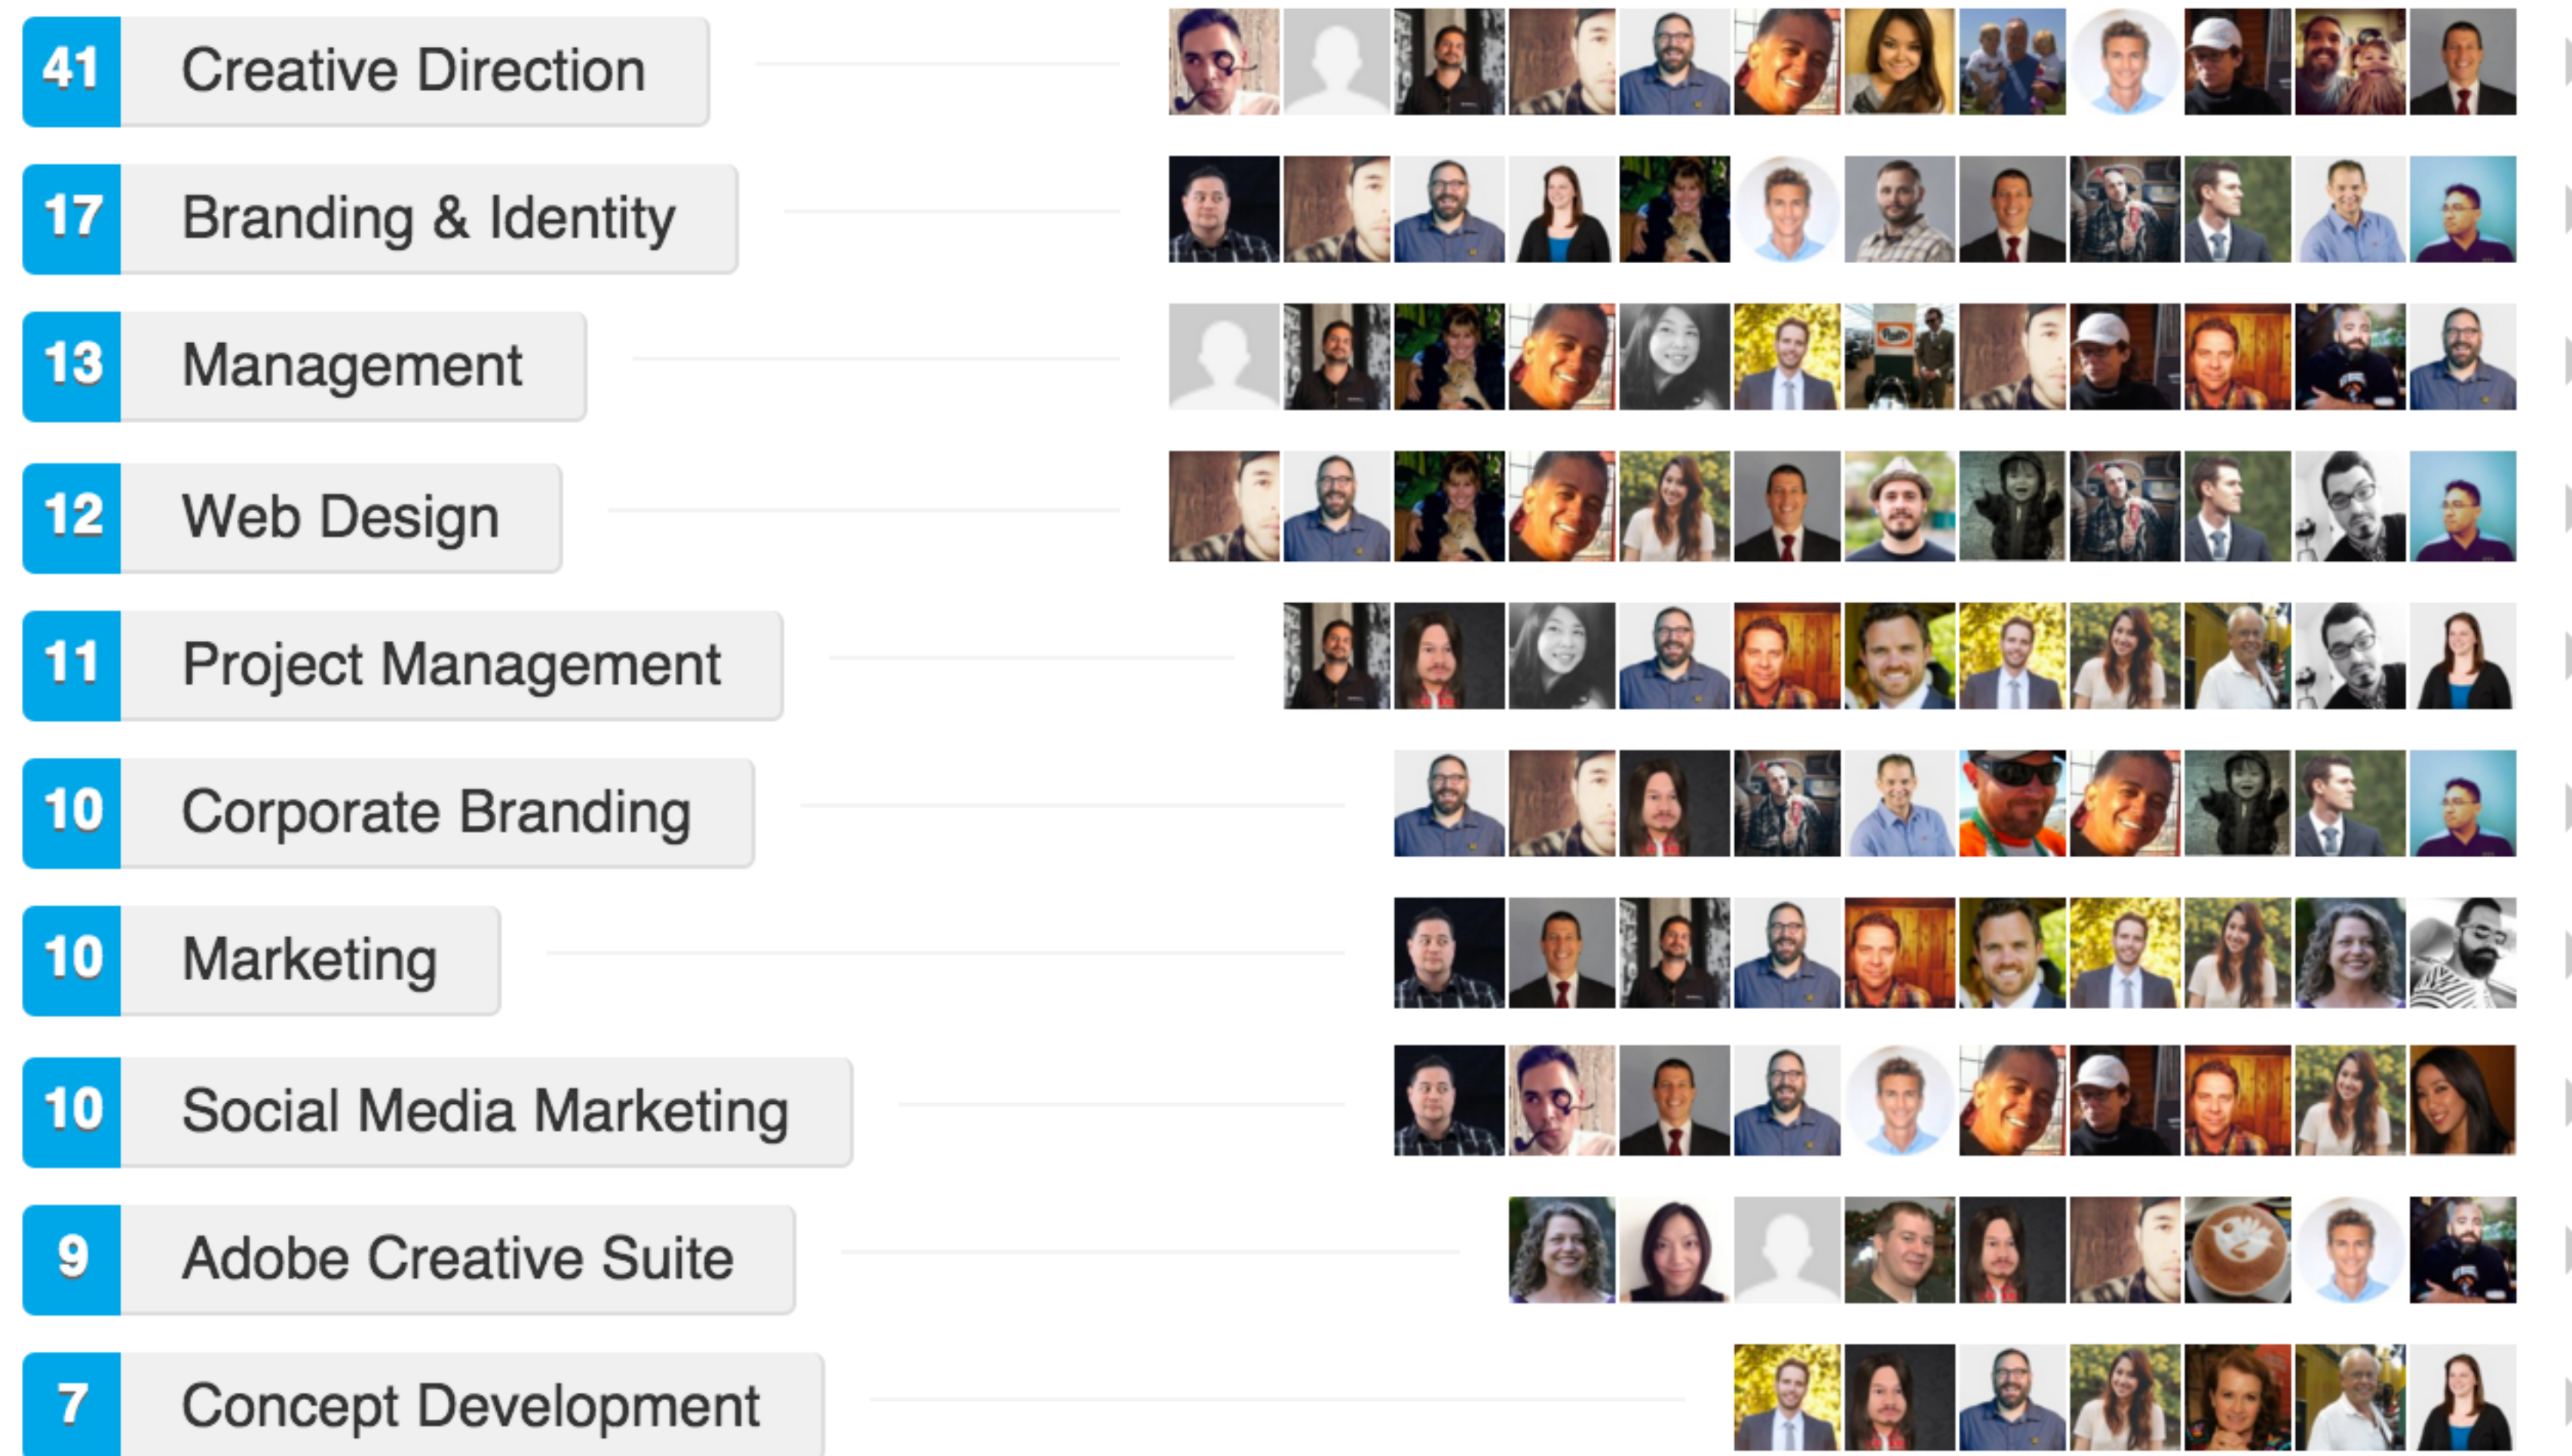
# Endorsements





# Skills & Endorsements

## Top Skills





8.6

---

Billion





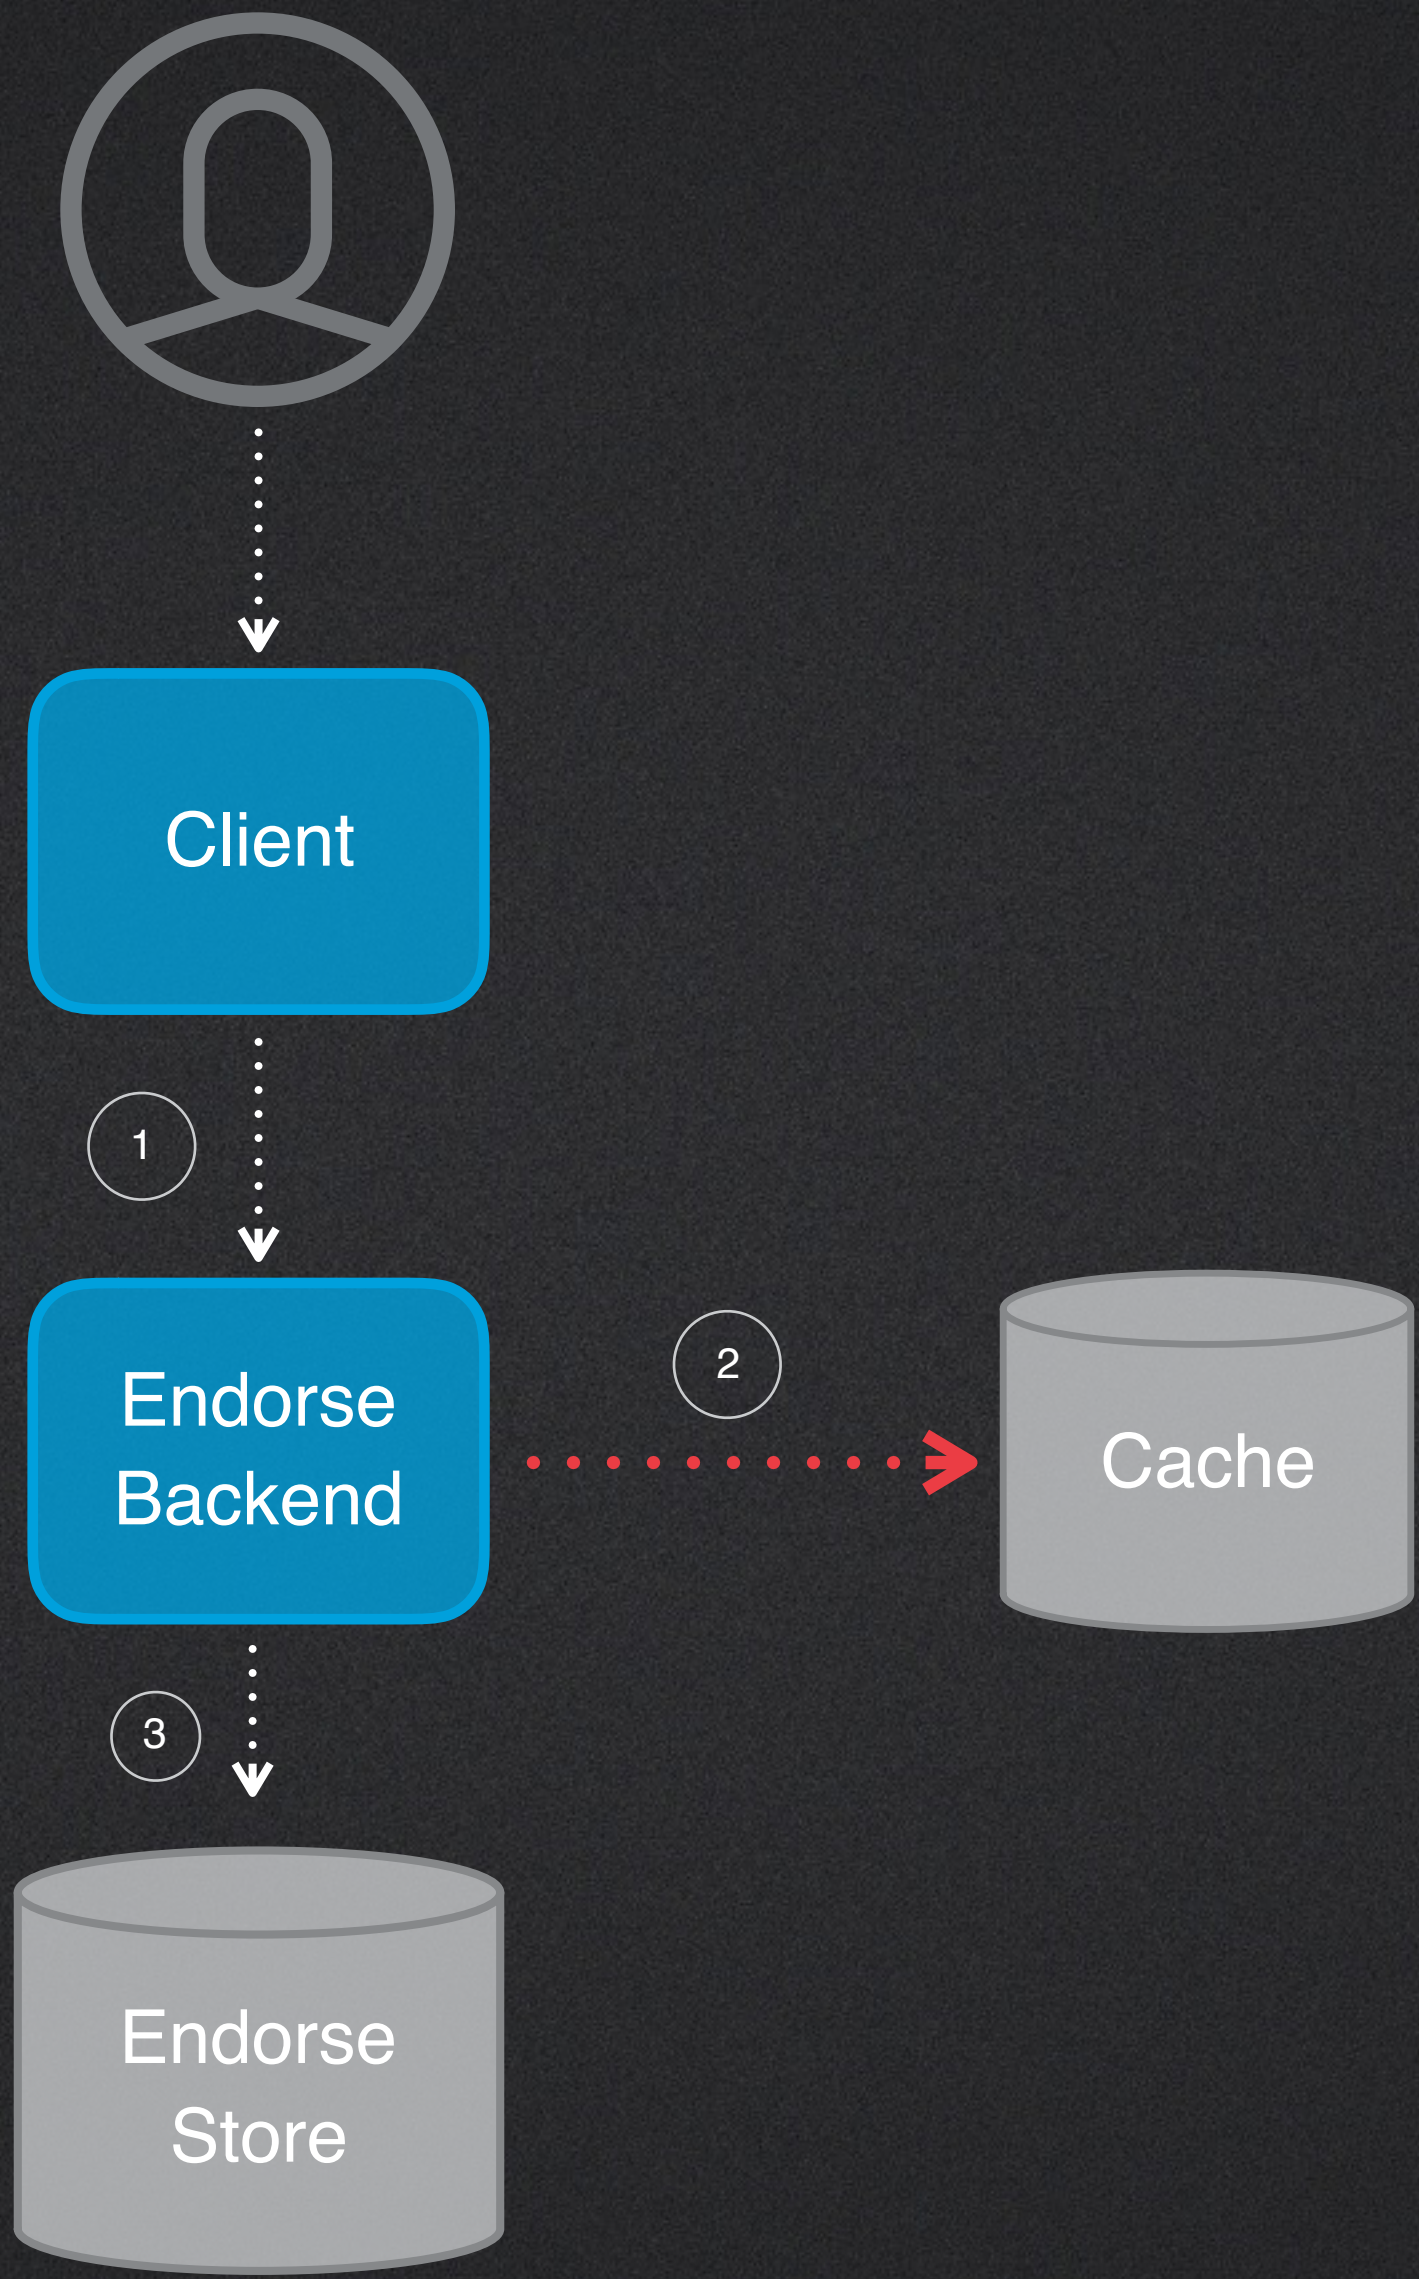
1



2





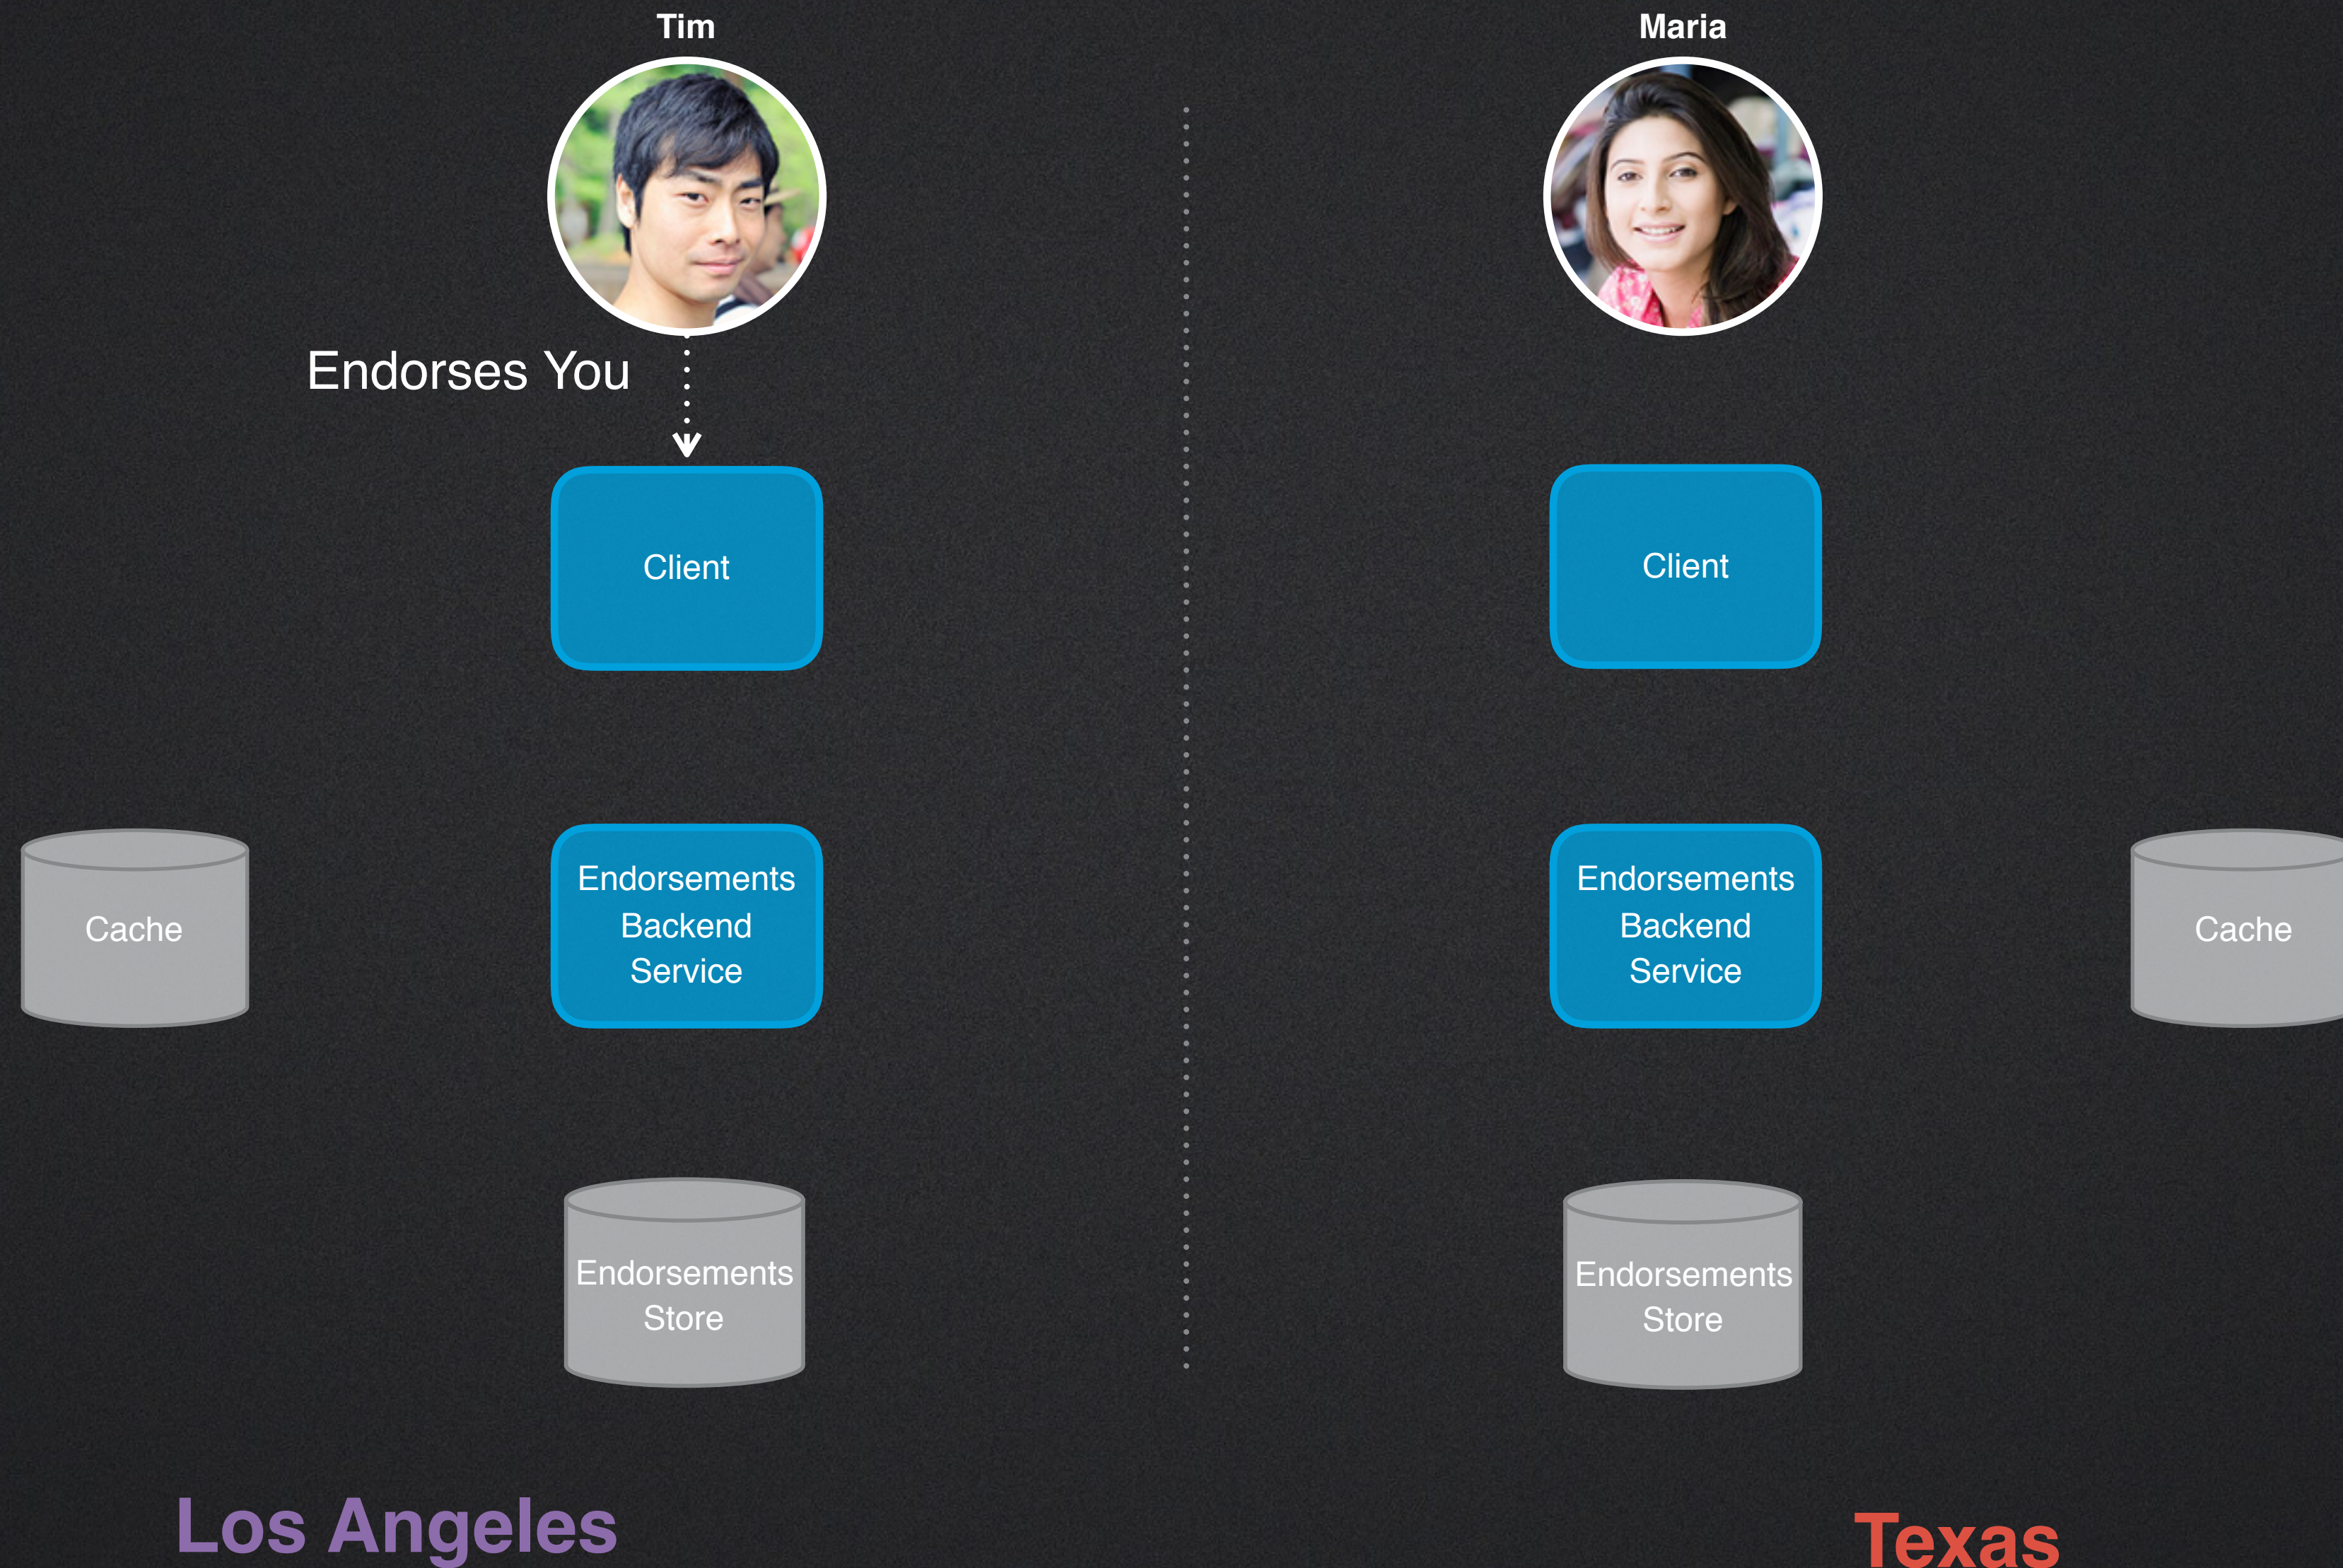




What could go wrong?

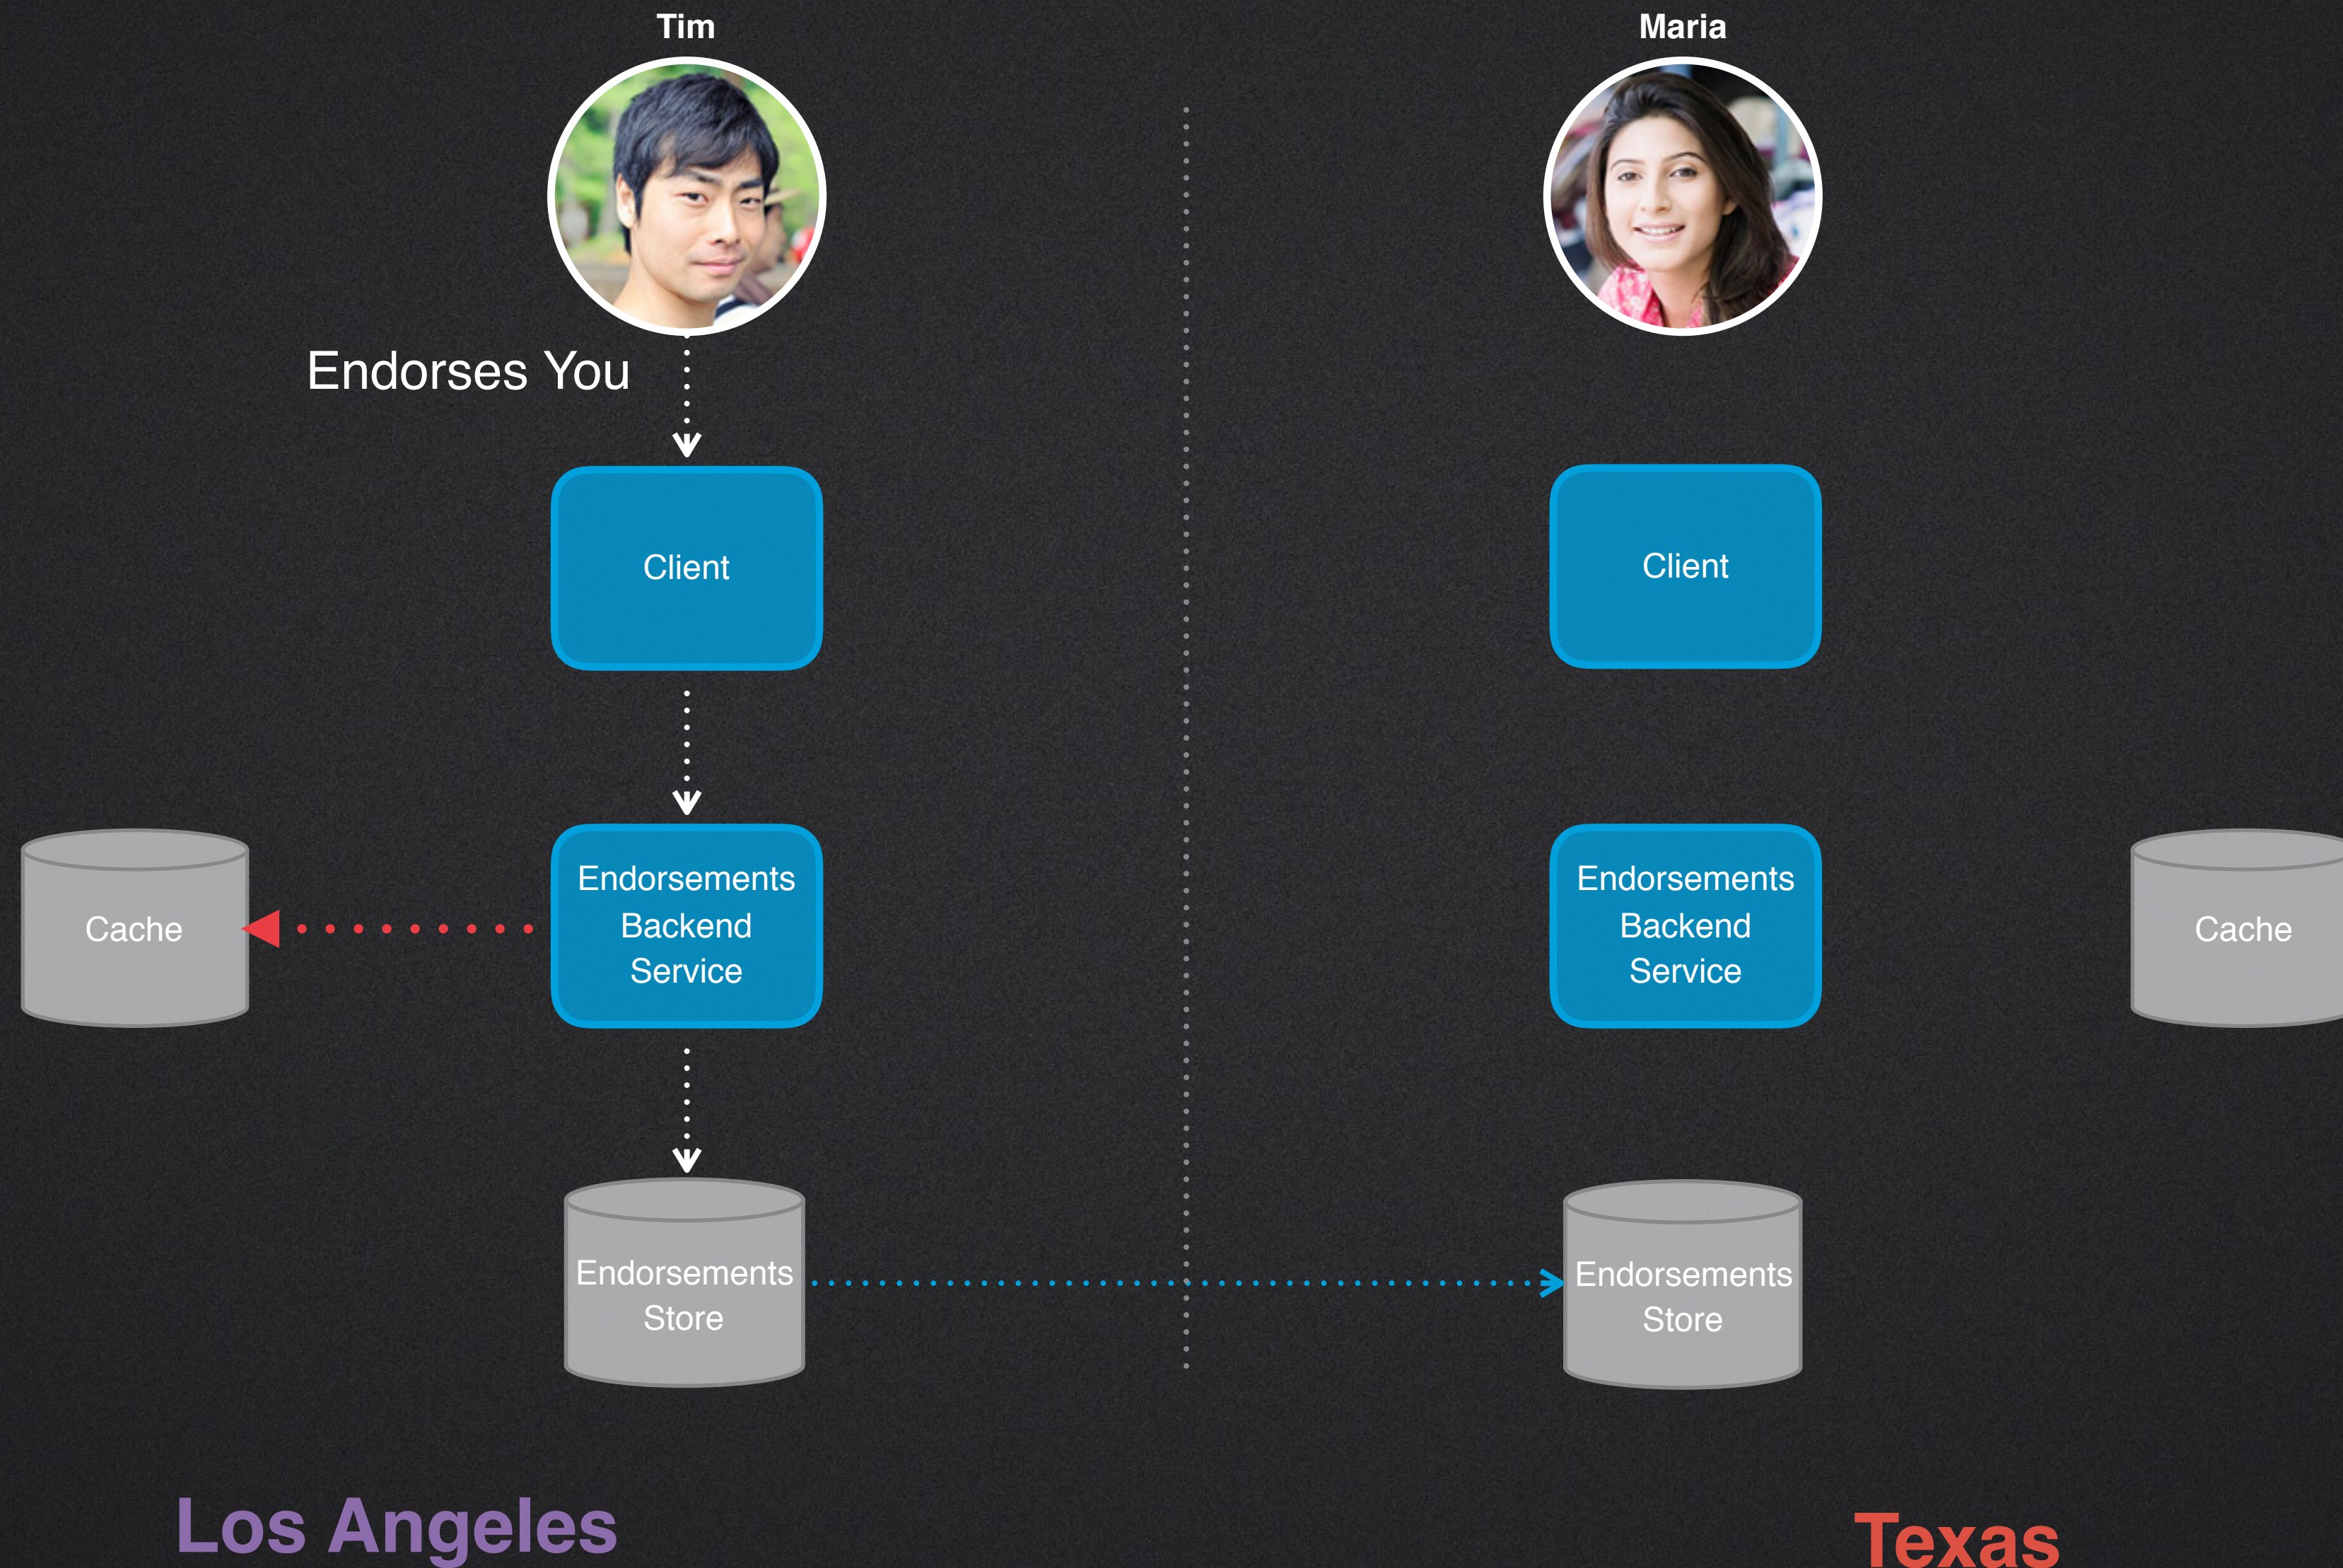


# The Problem with Cache Invalidation



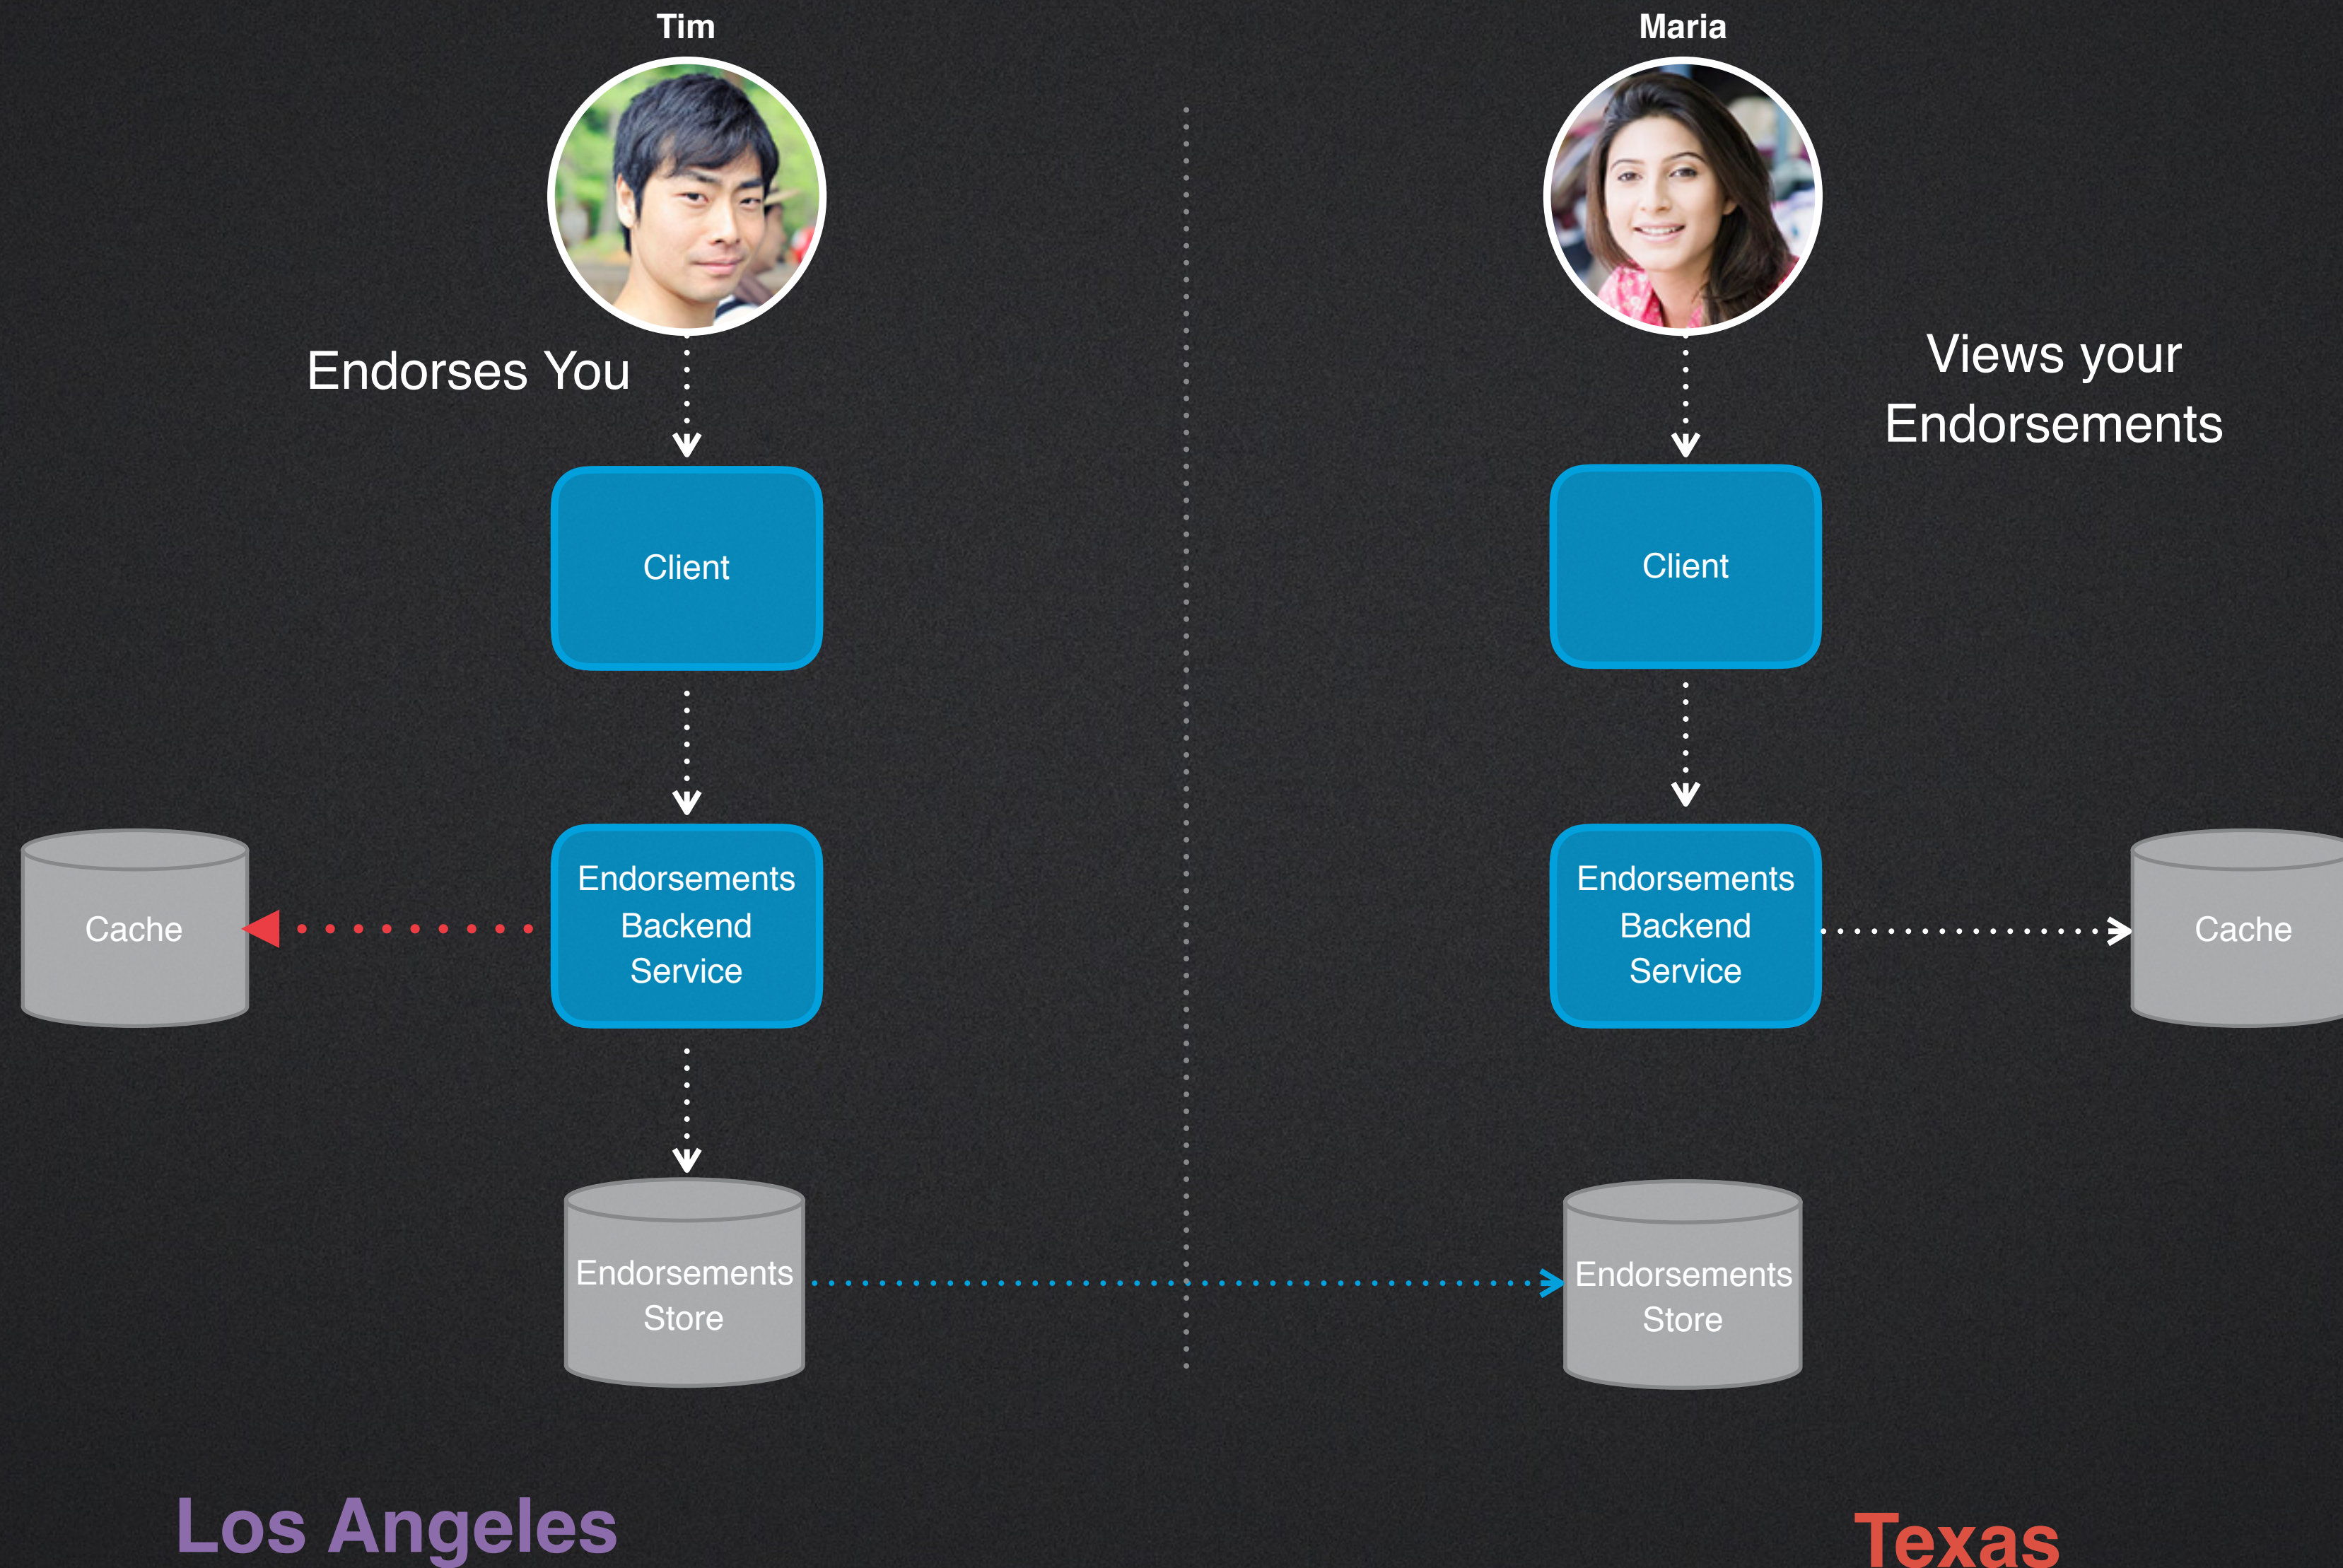


# The Problem with Cache Invalidation



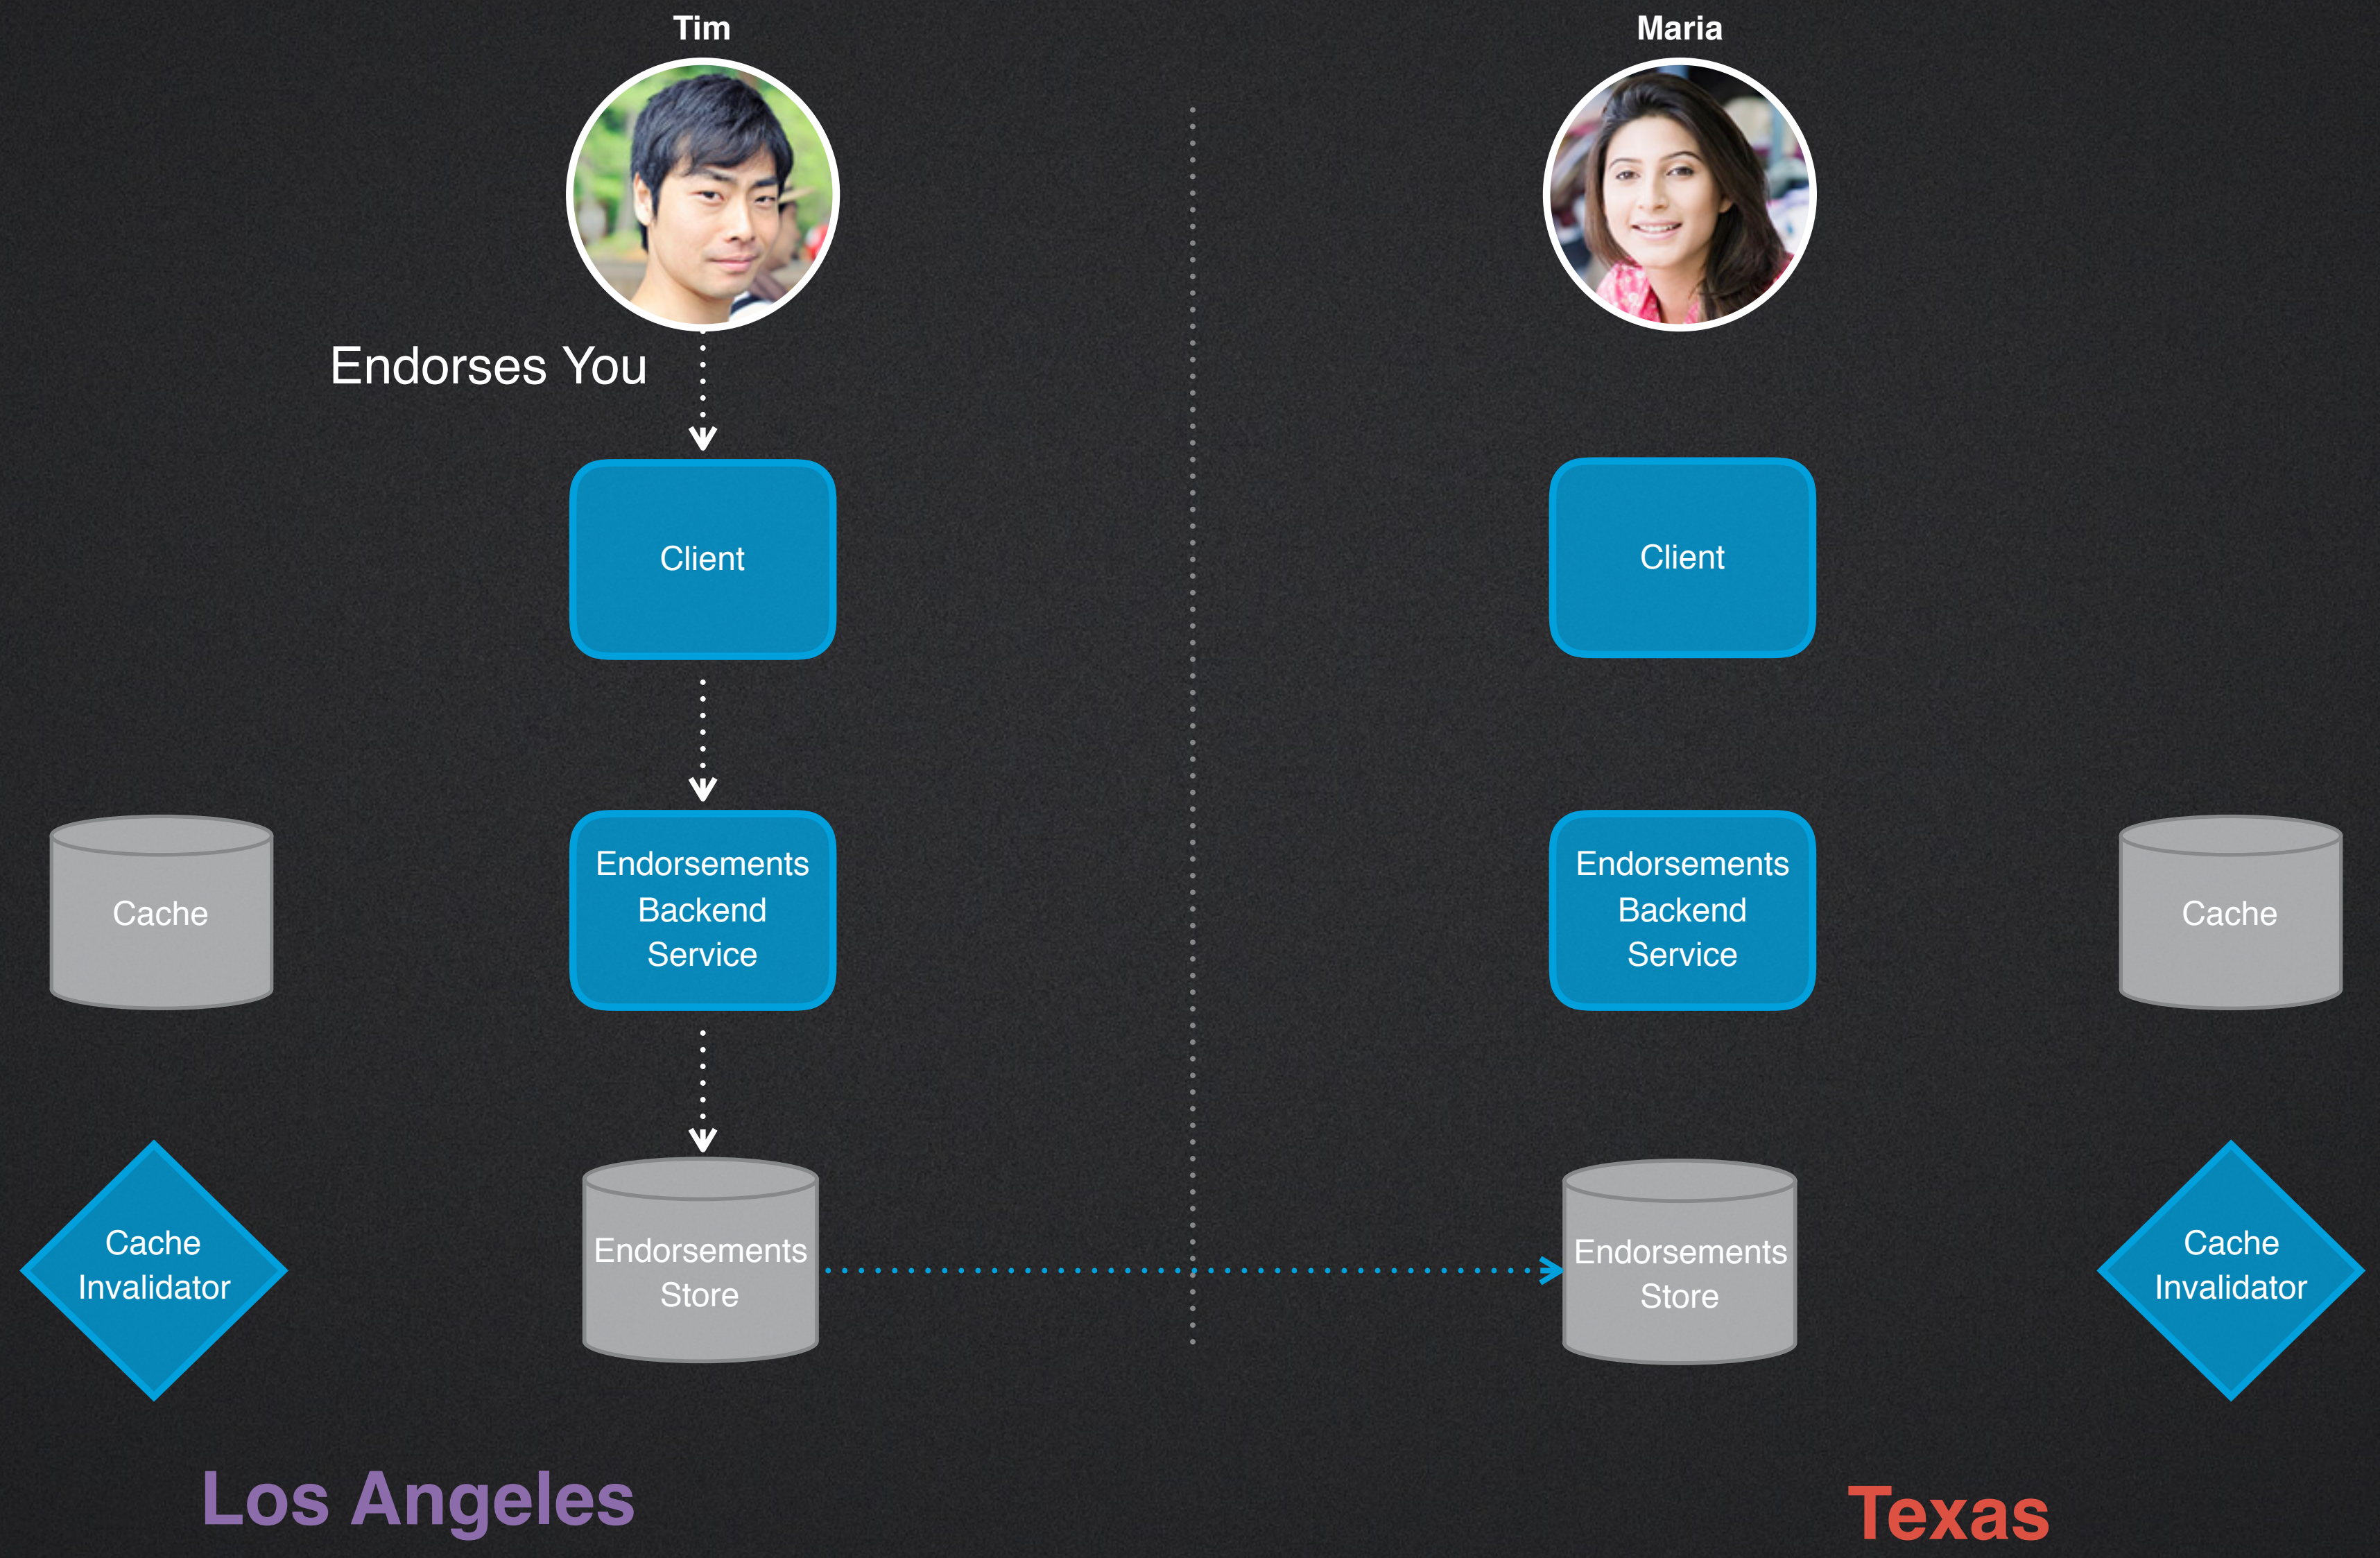


# The Problem with Cache Invalidation



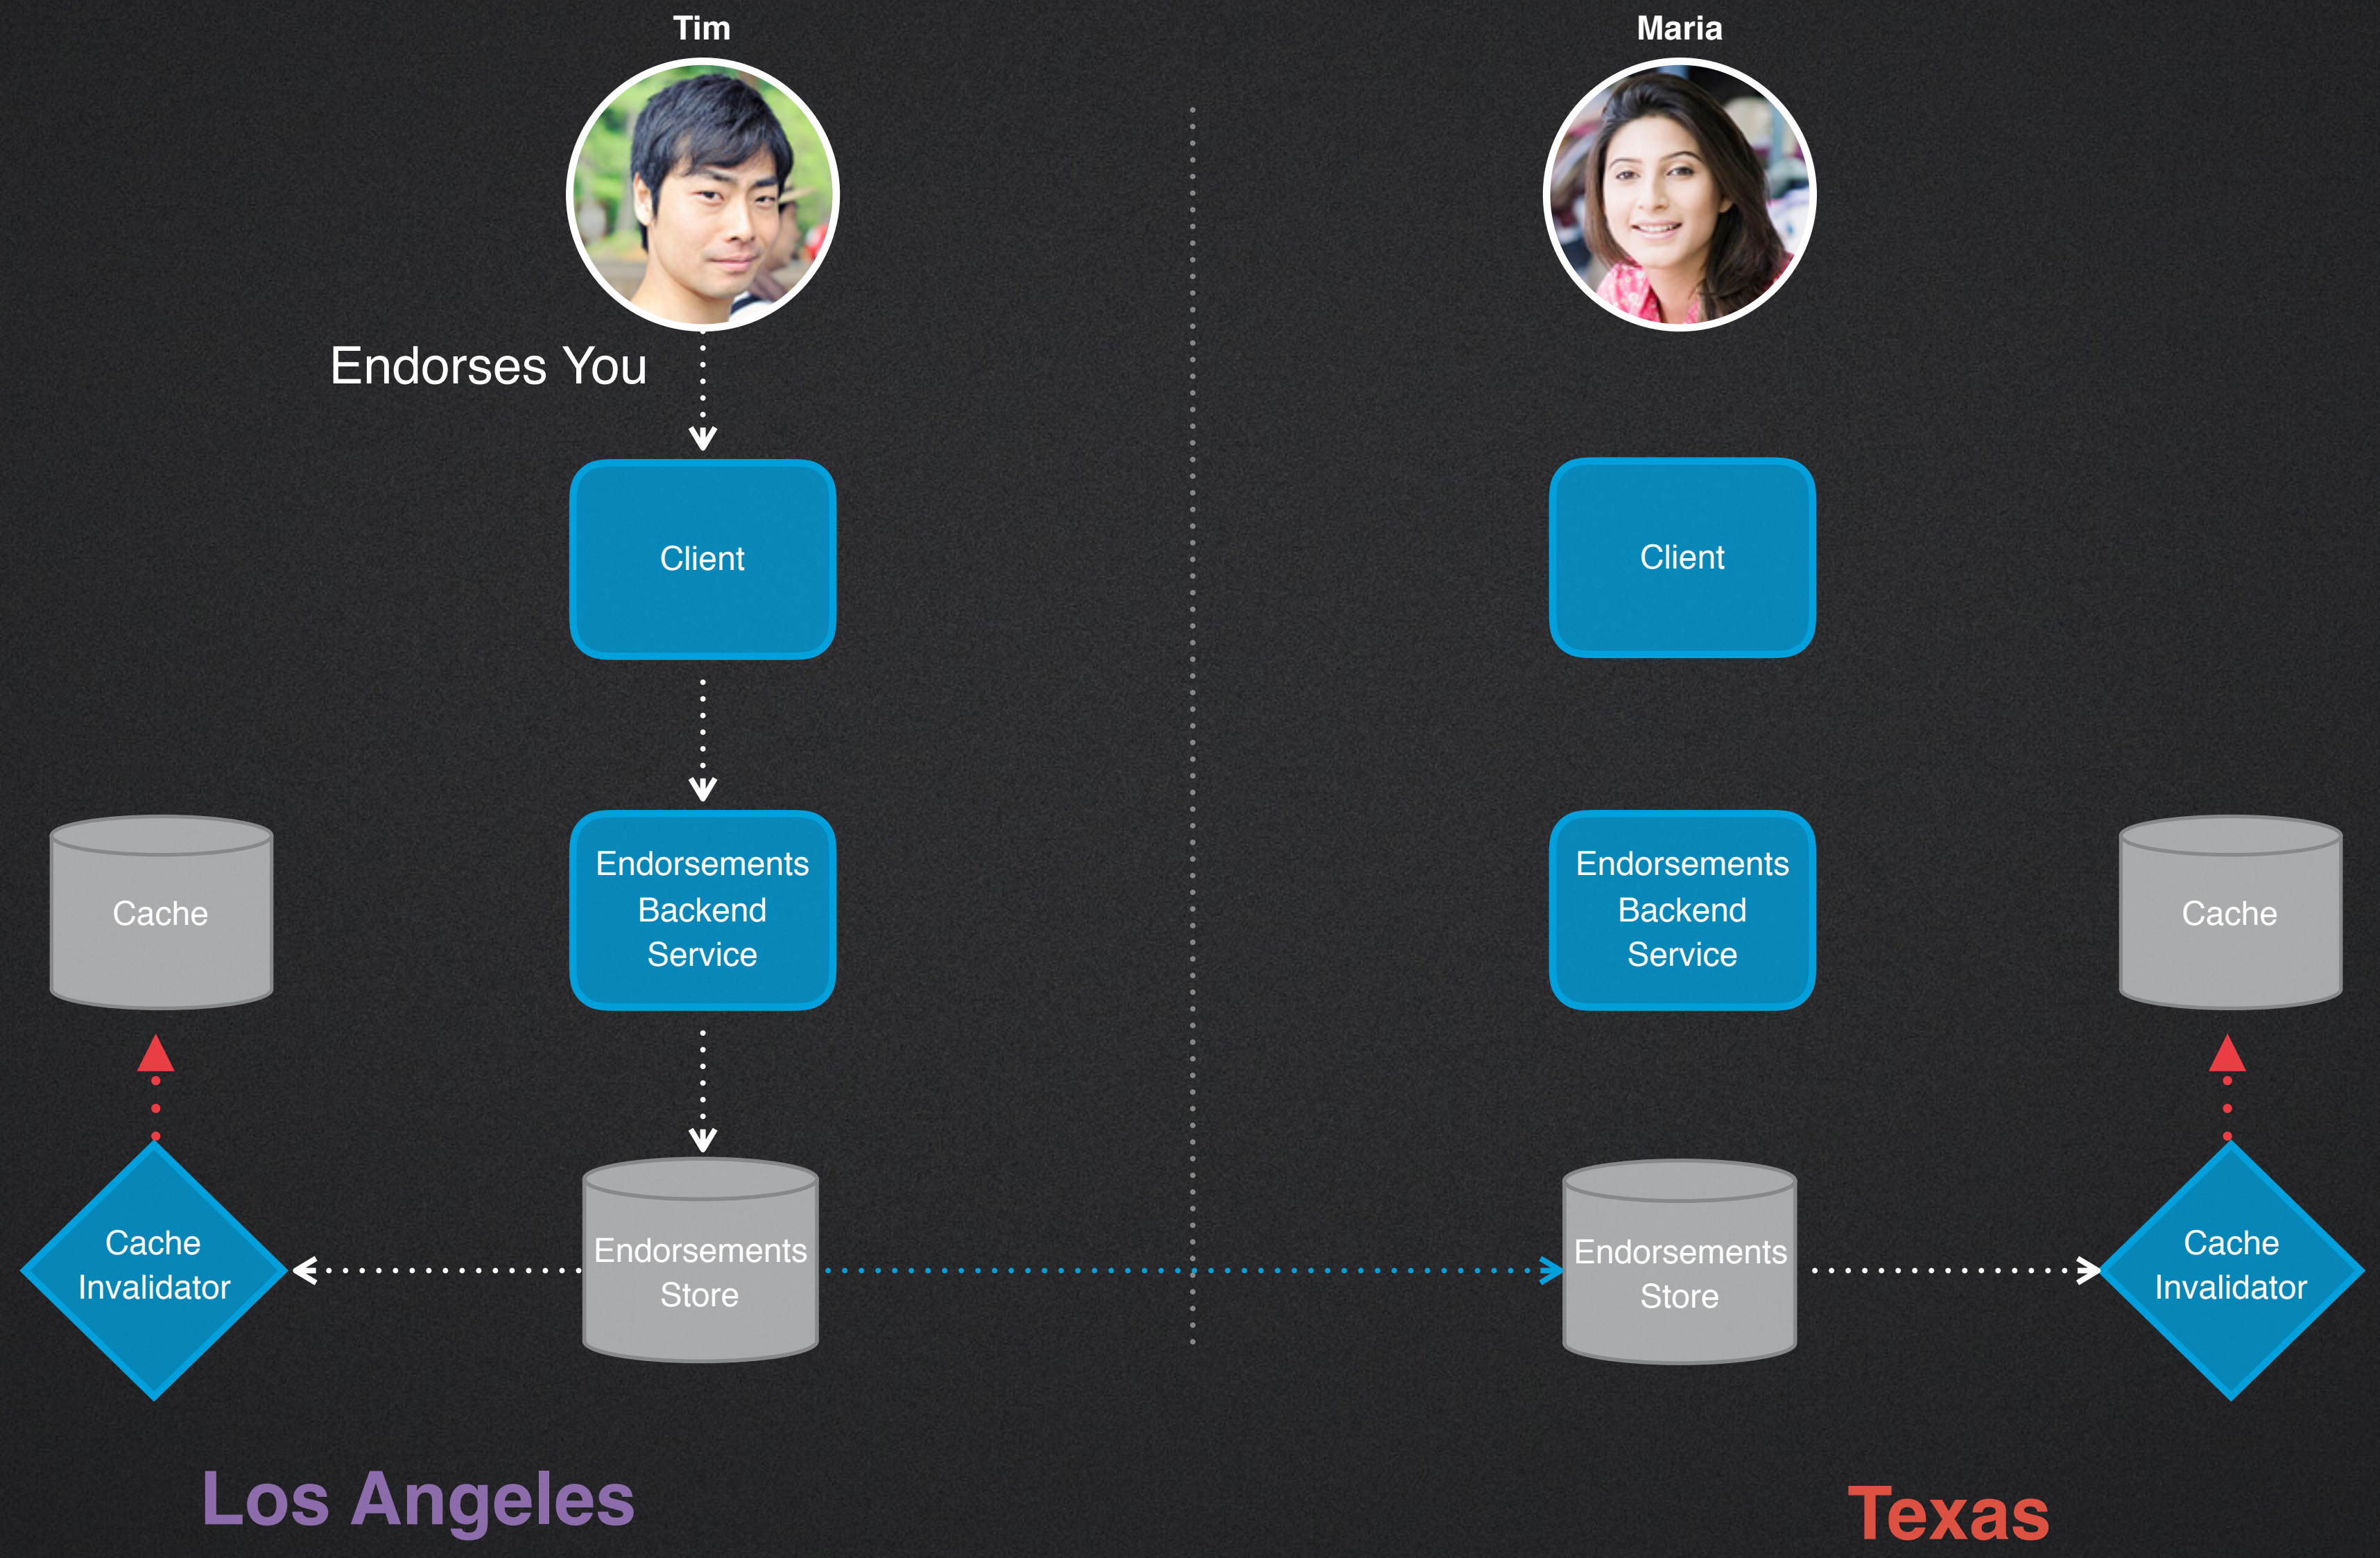


# The Solution



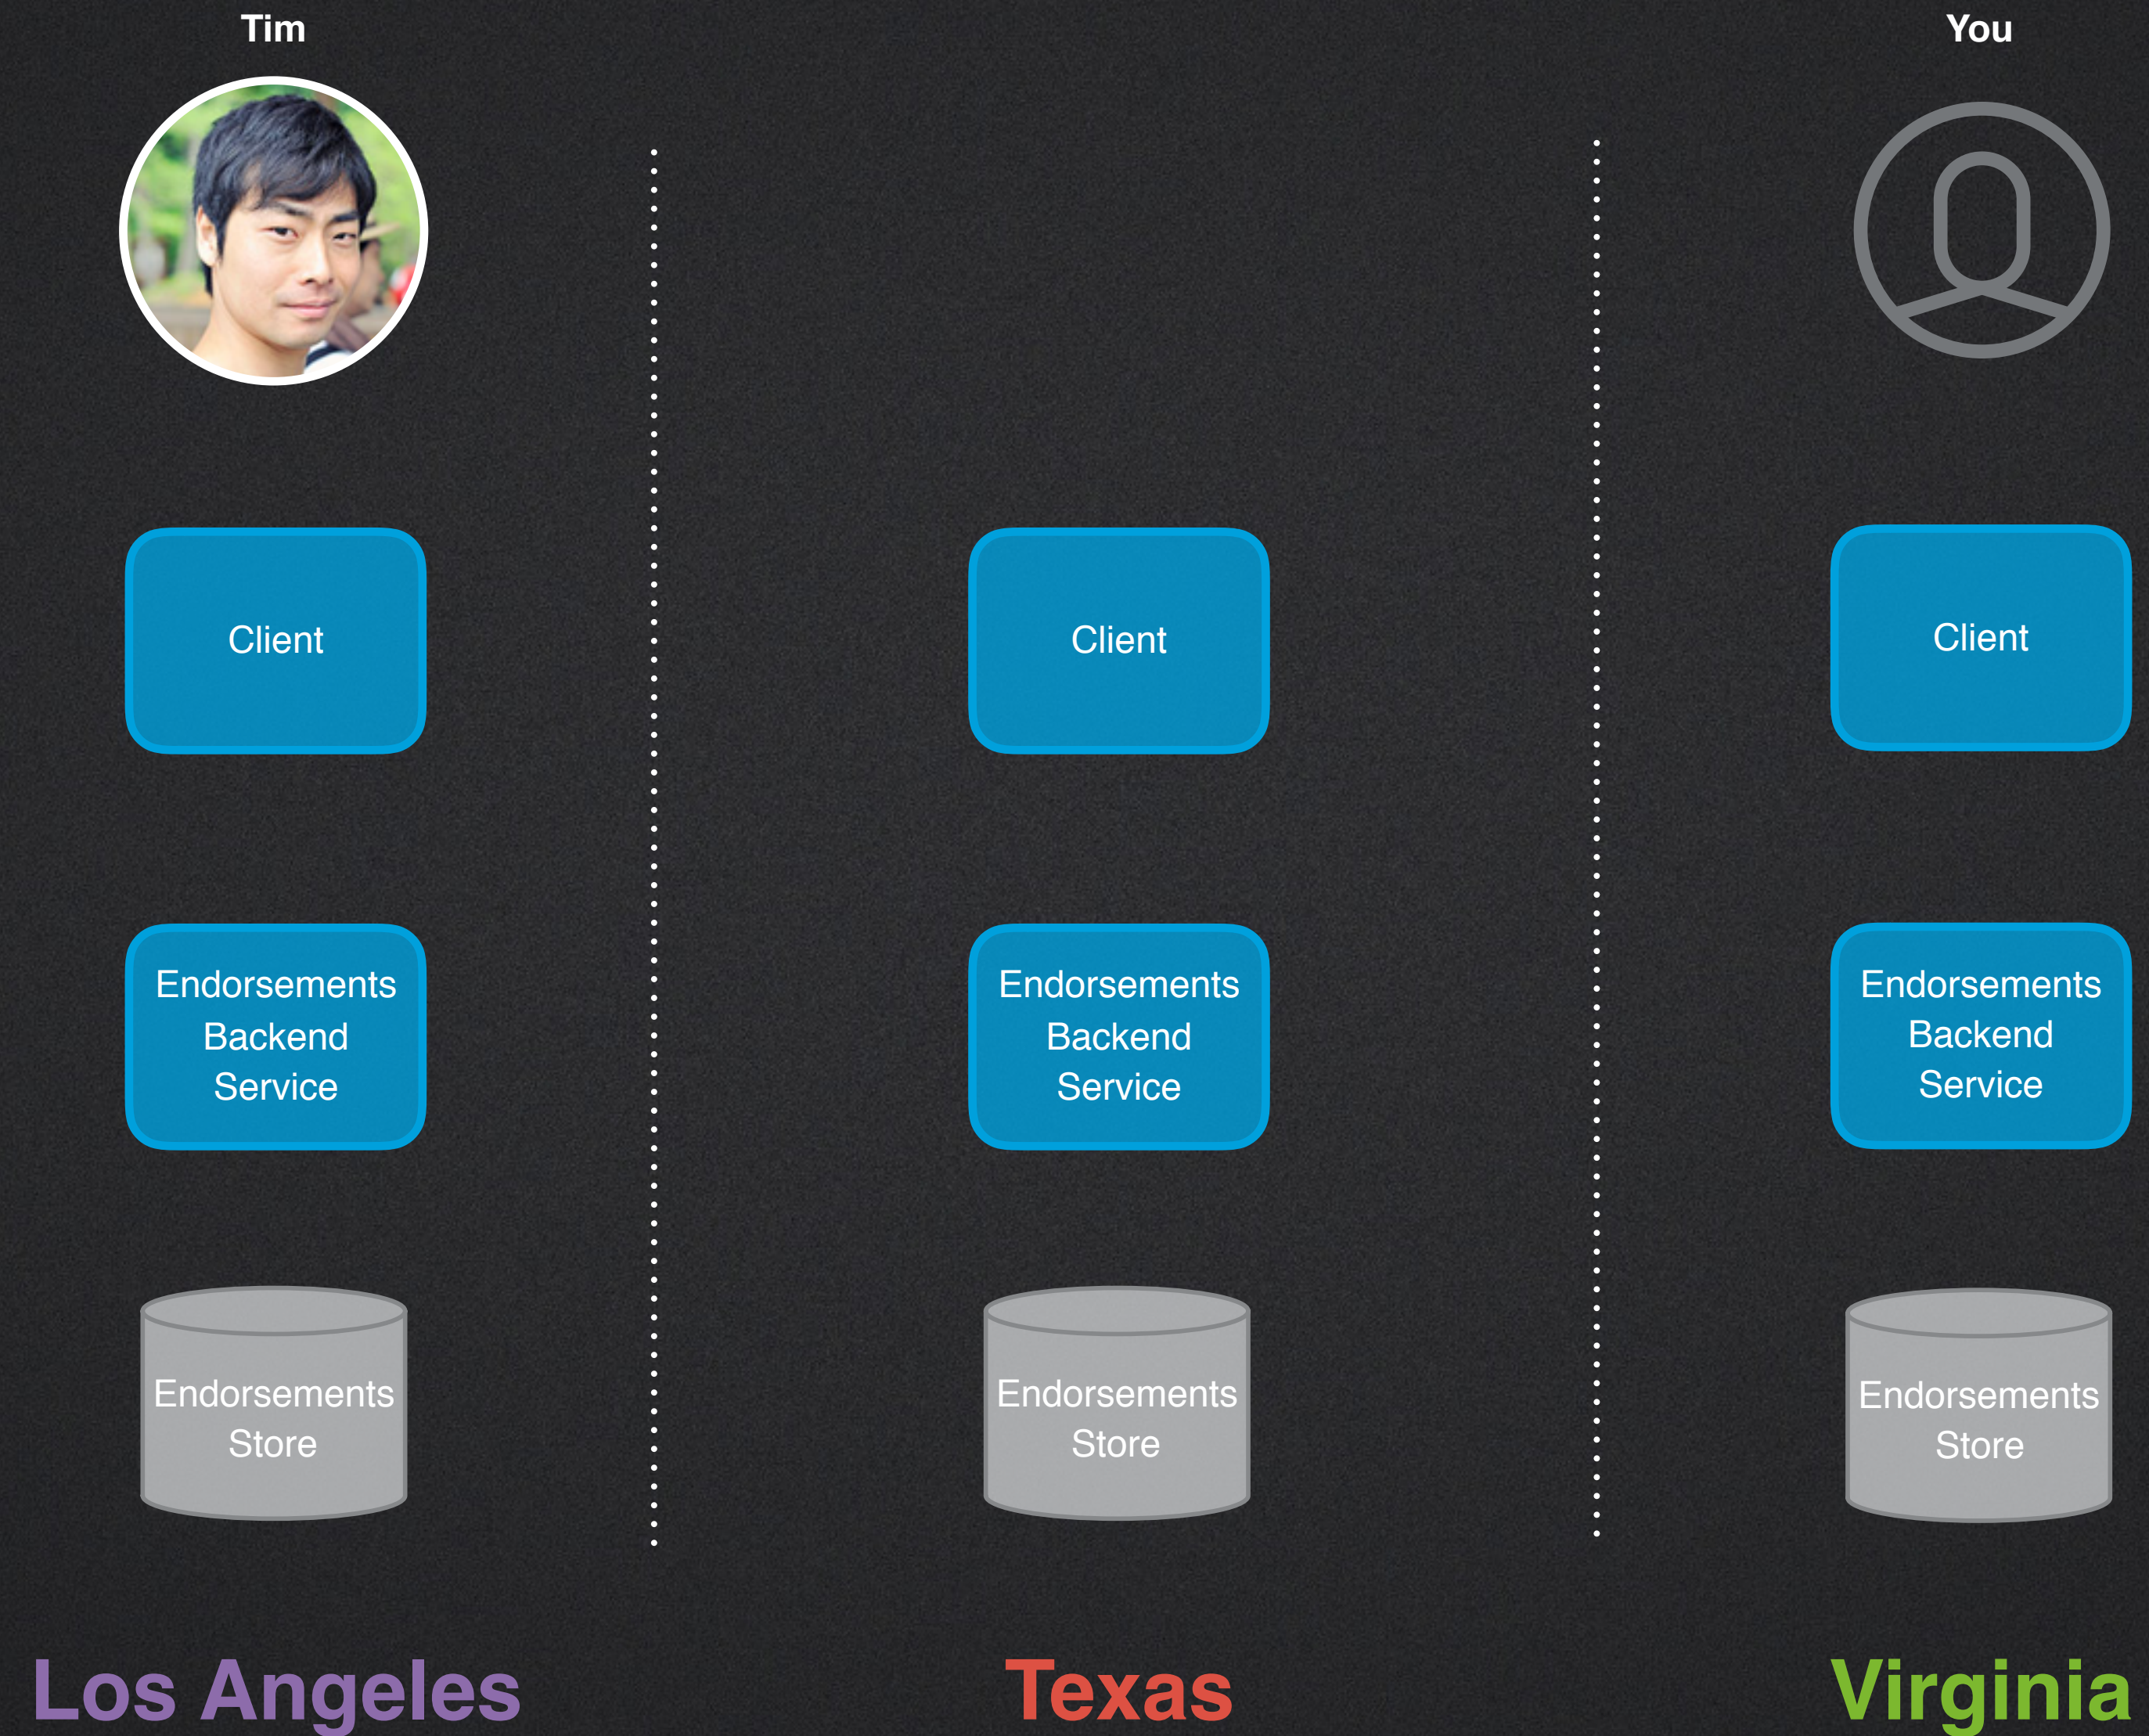


# The Solution



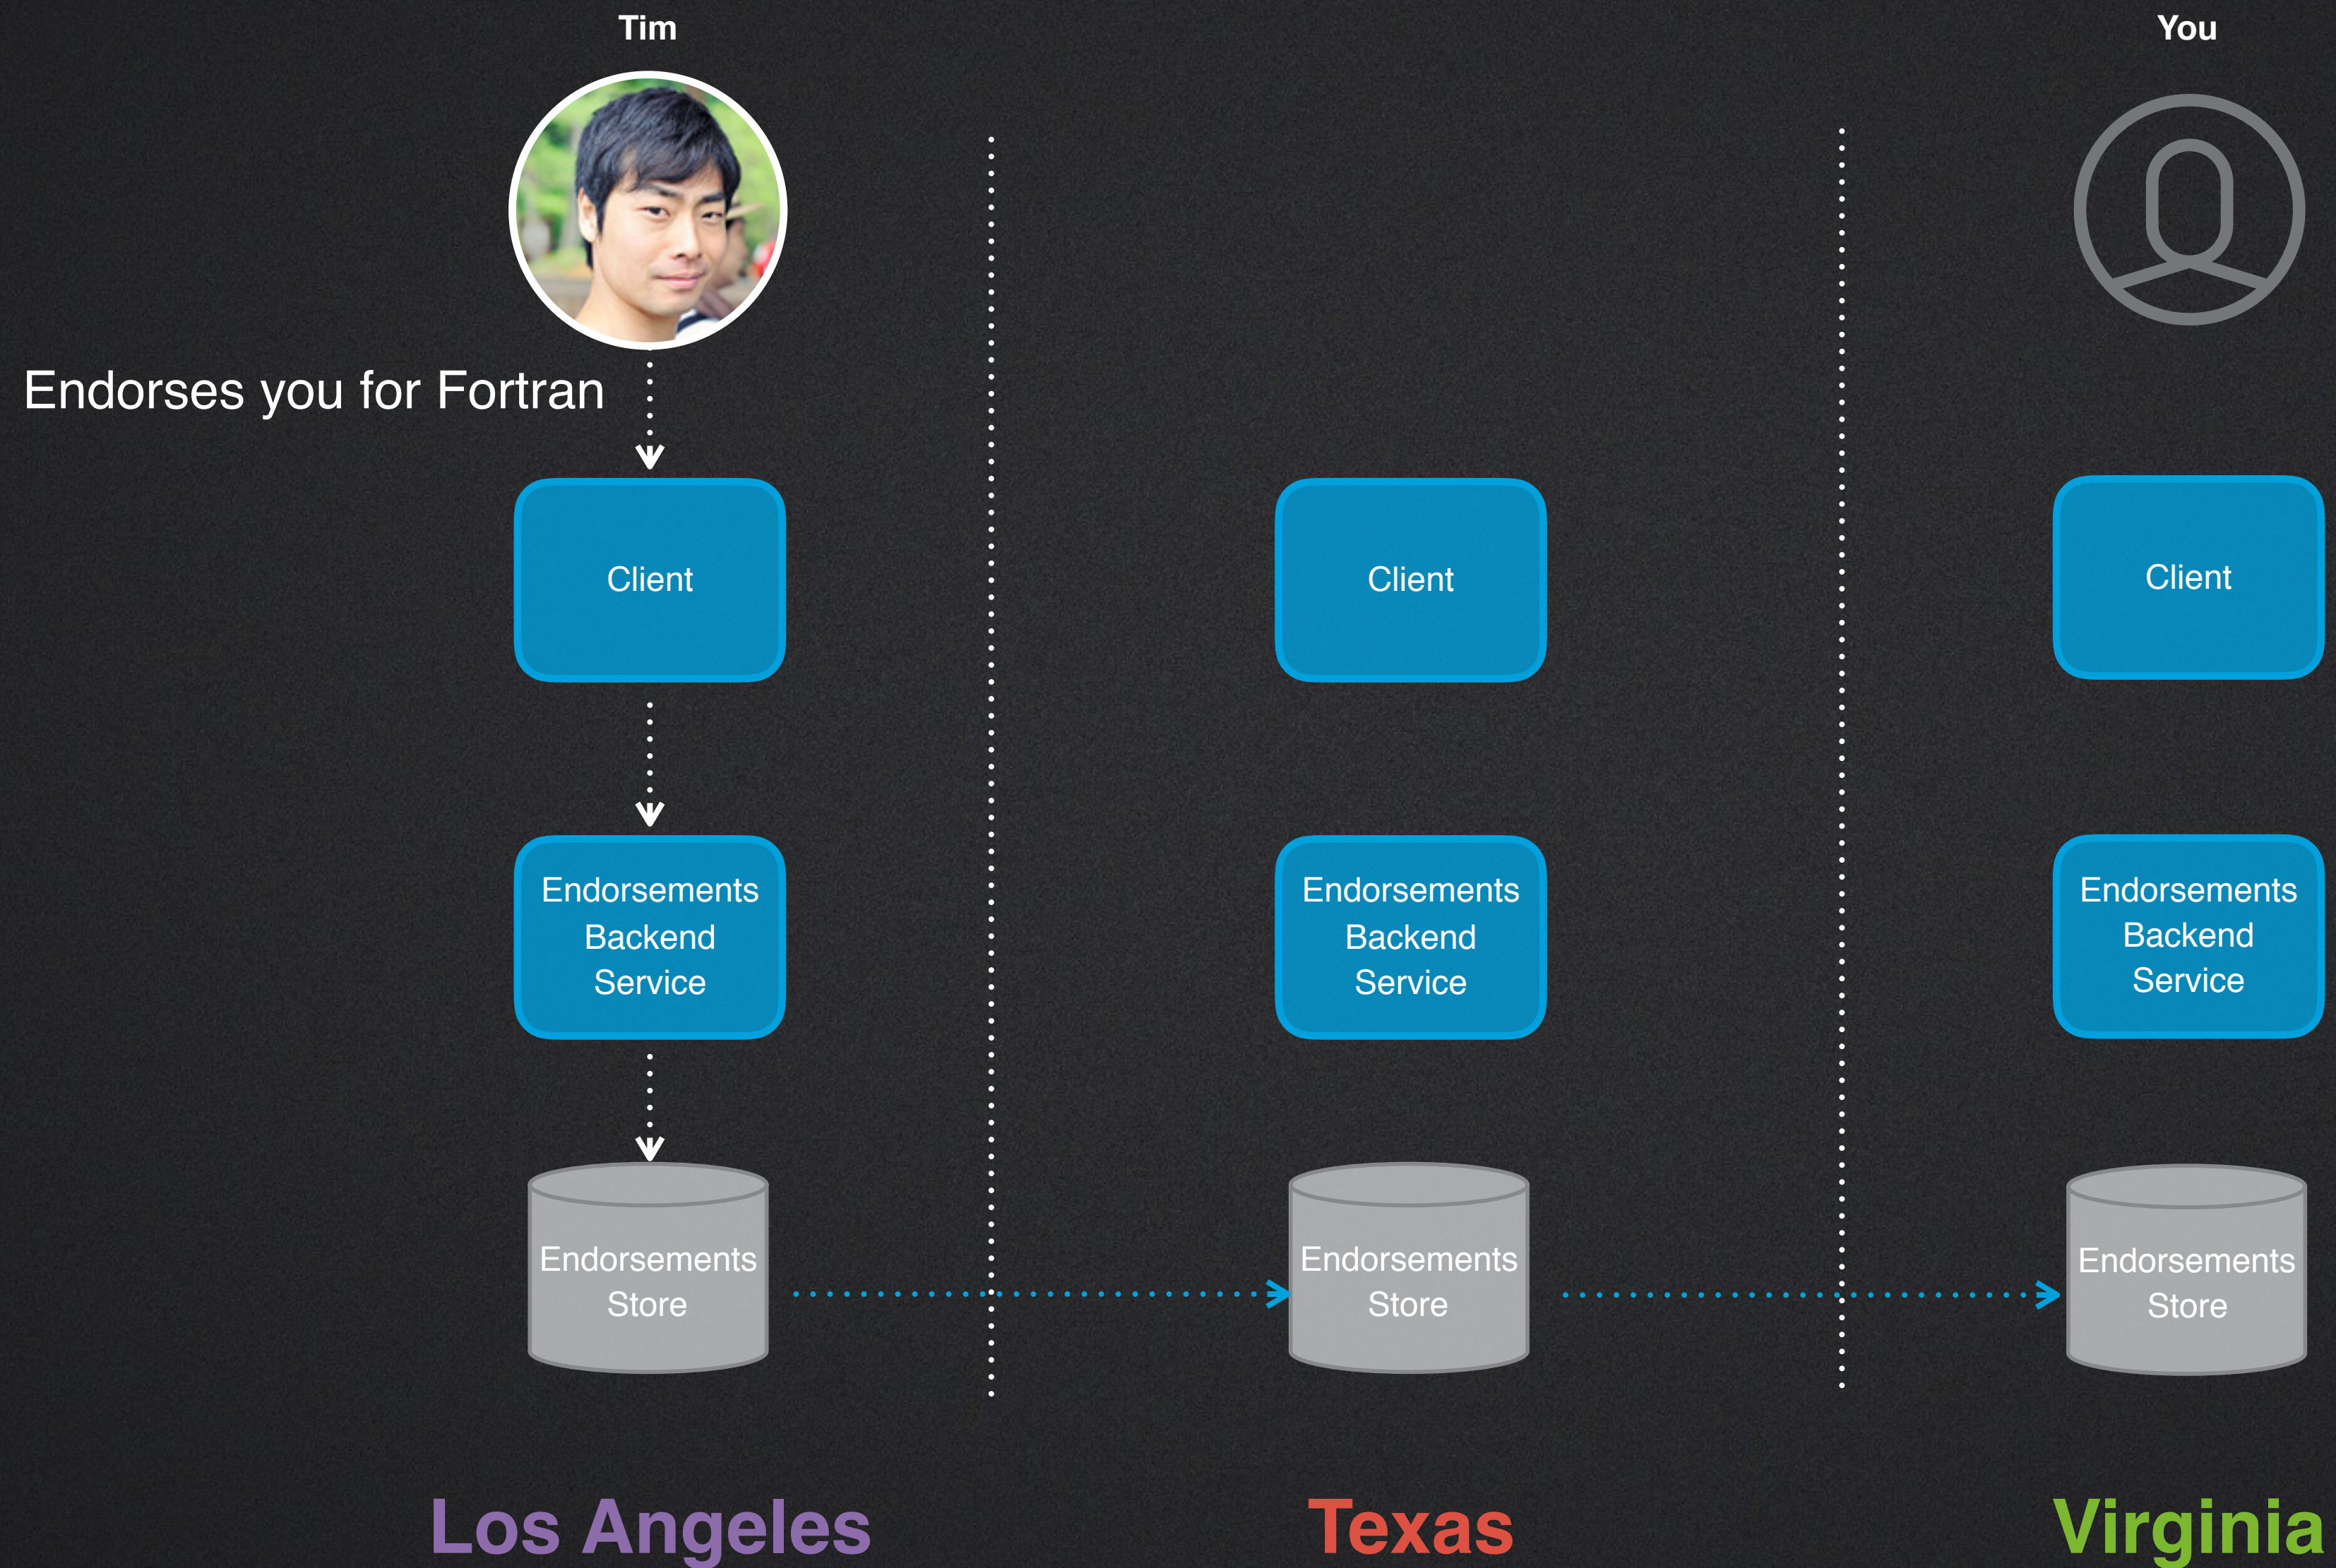


# The Problem with Data Replication



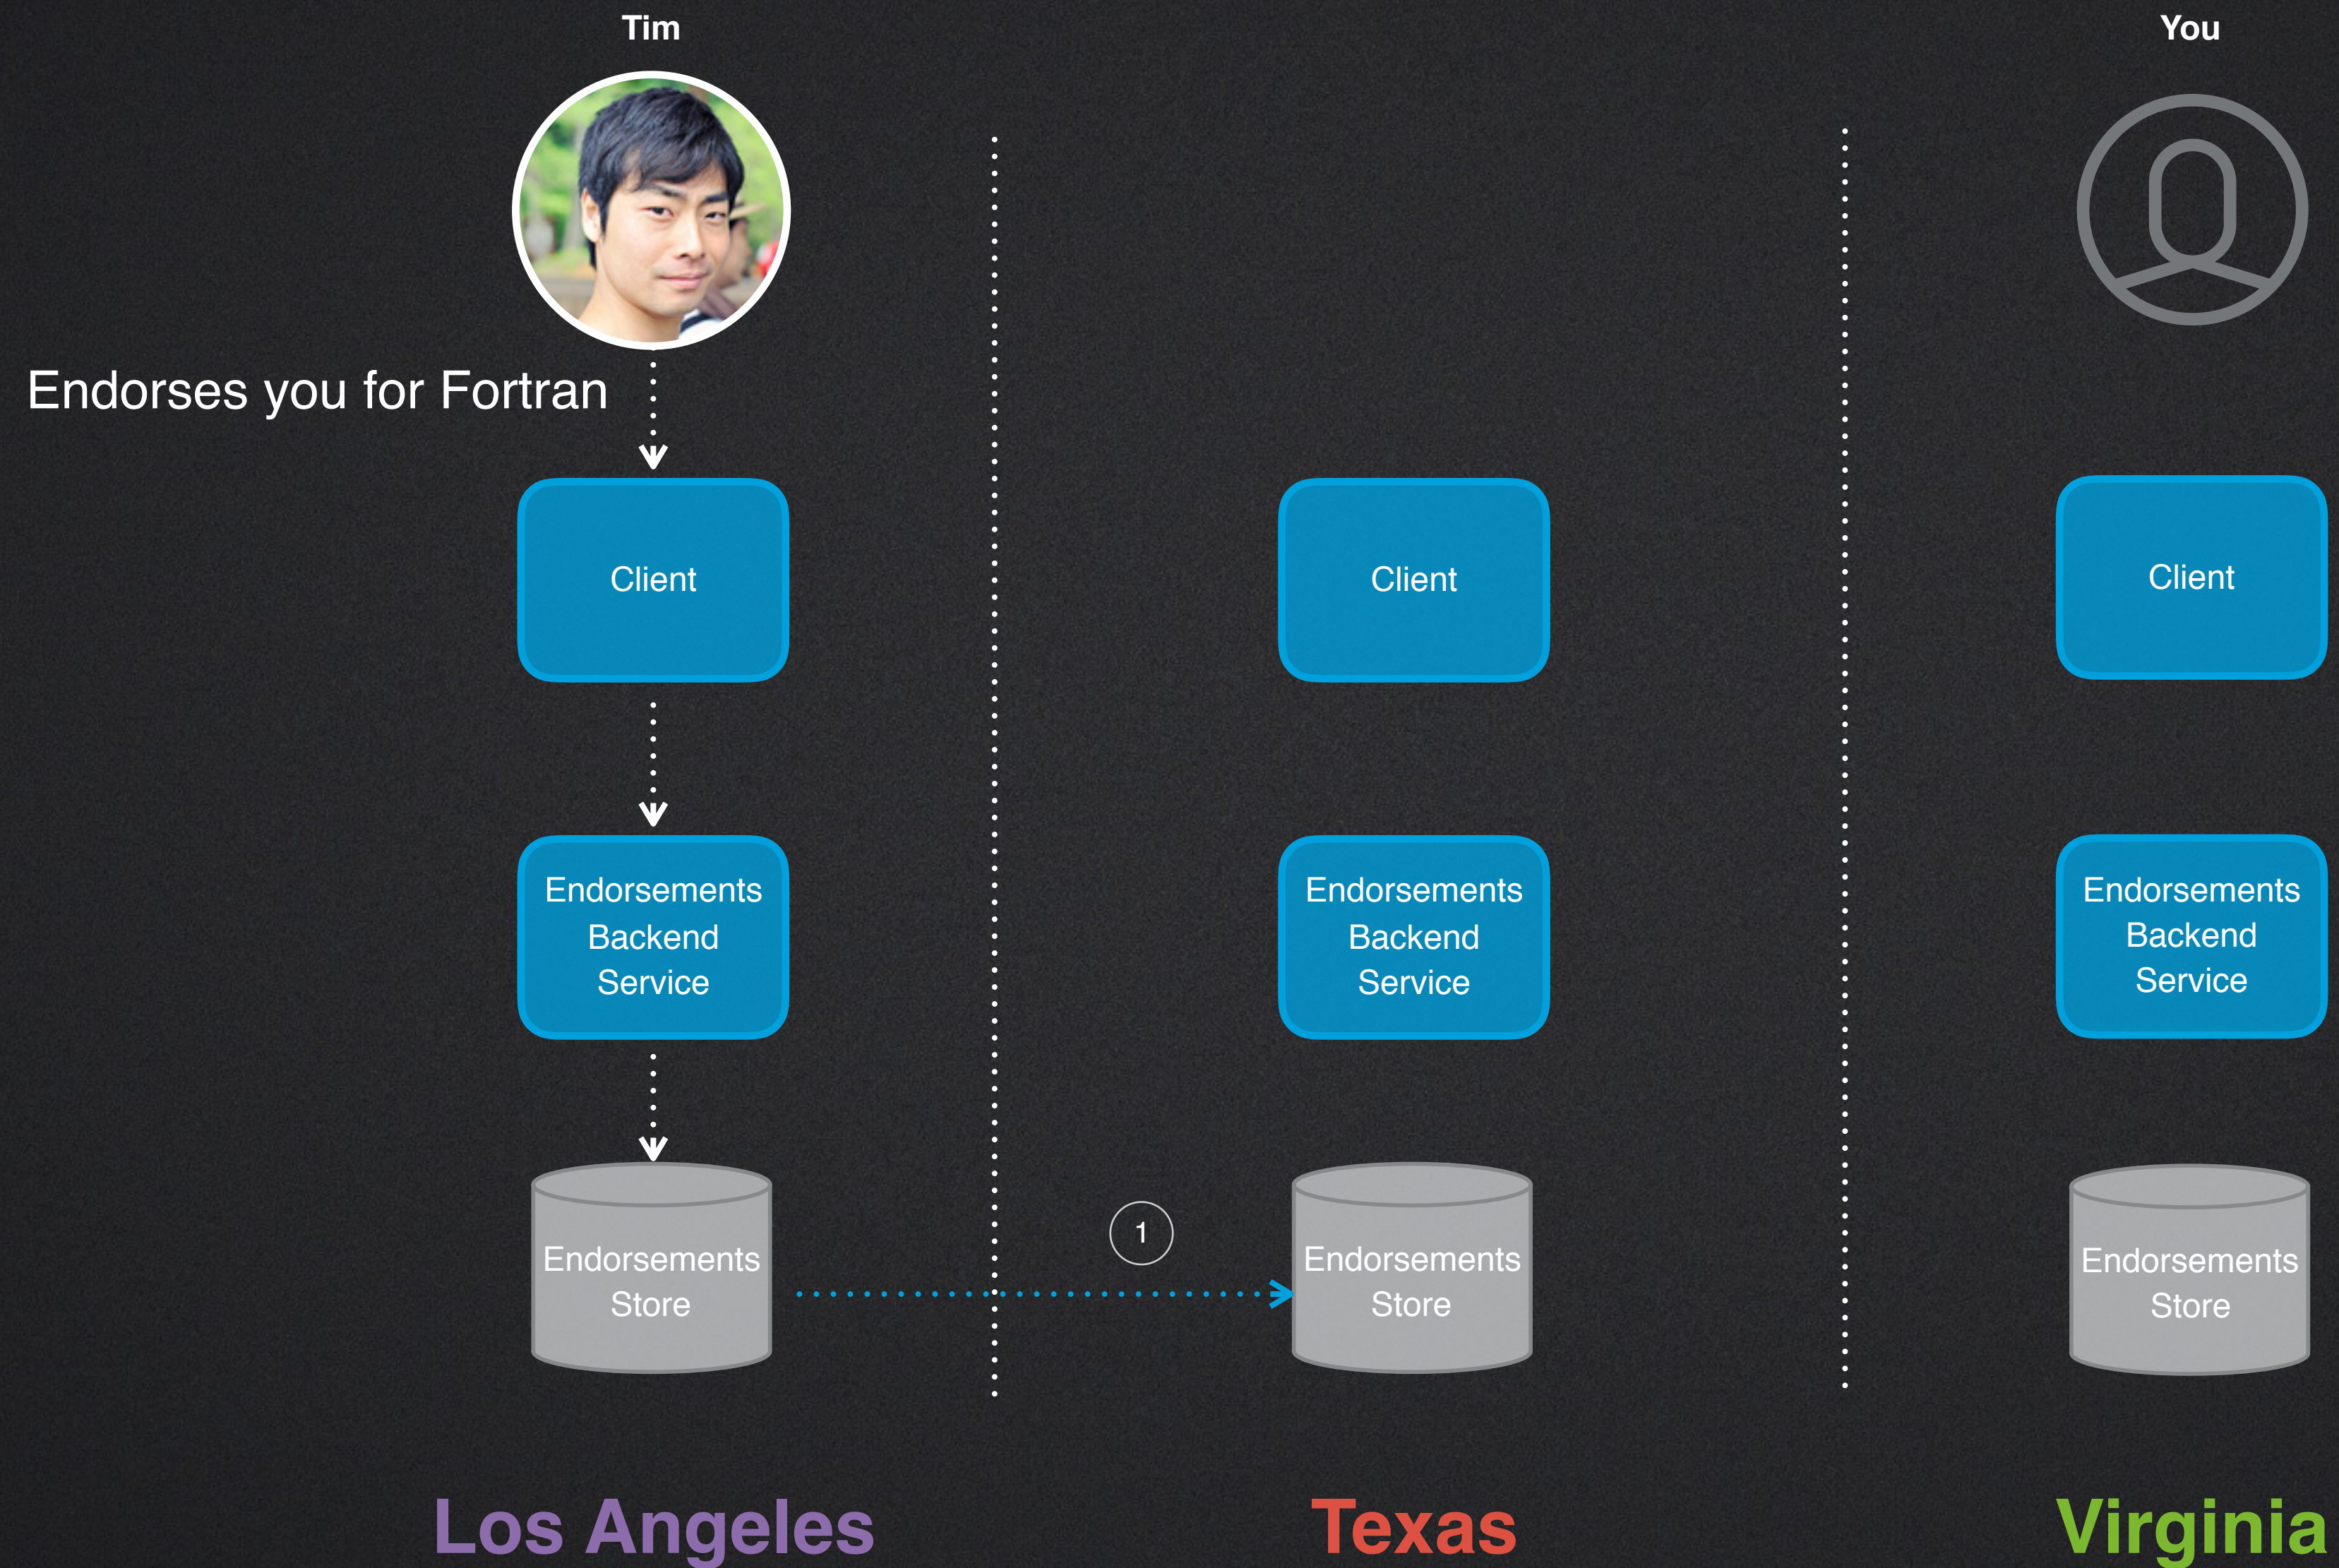


# The Problem with Data Replication



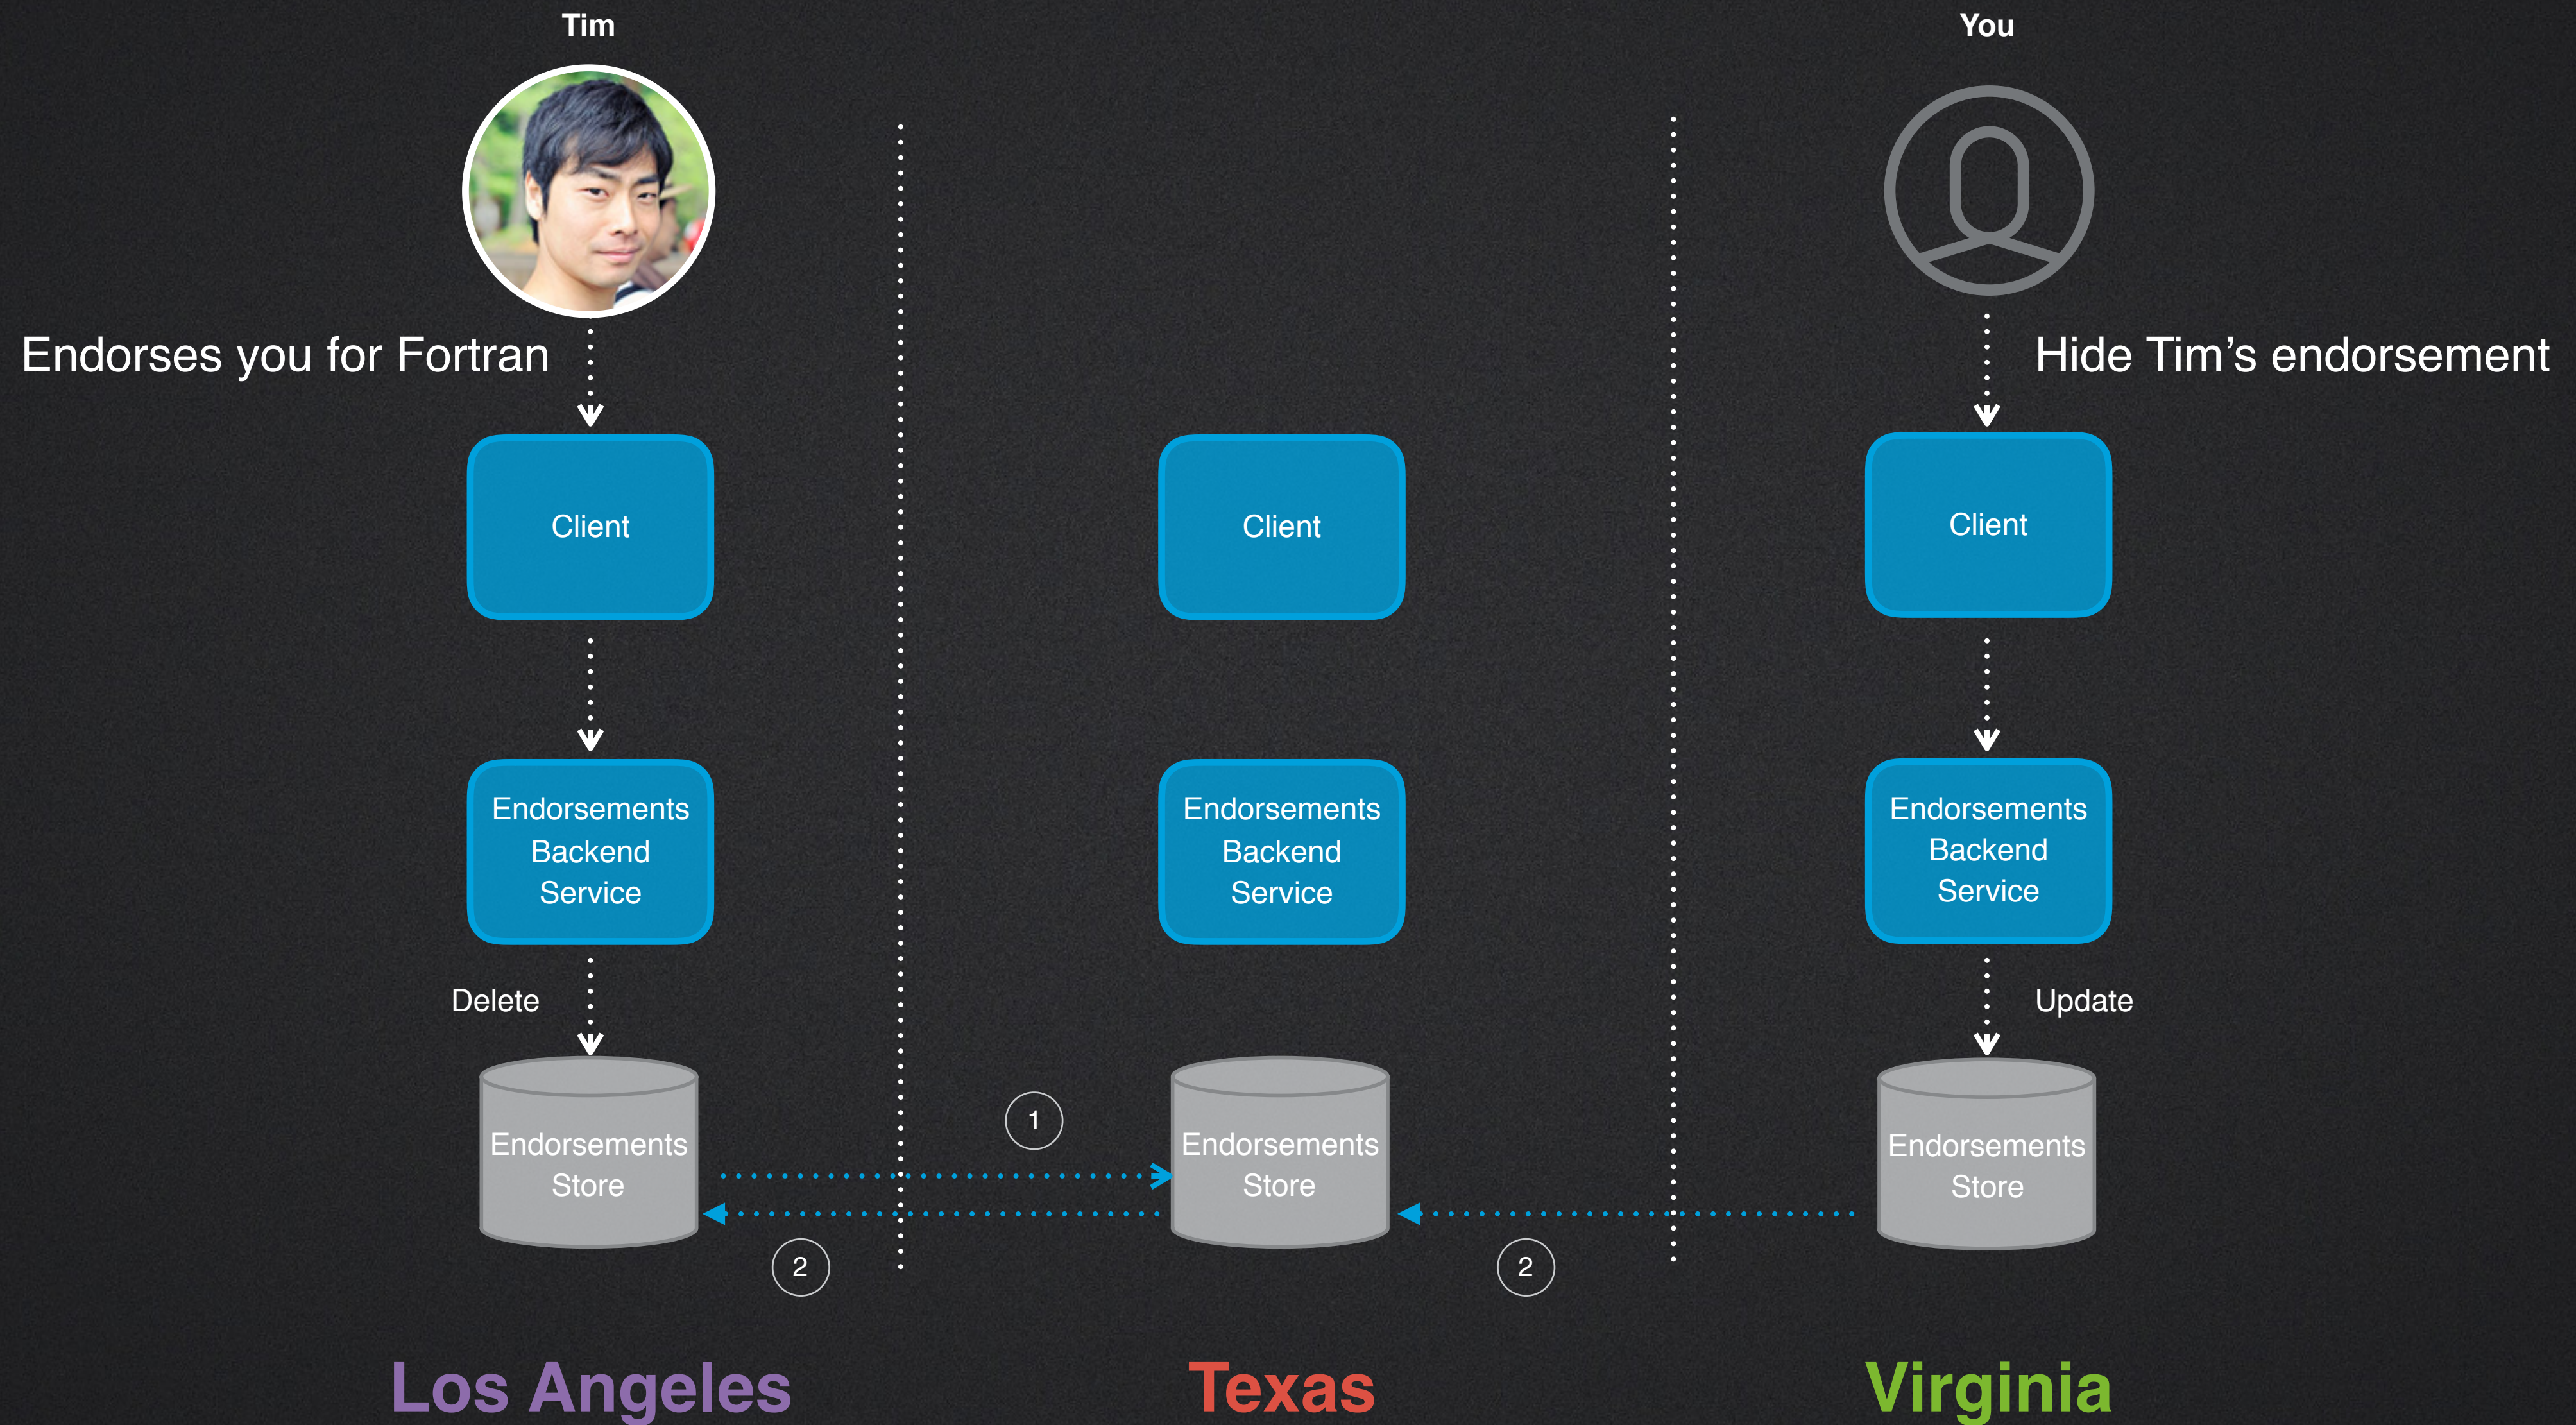


# The Problem with Data Replication





# The Problem with Data Replication





# Soft Deletes

Endorser	Endorsee	Skill	Is Visible	Is Deleted
Tim	You	Fortran	1	1



# Delete Priority Conflict Resolution Policy

Endorser	Endorsee	Skill	Is Visible	Is Deleted
Tim	You	Fortran	1	1

Los Angeles

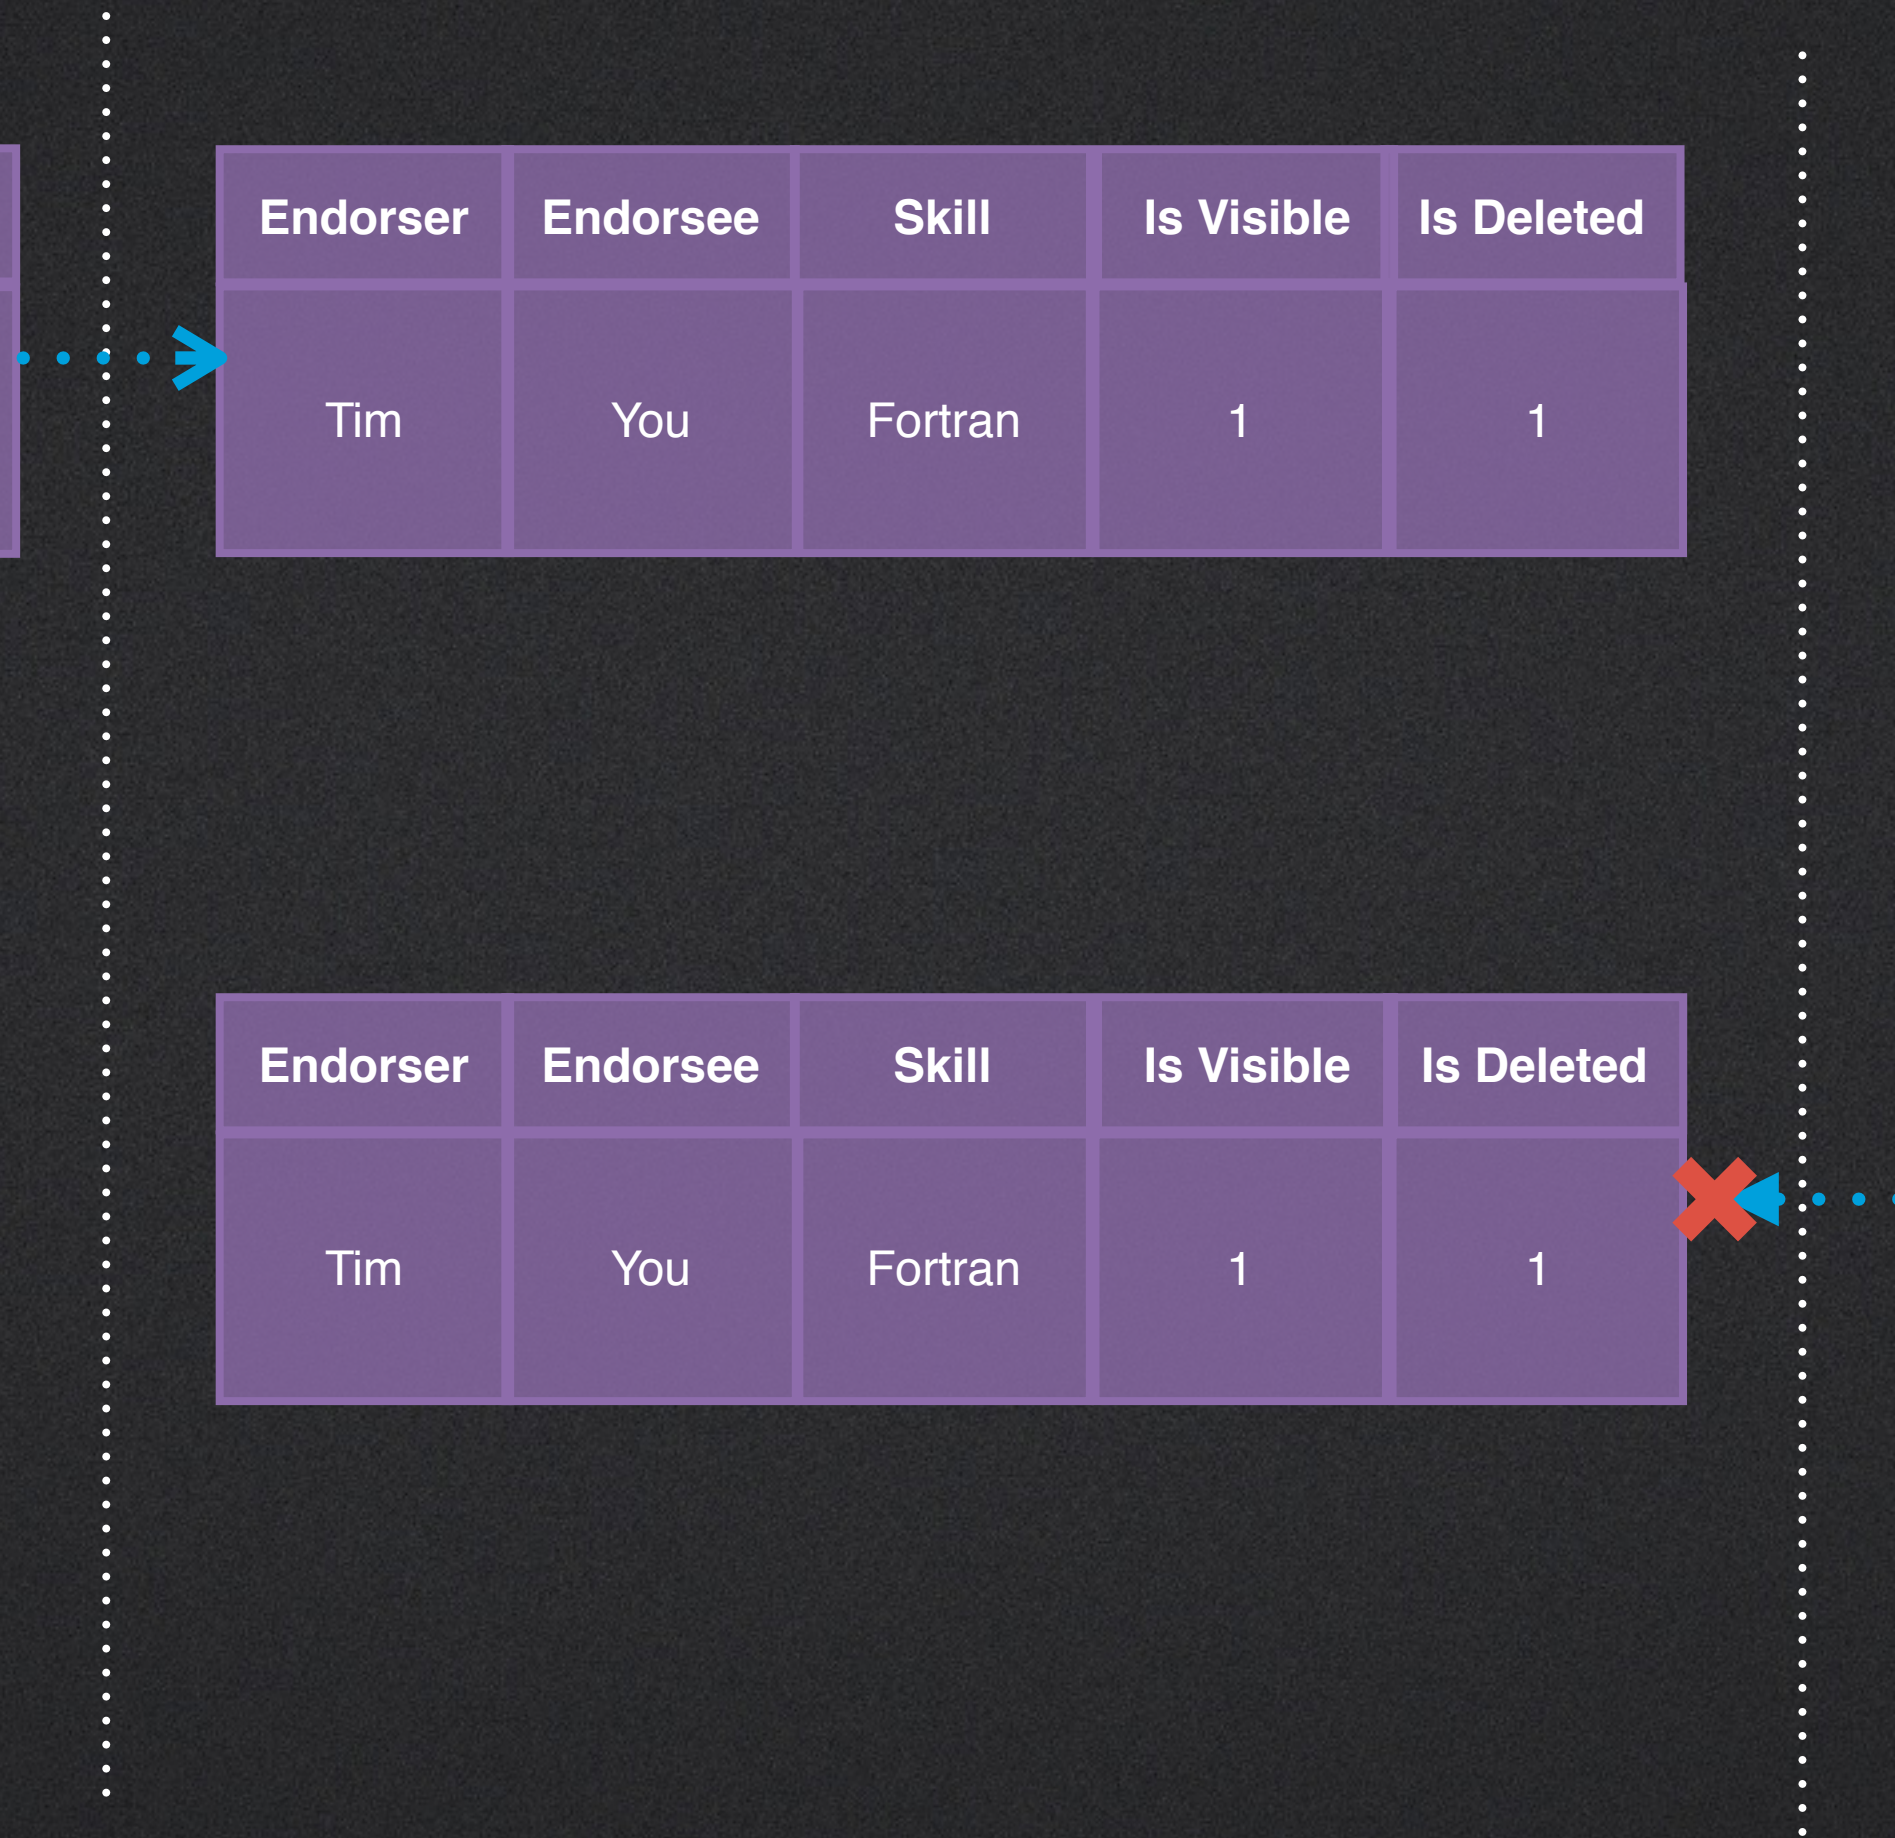
Endorser	Endorsee	Skill	Is Visible	Is Deleted
Tim	You	Fortran	1	1

Texas

Endorser	Endorsee	Skill	Is Visible	Is Deleted
Tim	You	Fortran	1	1

Endorser	Endorsee	Skill	Is Visible	Is Deleted
Tim	You	Fortran	0	0

Virginia





# Single Role Tables

## Endorsements Received Table

Endorsement ID	Skill	Is Visible	Endorsee
1	Fortran	1	Me

## Endorser Given Table

Endorsement ID	Endorser	Is Deleted
1	Tim	1



# Takeaways





# Multicolor From the Get-Go





1

# Stream Processing



2

Personal Data



3

# Cache Invalidation & Replication Conflicts











# Q&A



Erran Berger

Head of Engineering, Content Products, LinkedIn  
[www.linkedin.com/in/erranberger](http://www.linkedin.com/in/erranberger)  
@eberger45