



Being Open:

How Facebook Got Its Edge

James Quinn
Network Engineer, facebook



facebook



1.71 Billion Users


1.1+ Billion Daily Users

2 Billion Photos daily

60 Billion Messages daily

15% of global egress
is already IPv6

Our edge connects to the world



What is Facebook?

- **dynamic**

- newsfeed
- likes
- status updates

- **static**

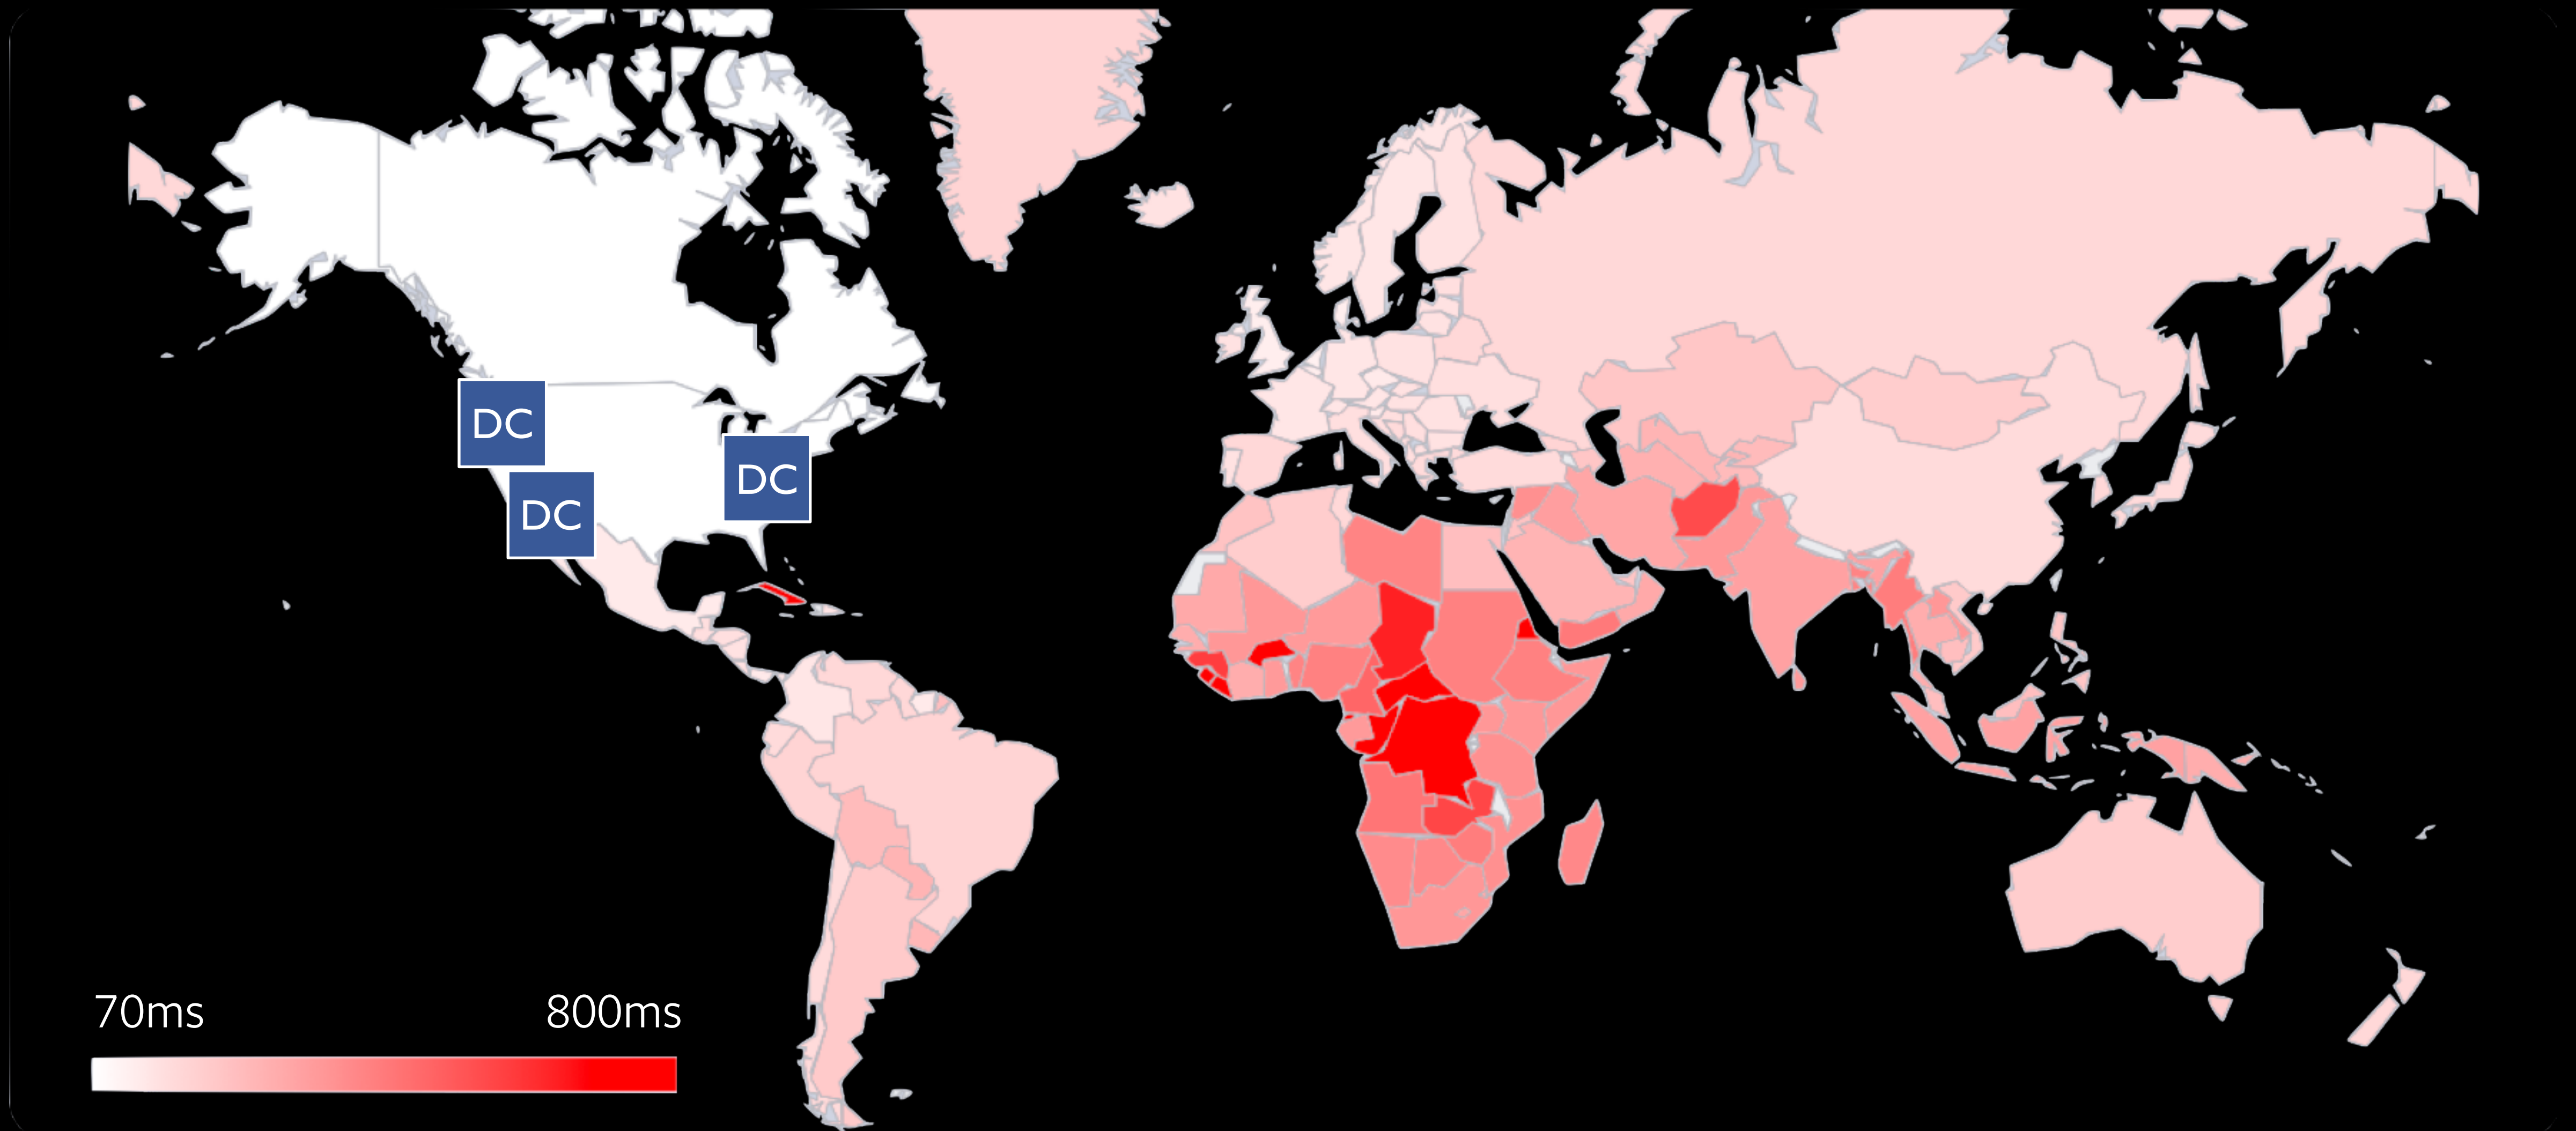
- images
- CSS
- js






International RTT

circa 11/2011



Asia -> Oregon



TCP Connect: 150ms

DC

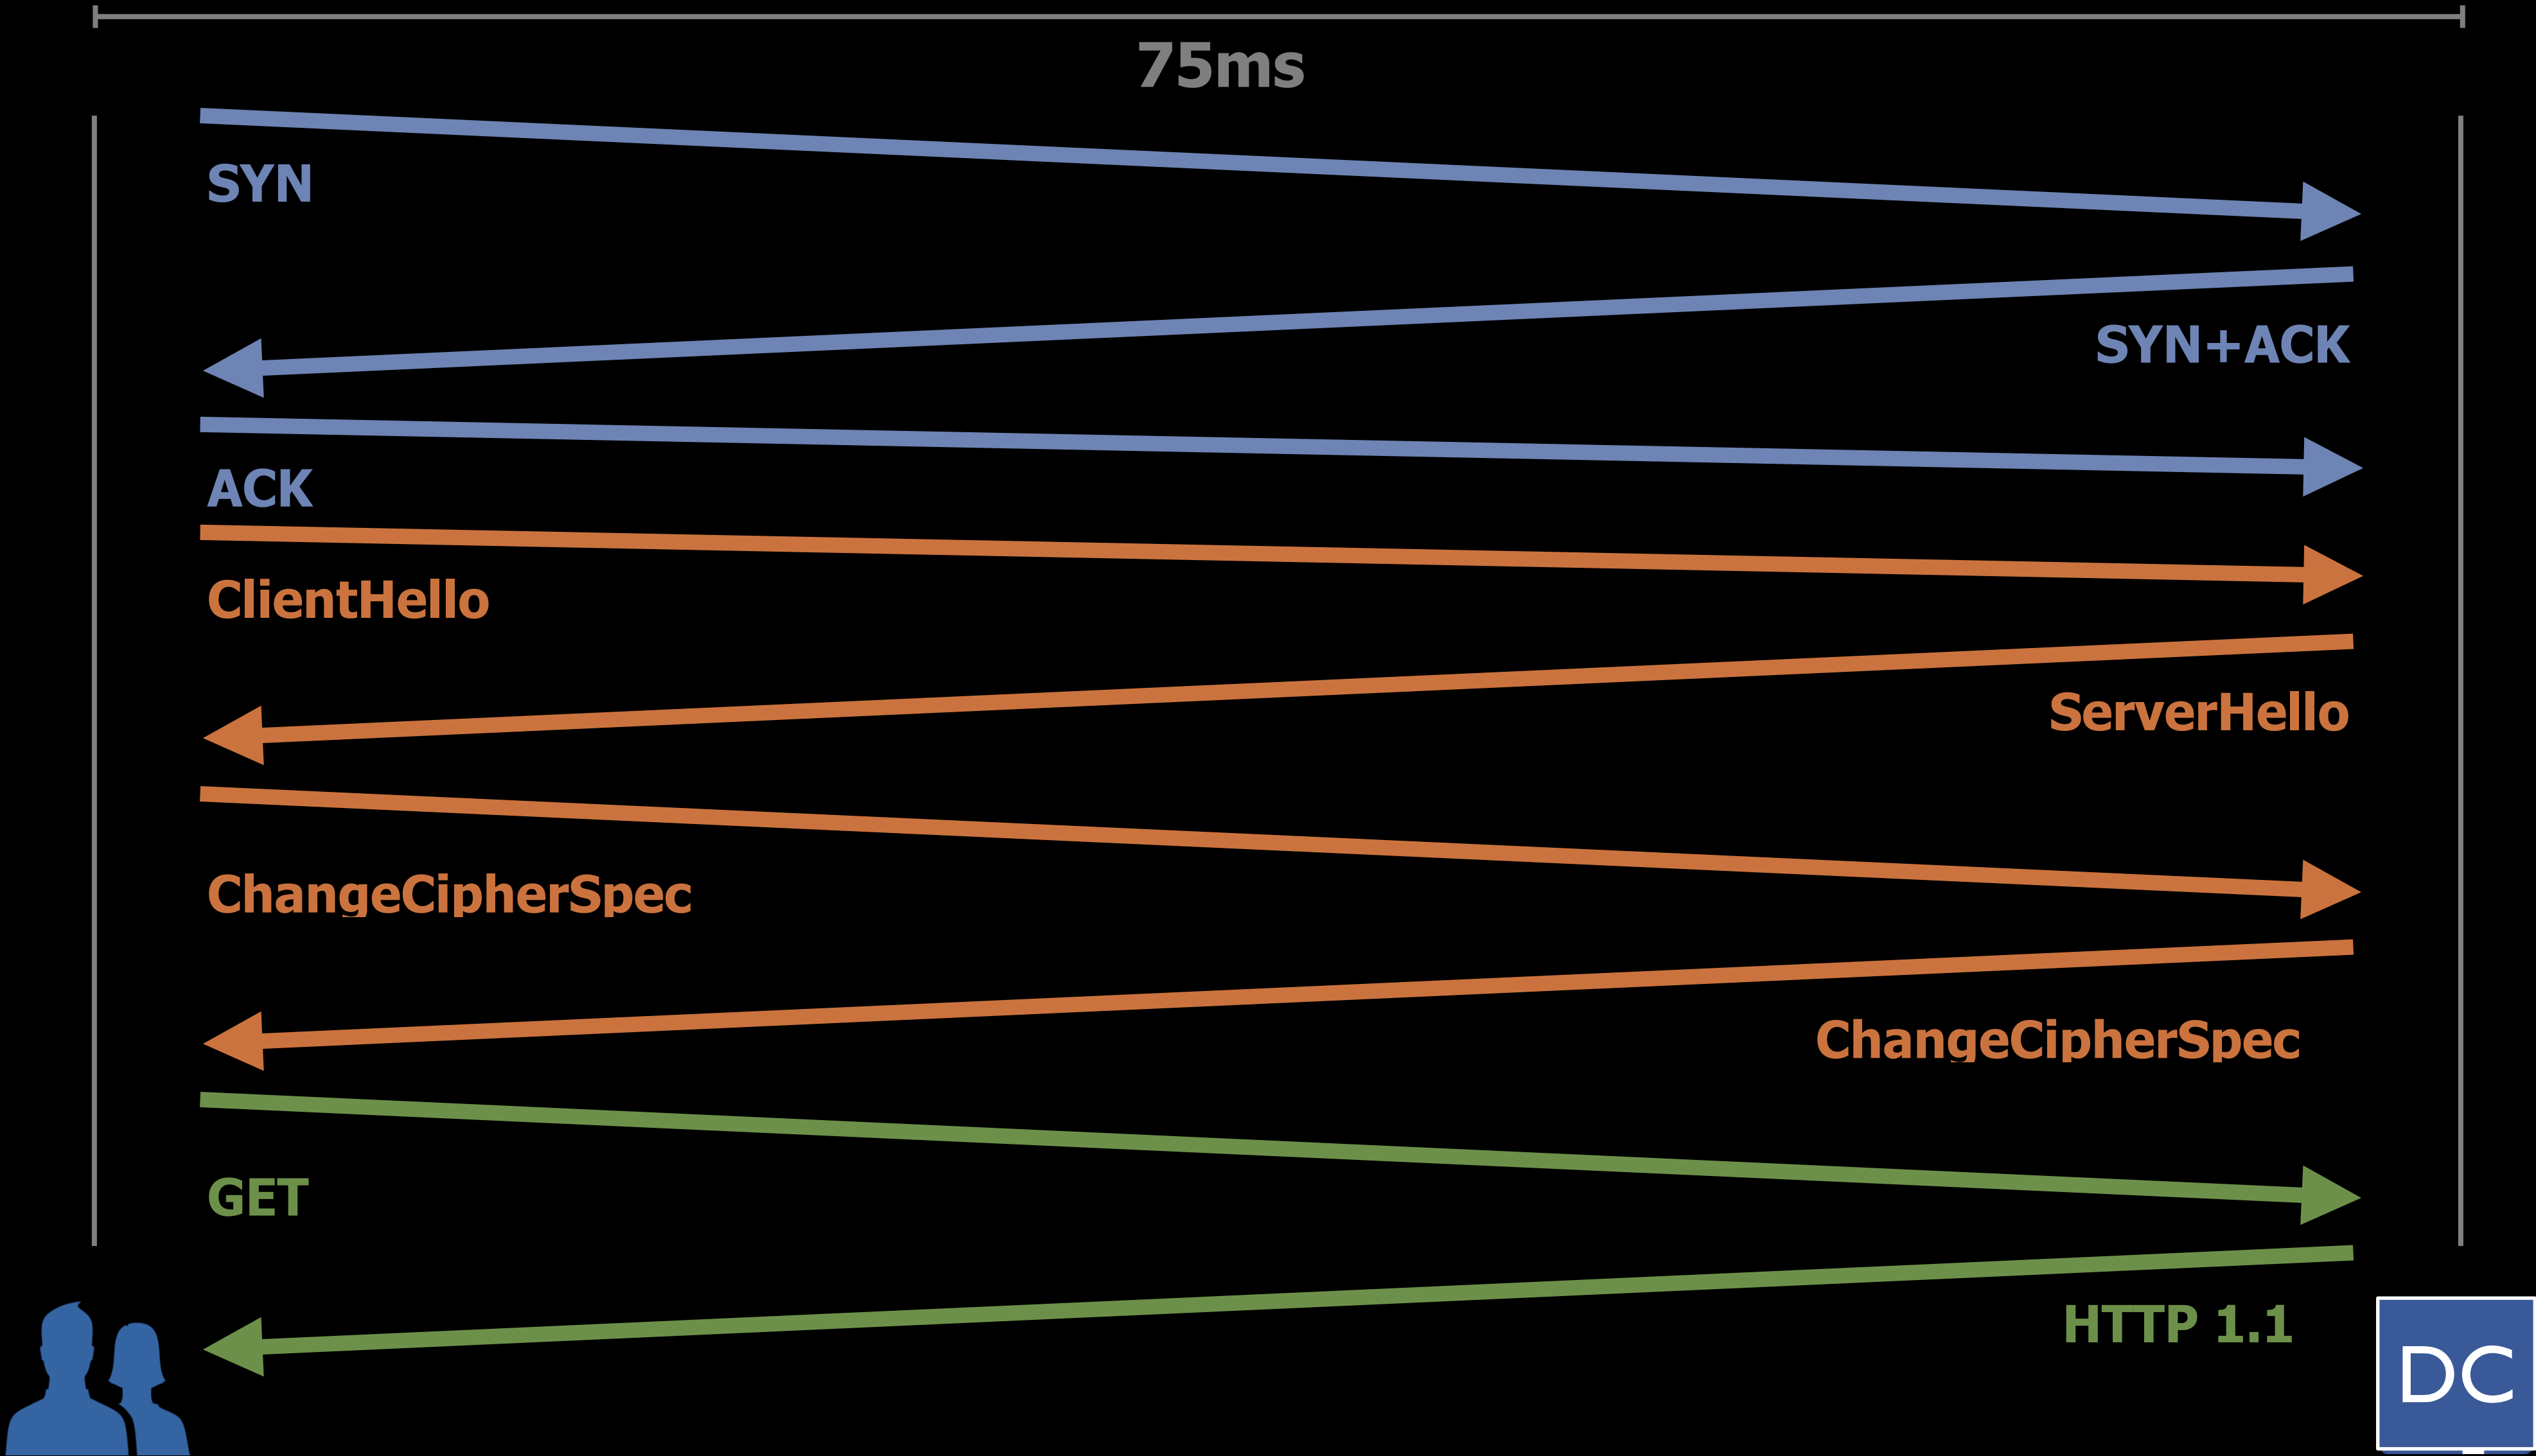


HTTPS Asia -> Oregon

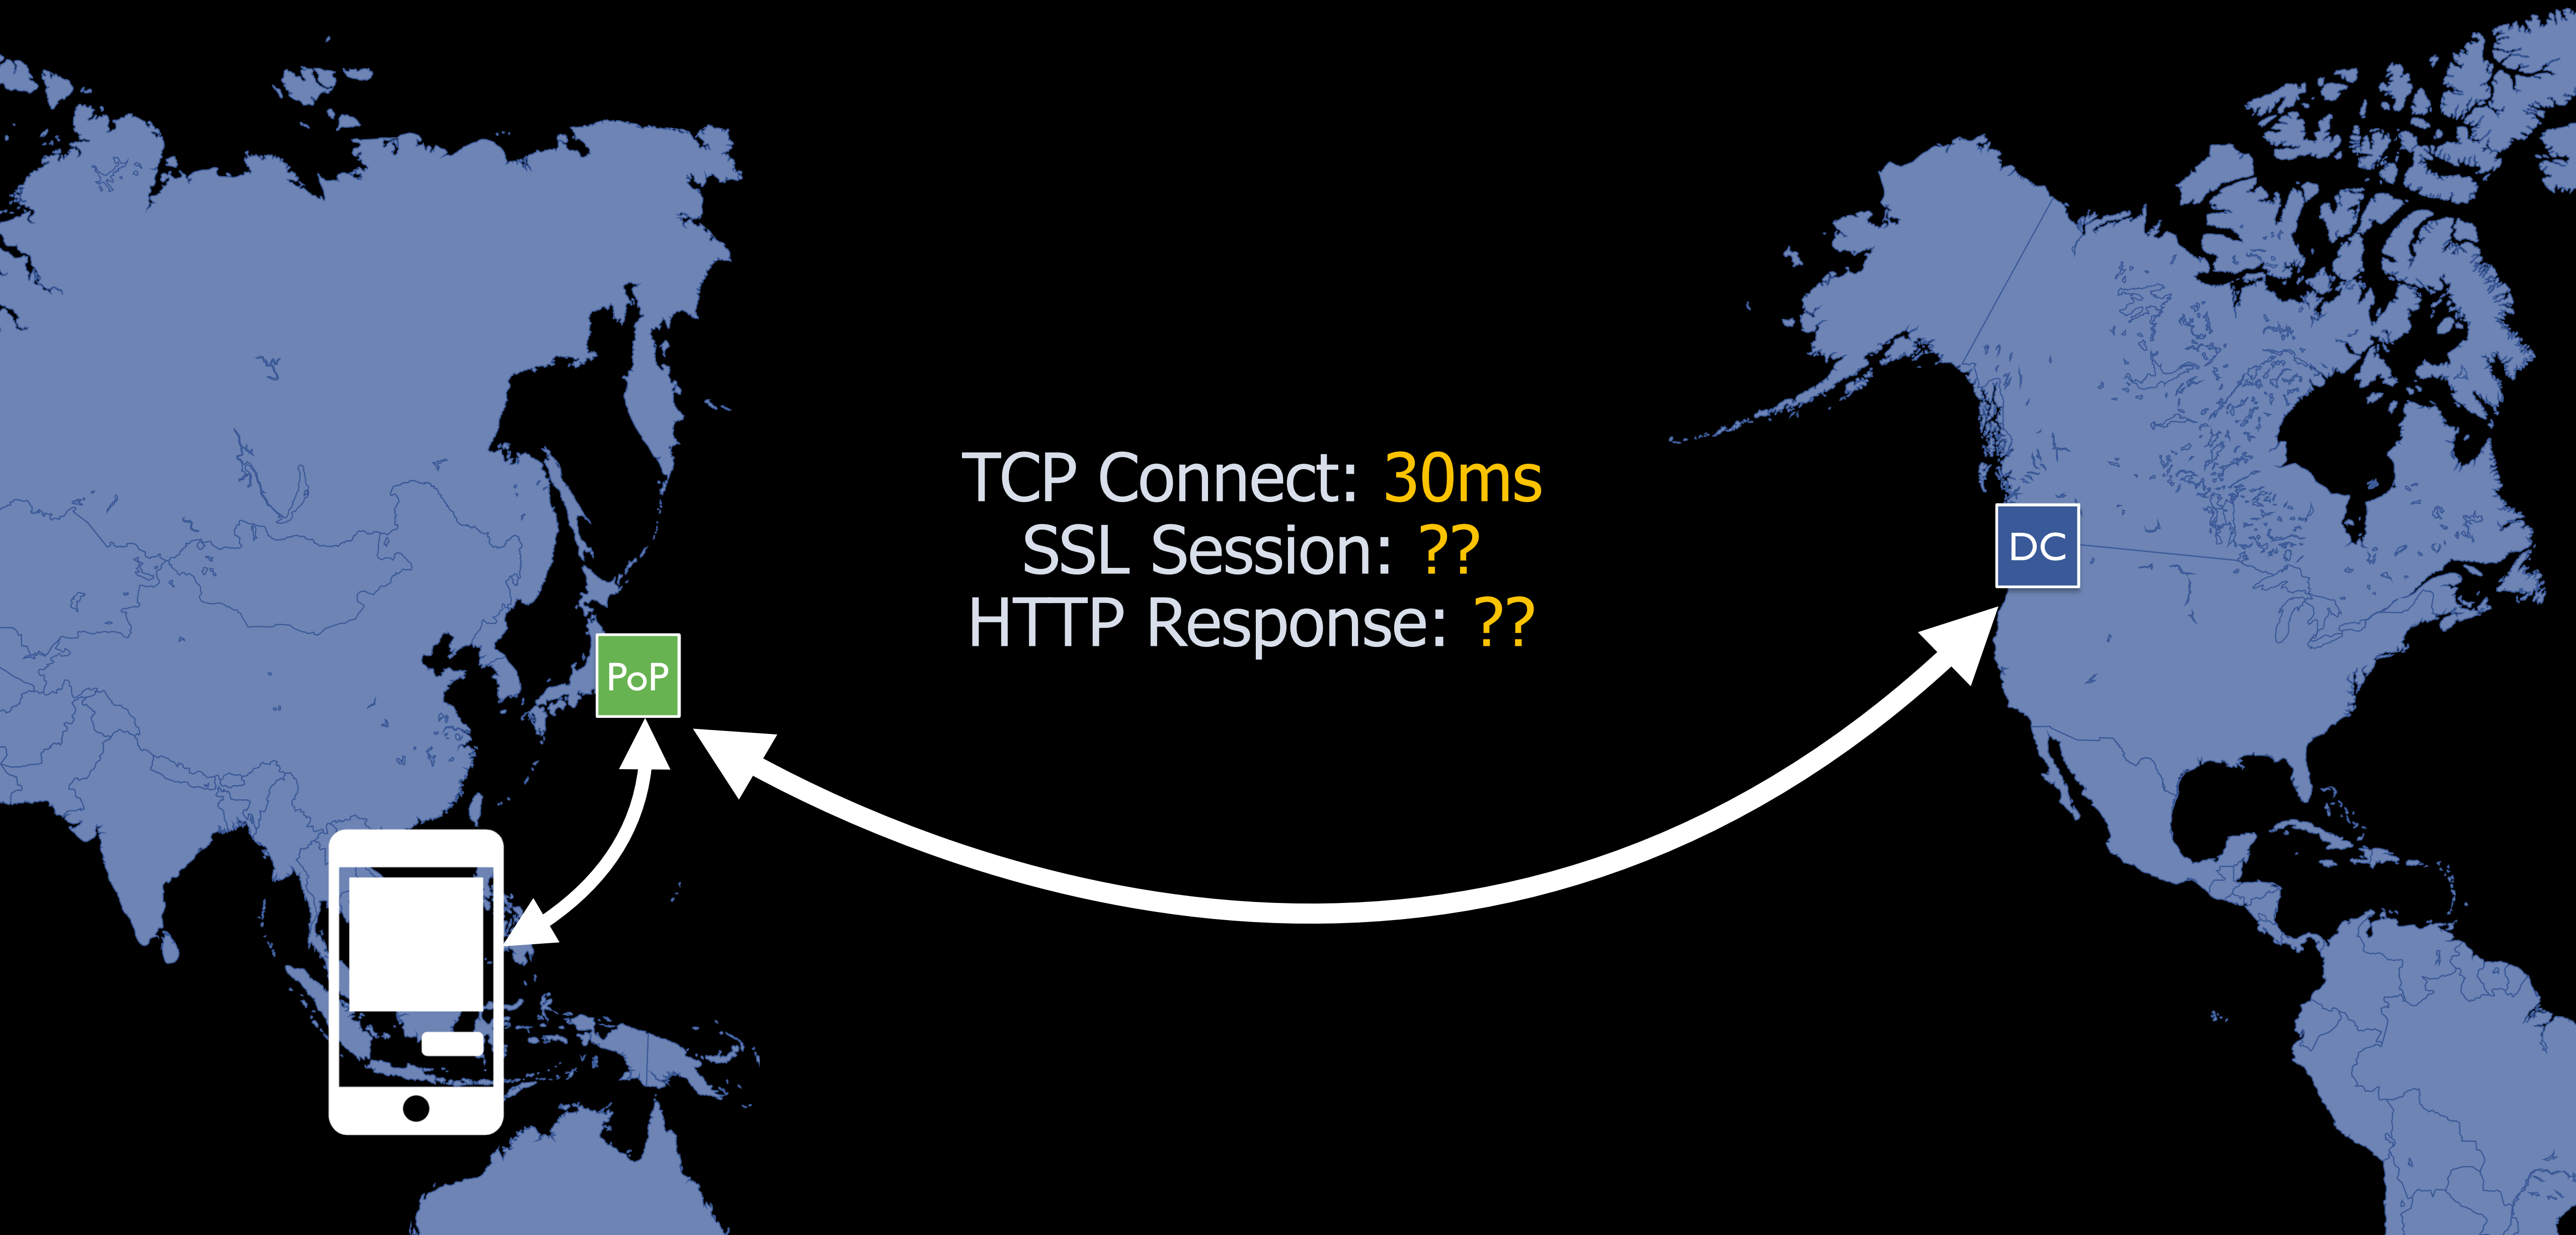
TCP conn
established:
150 ms

SSL session
established:
450 ms

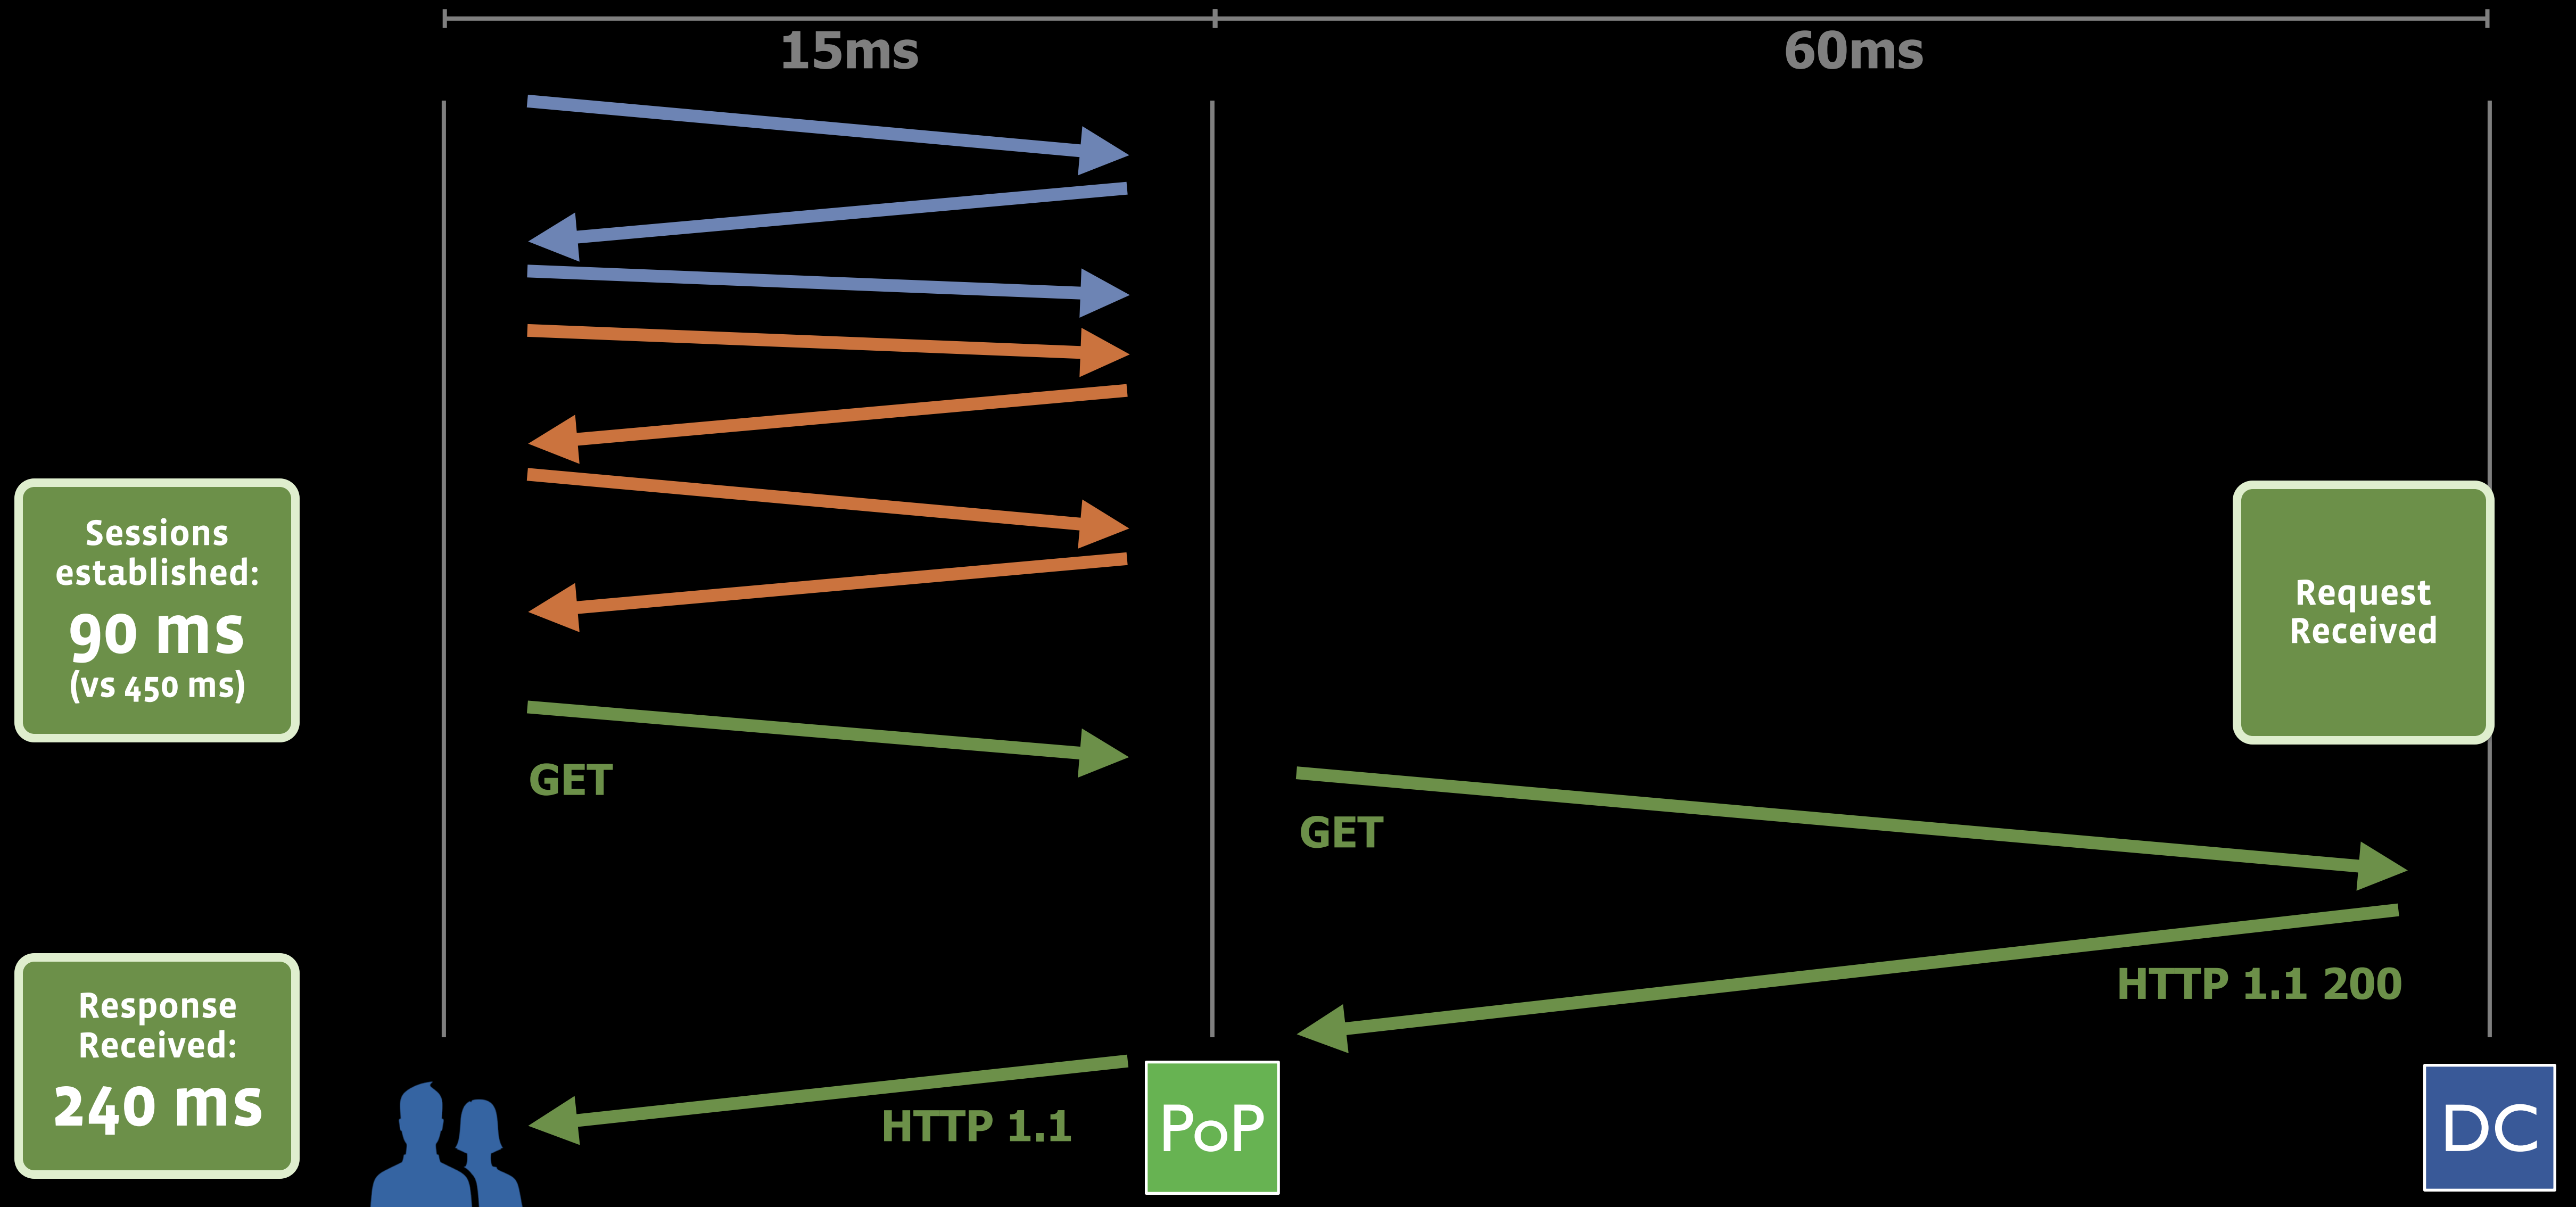
Response
Received
600 ms



Asia -> Oregon



HTTPS Asia -> POP -> Oregon



Asia -> Oregon

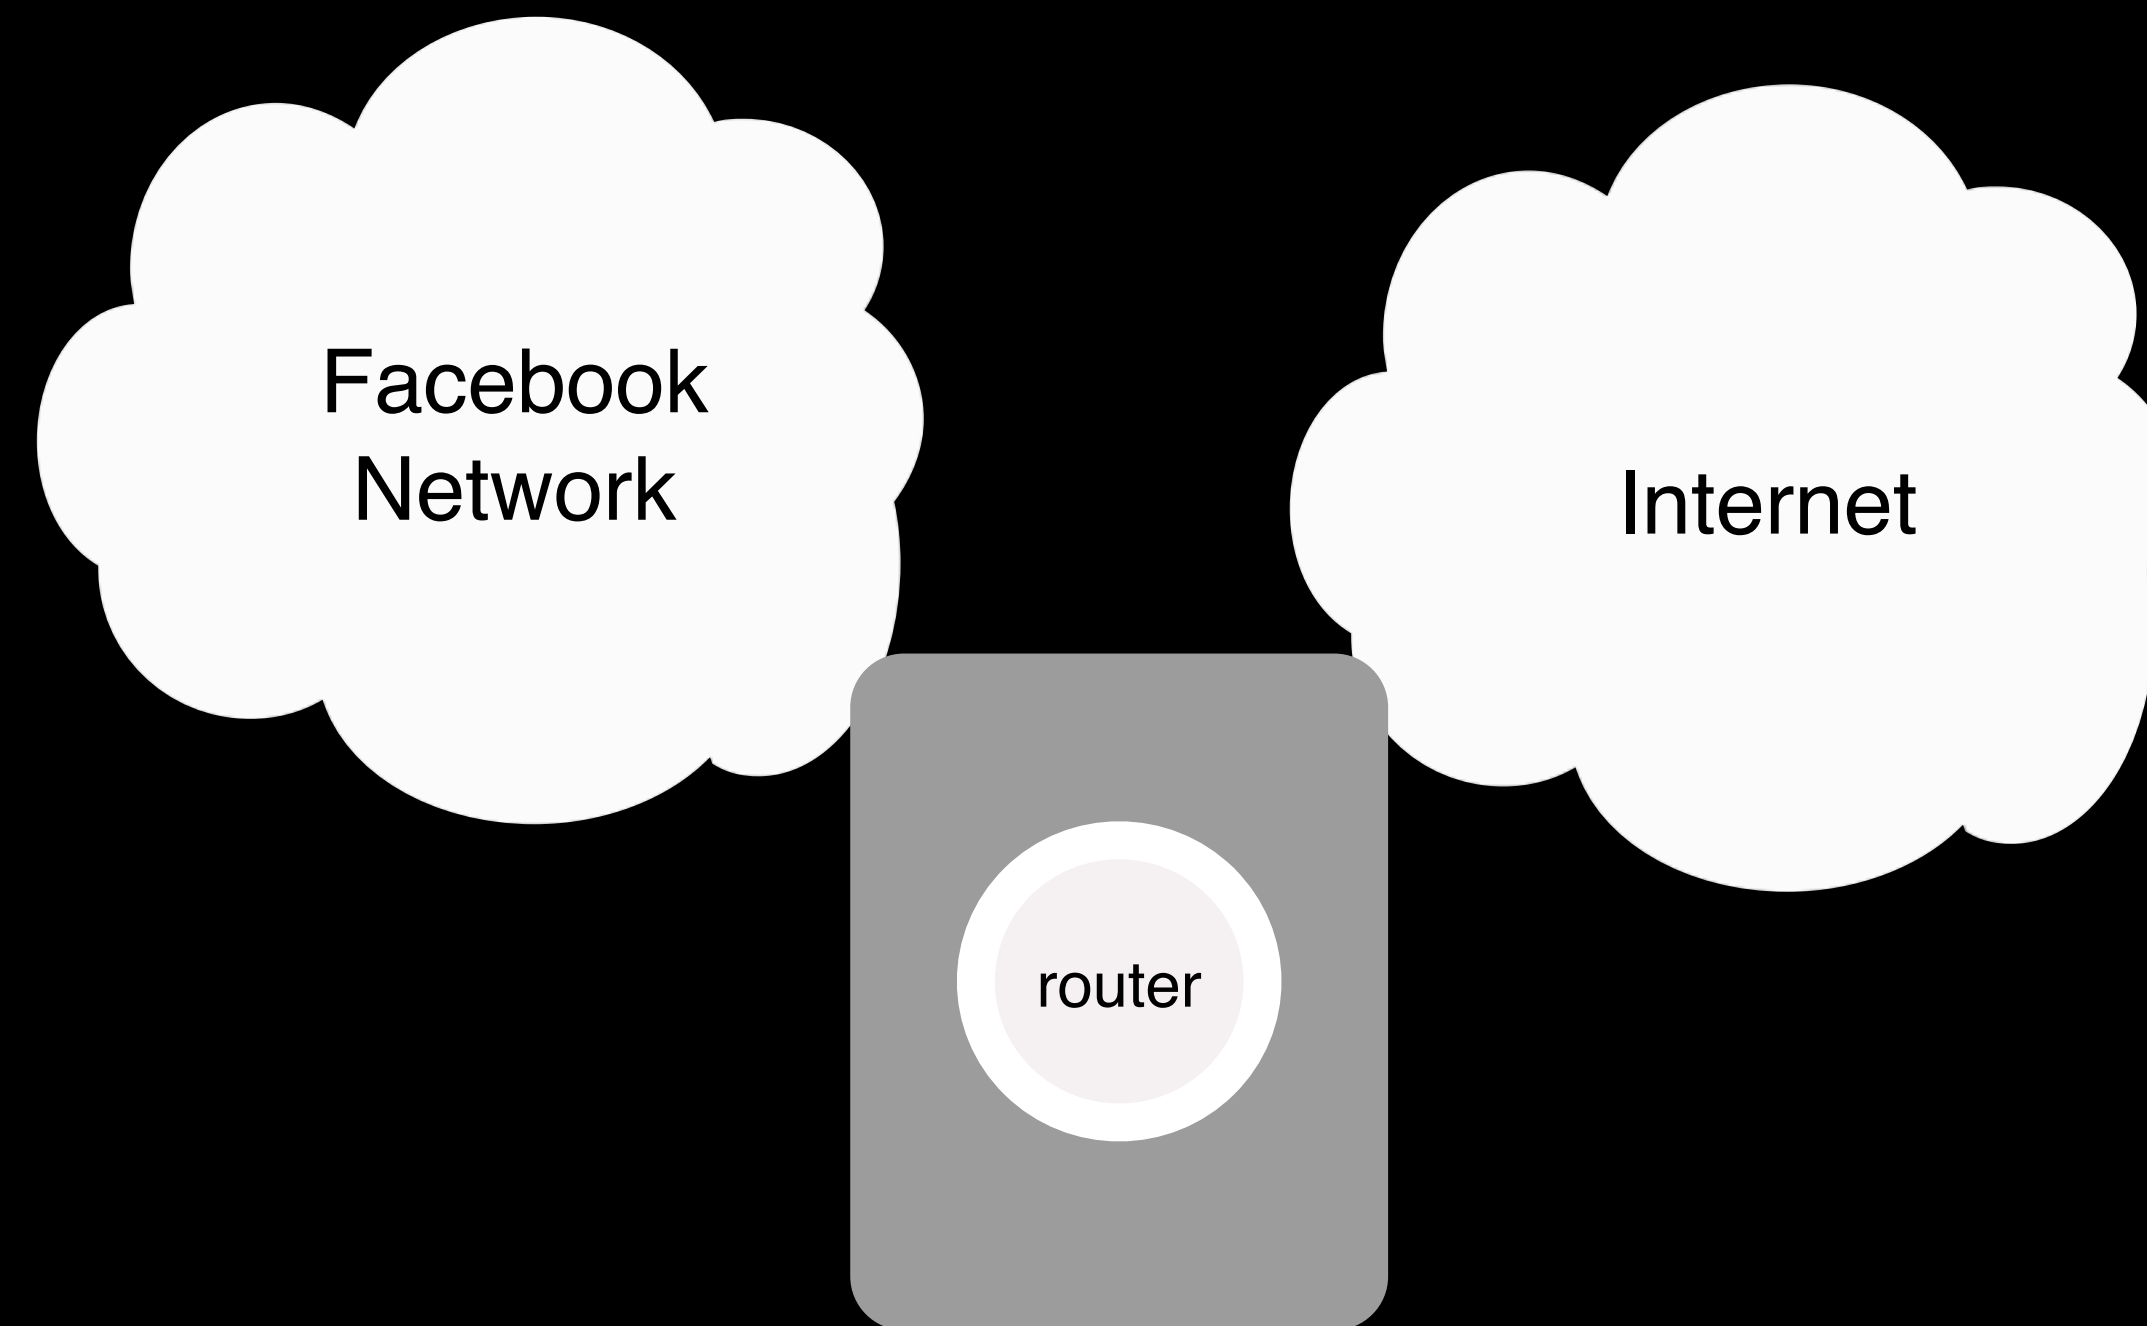


PoP

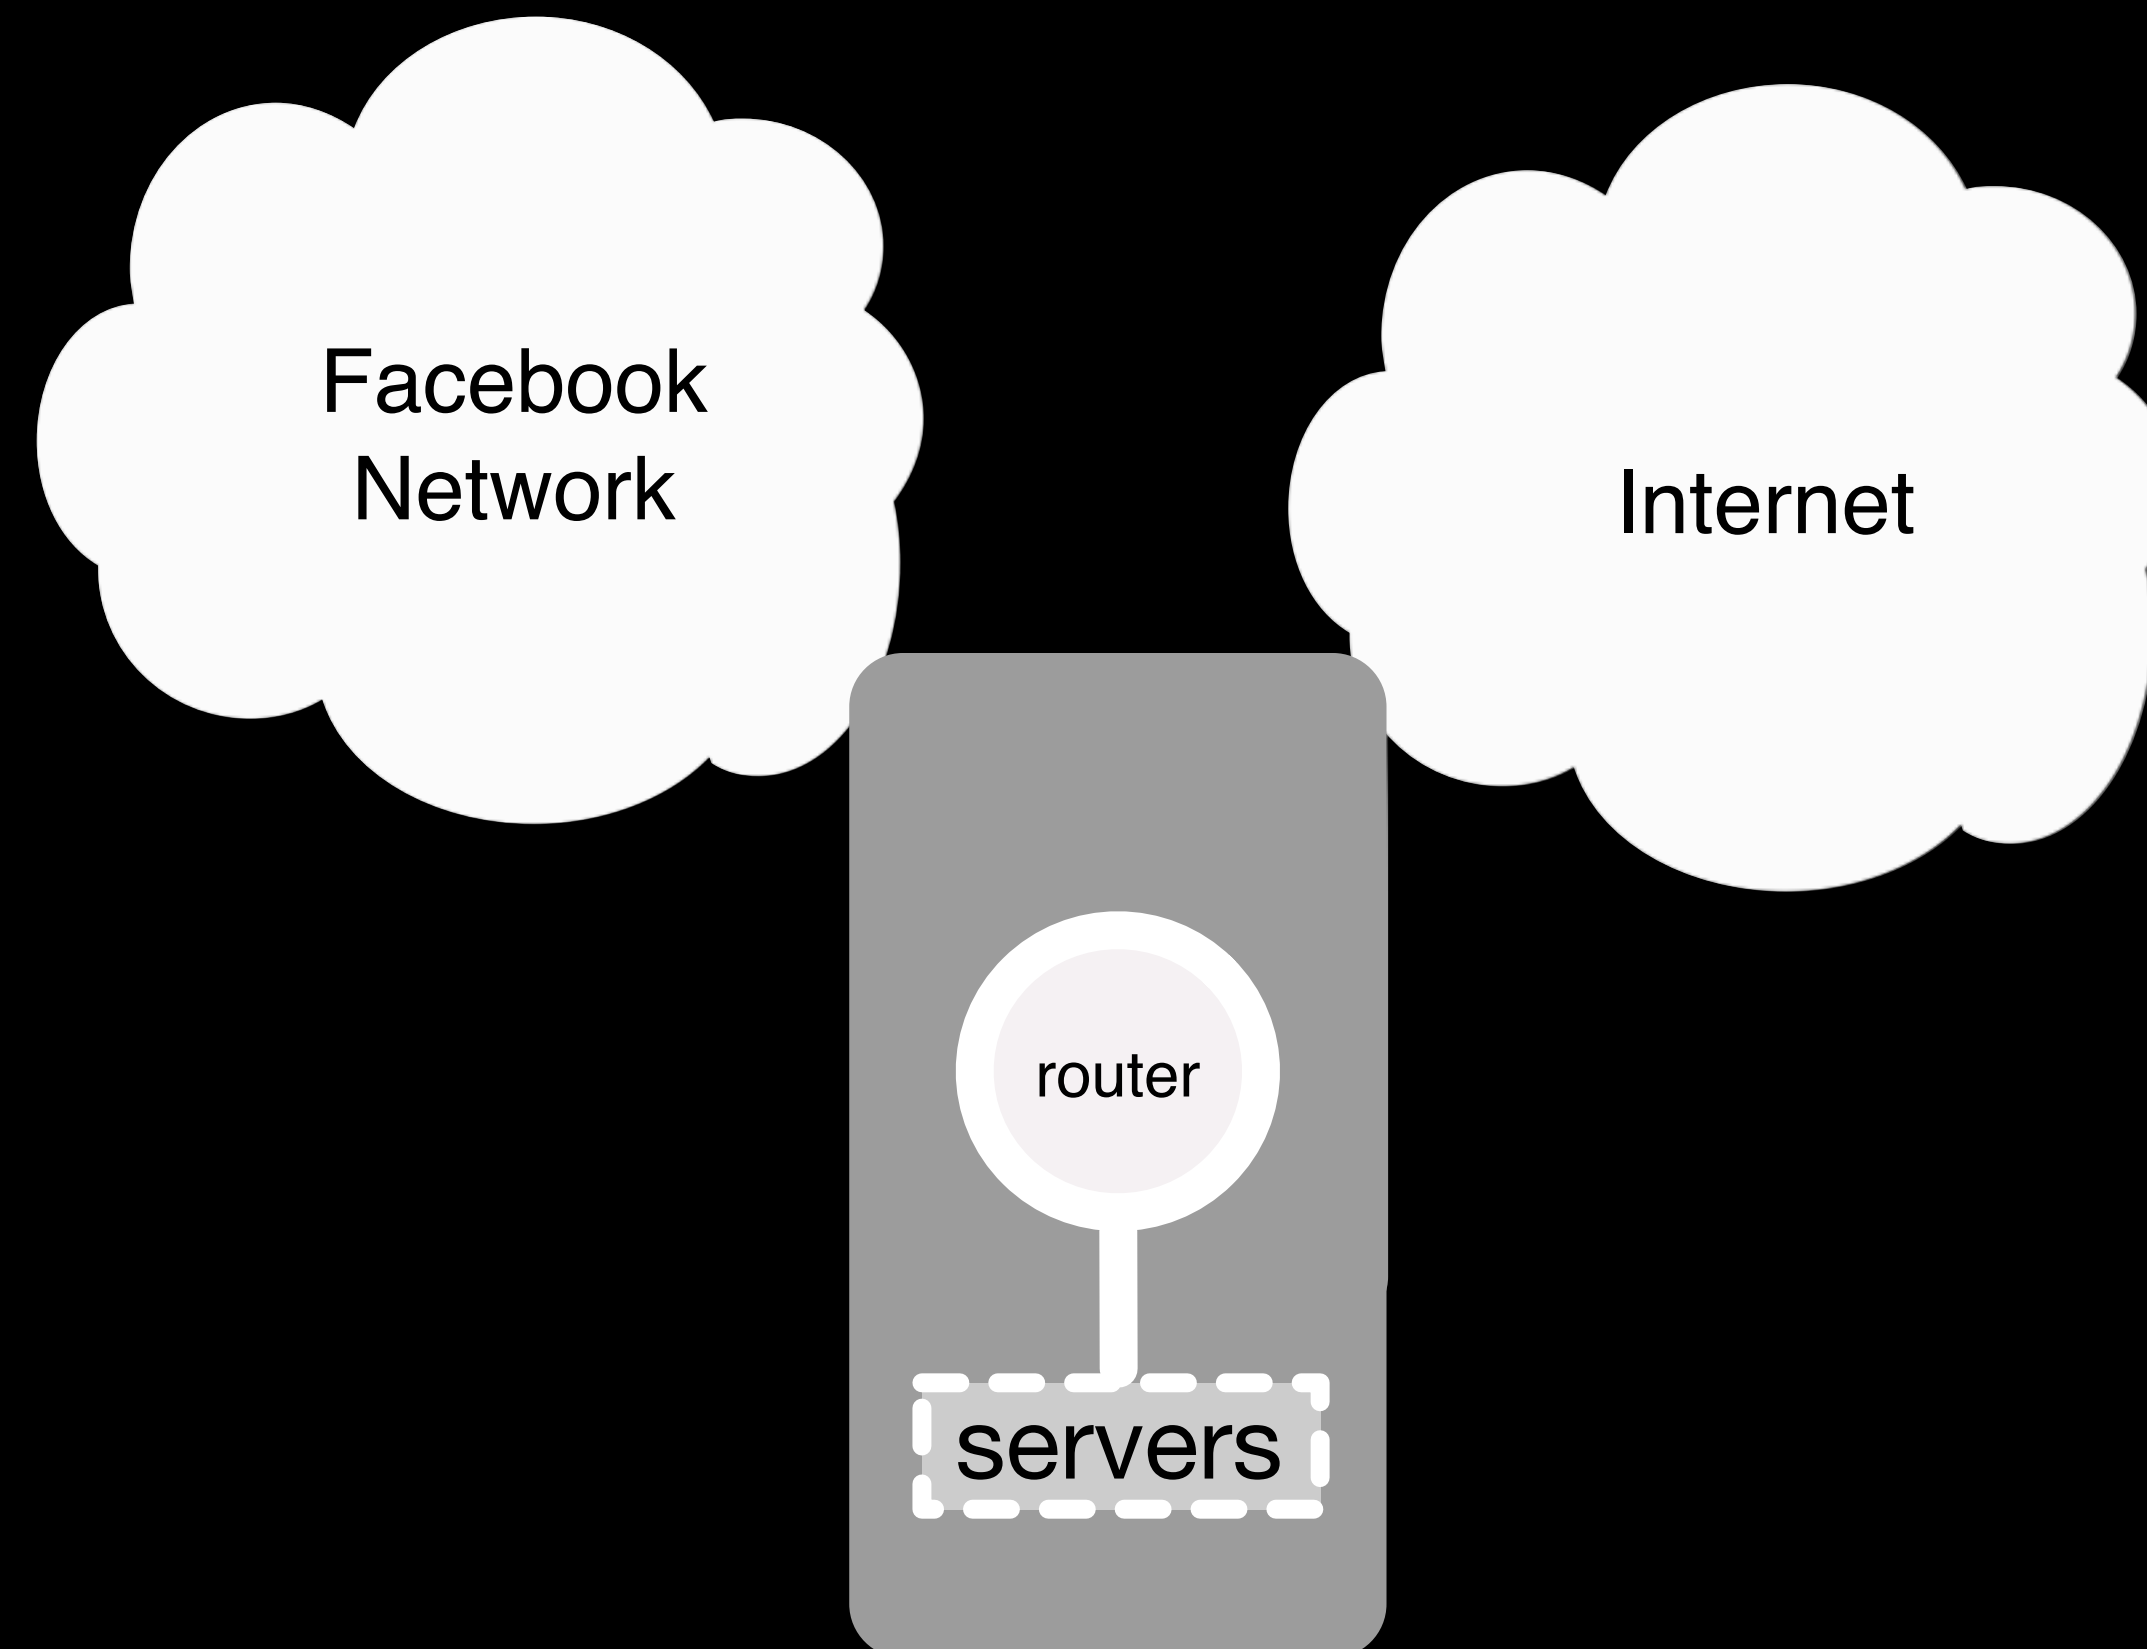
DC

TCP Connect: ~~150ms~~ **30ms**
SSL Session: ~~450ms~~ **90ms**
HTTP Response: ~~600ms~~ **240ms**

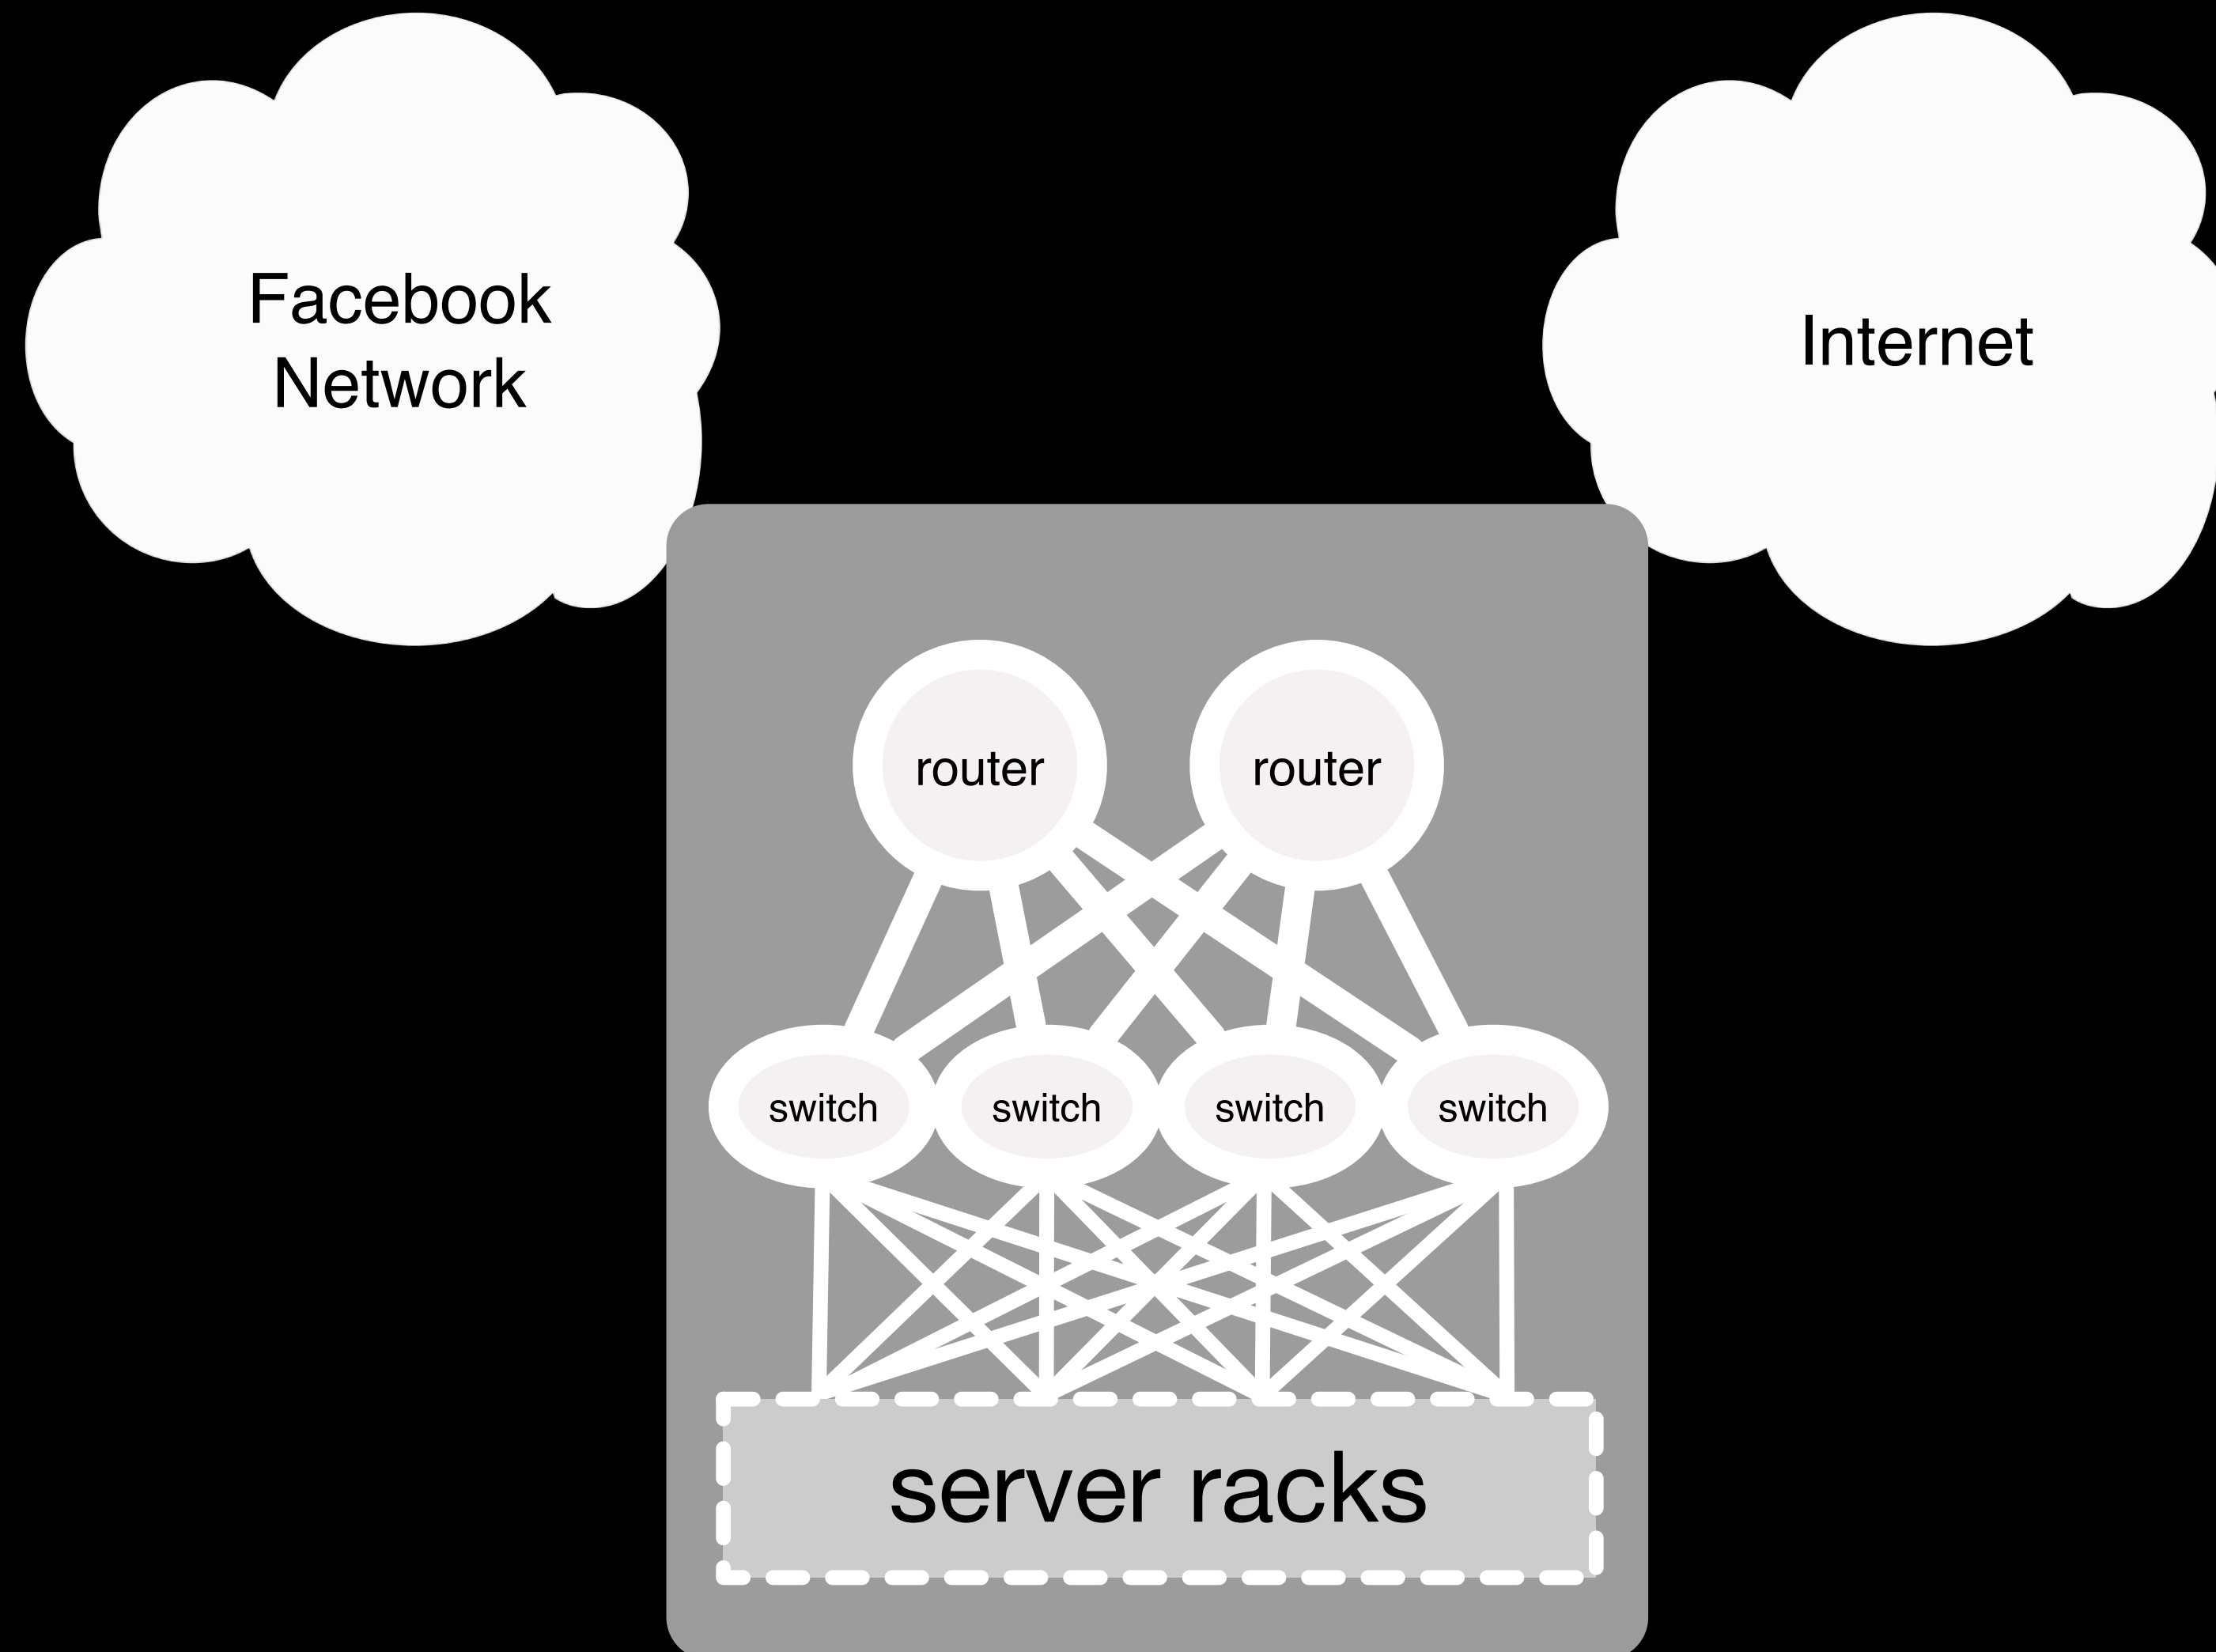
edge routers -> edge clusters



edge routers -> edge clusters




edge routers -> edge clusters



LIVE

Explore to see who's live around the world





Mark Zuckerberg is live now.
Public • 2 mins
June 1st 2016



 Like

 Comment

 Share






9.7K

277 shares




يمّا ال انكزي
Hey mark, congrats on the space chat initiative :)
Just now • Like • Reply

Video Is Growing



Video Discussion



Share Your Thoughts
What do you think about the presidential debate? Share your thoughts with friends.

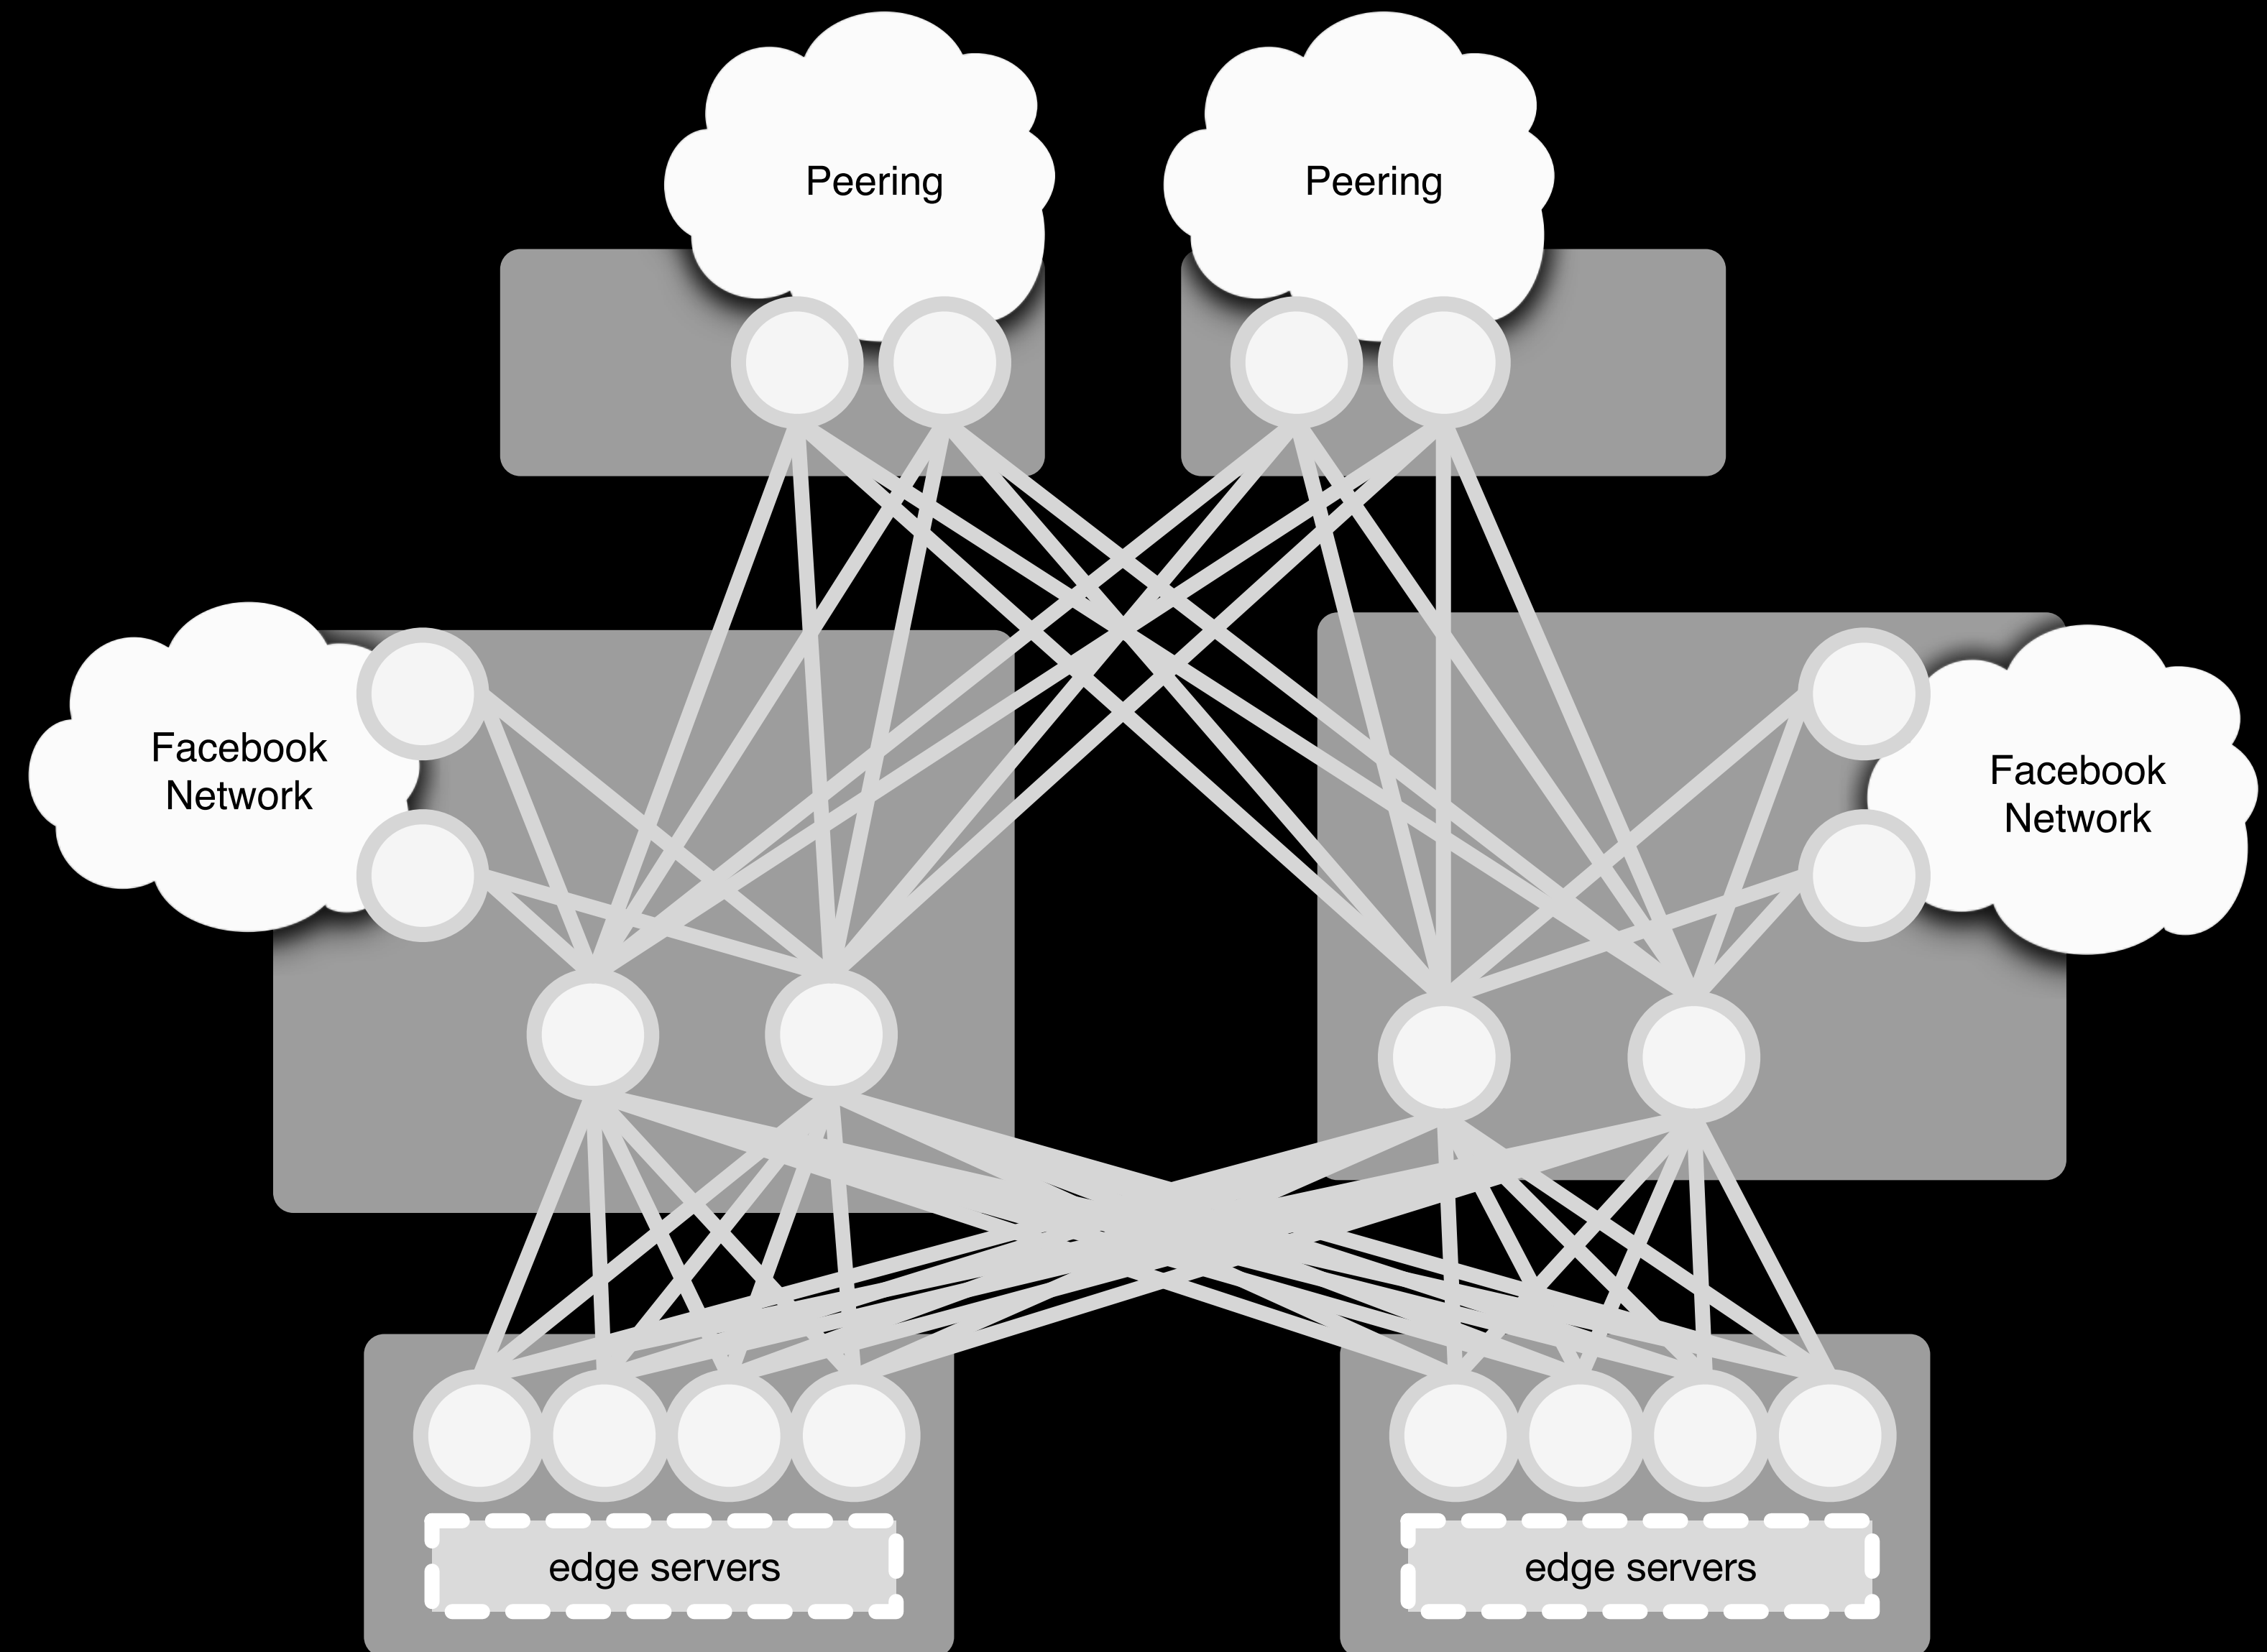
Post



Christopher Kusek
Just now • Twitter • 

My favorite part about the Sony hack investigation by the FBI was where it was definitely without a doubt North Korea. Or, not. #Debates

-> edge metro topology

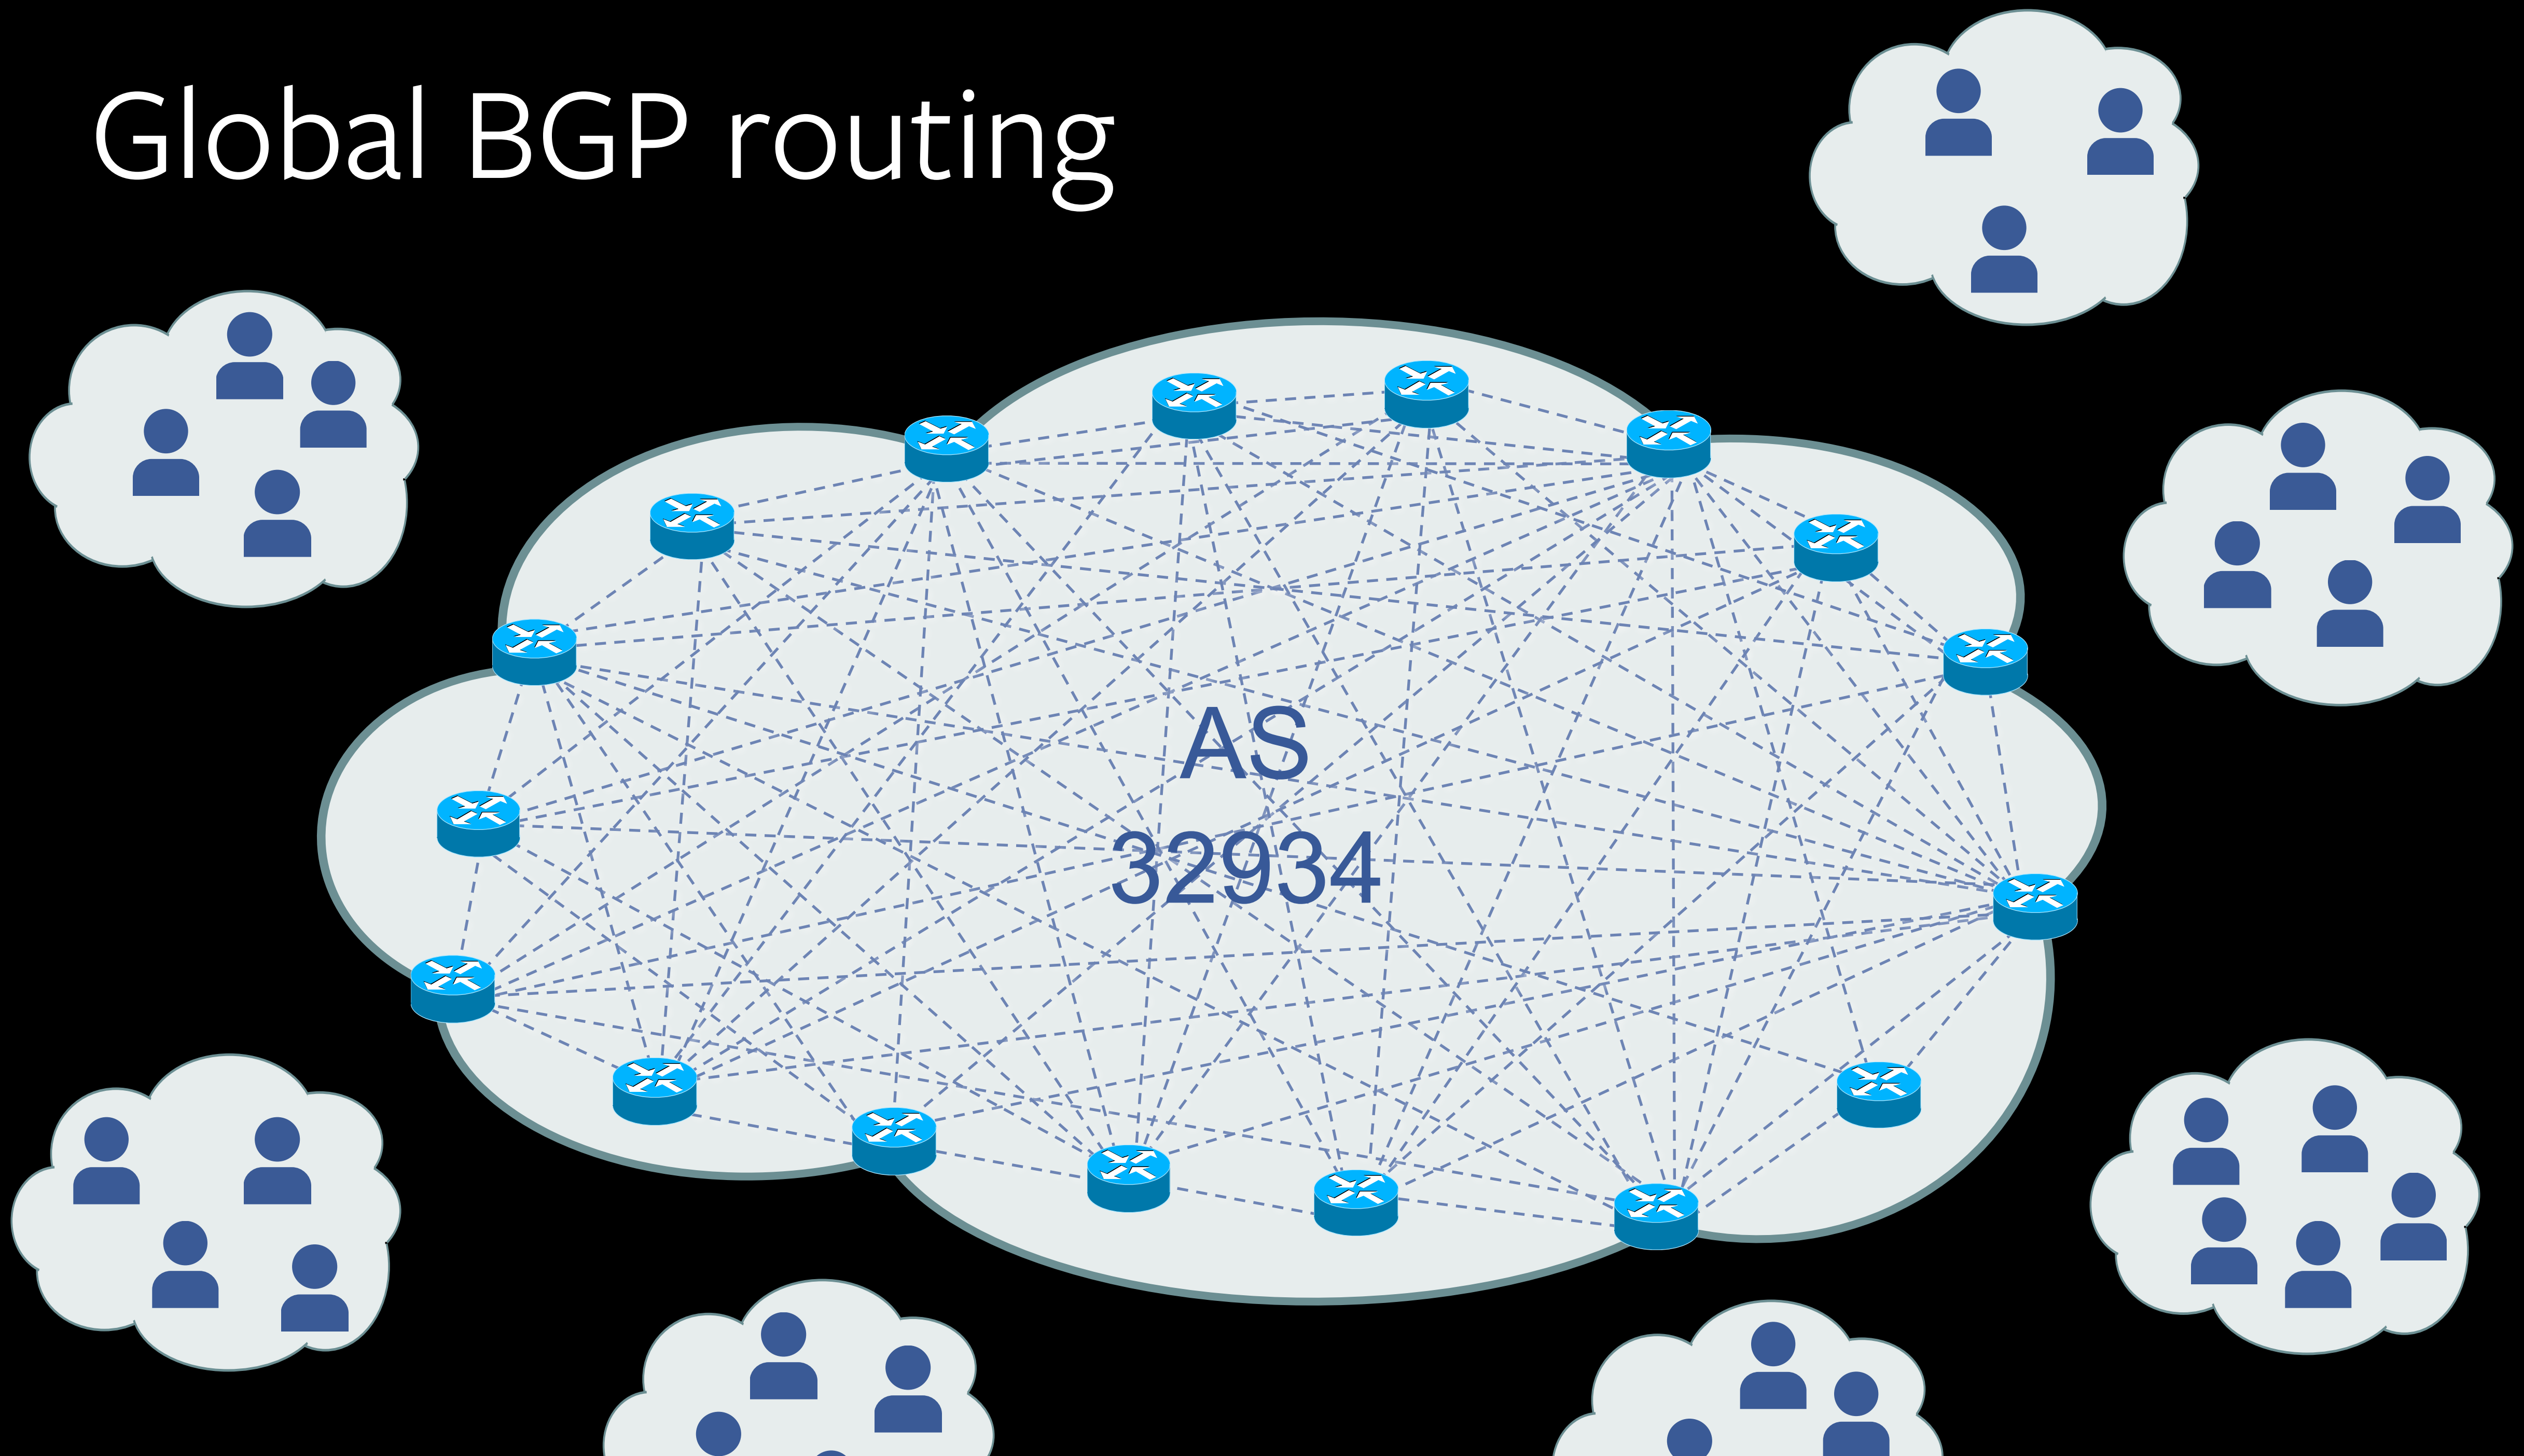


100G

Everywhere!

Evolving beyond BGP

Global BGP routing



Which POP is best?

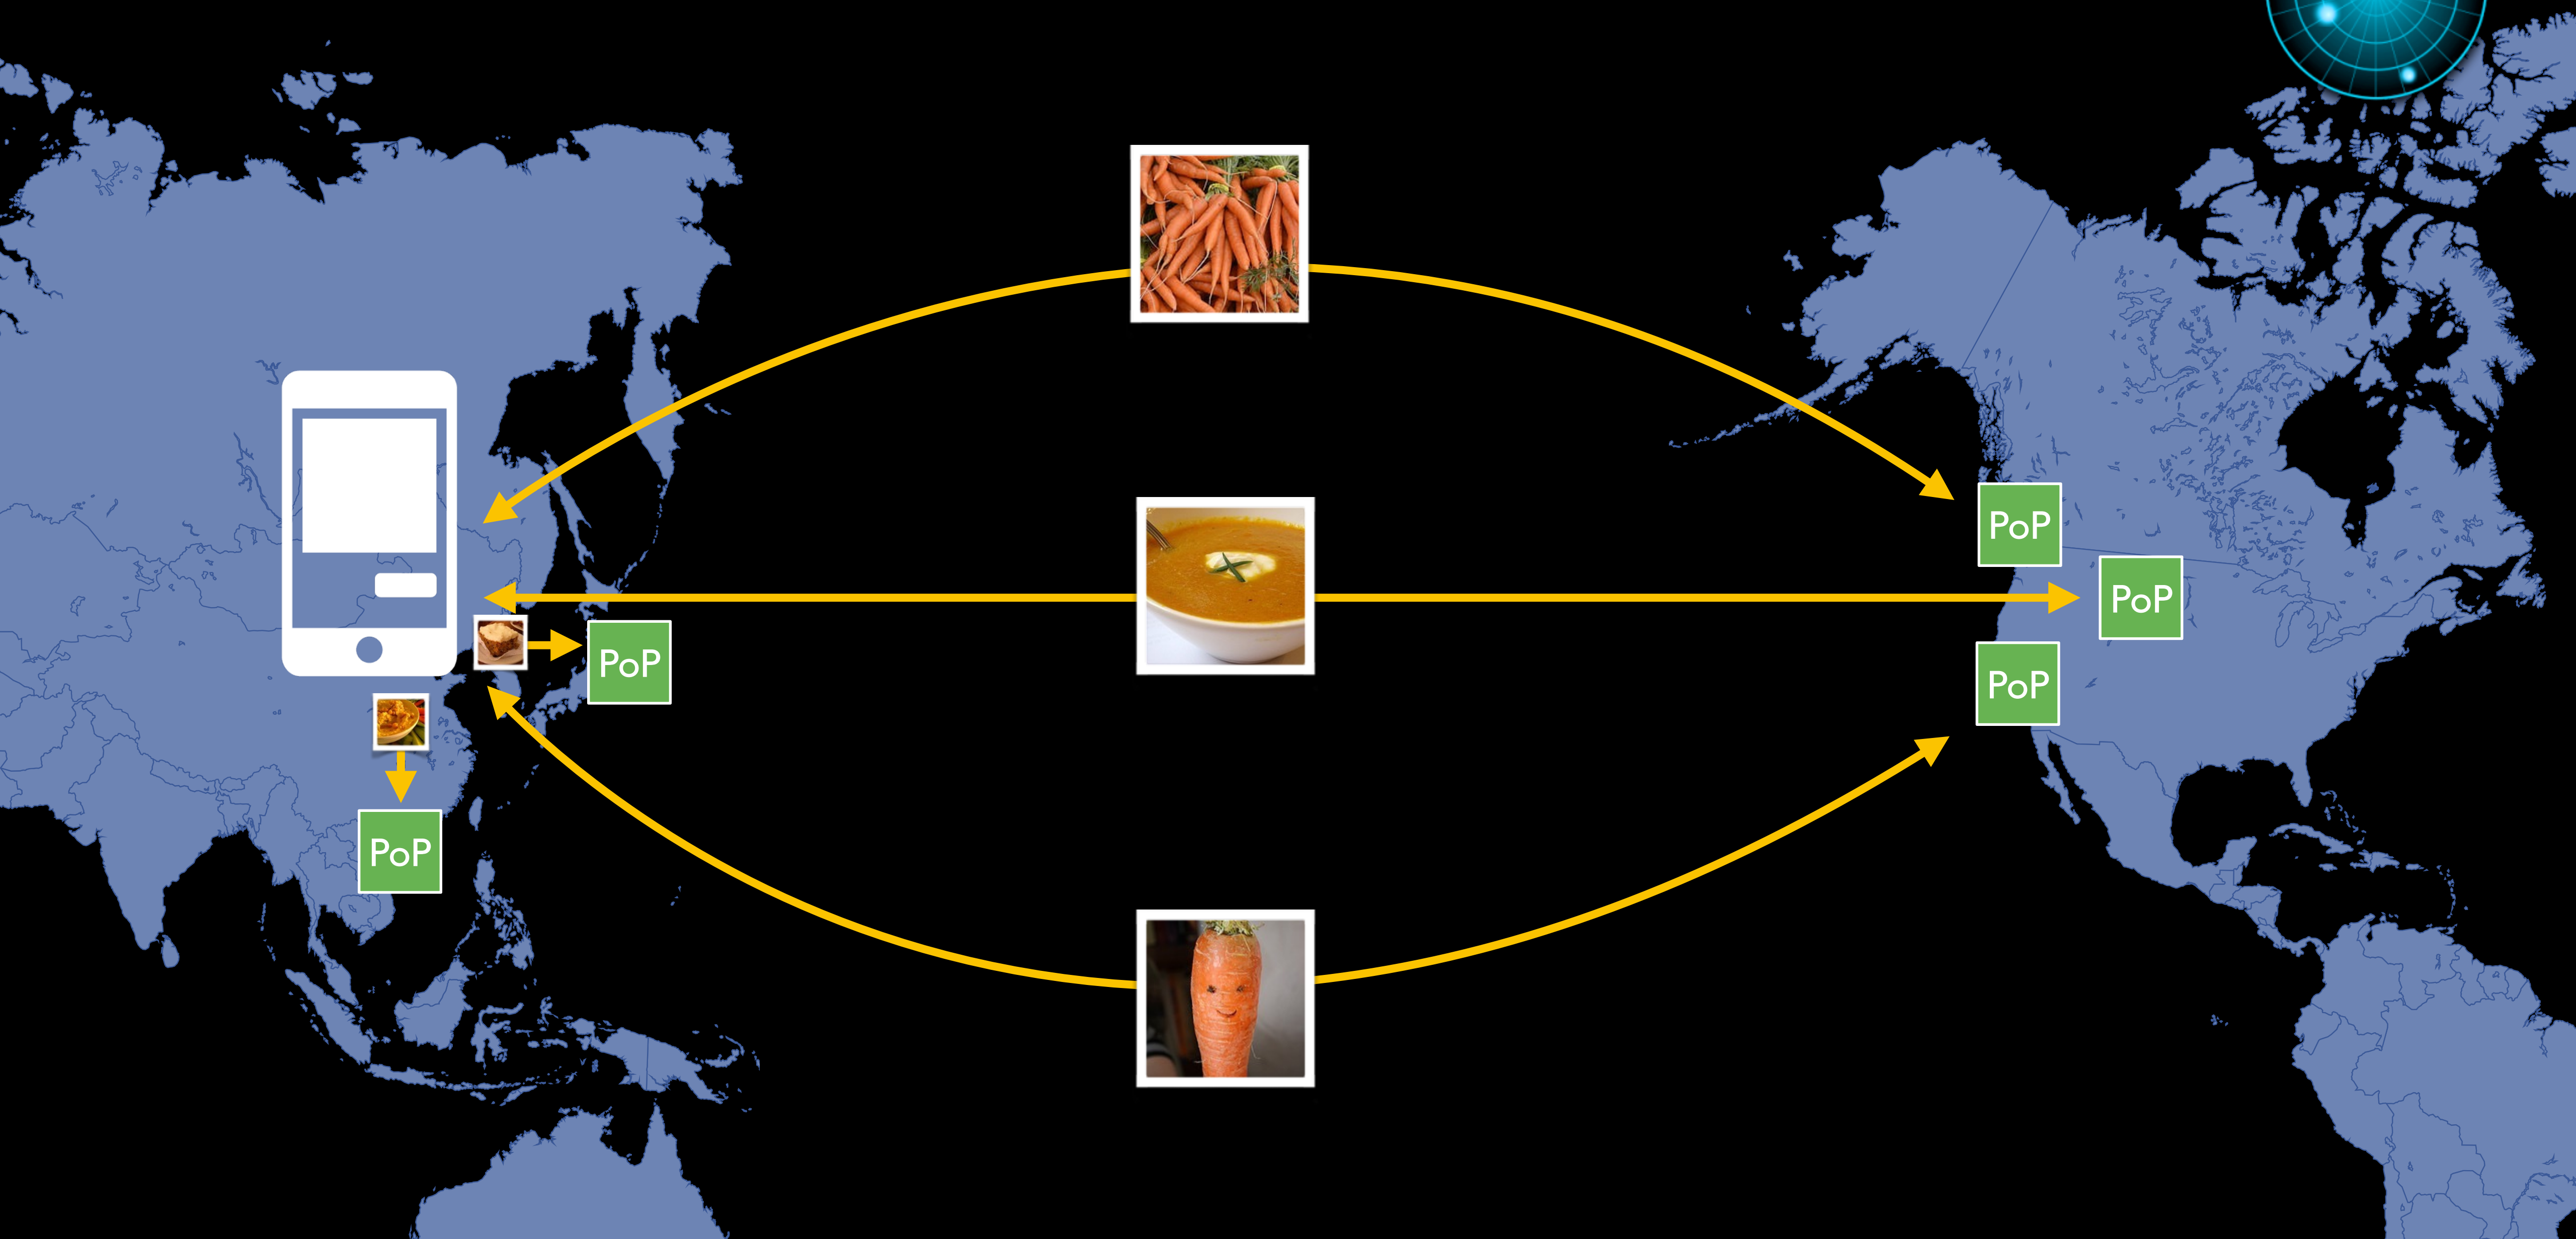
Considerations:

- Closest Edge to user

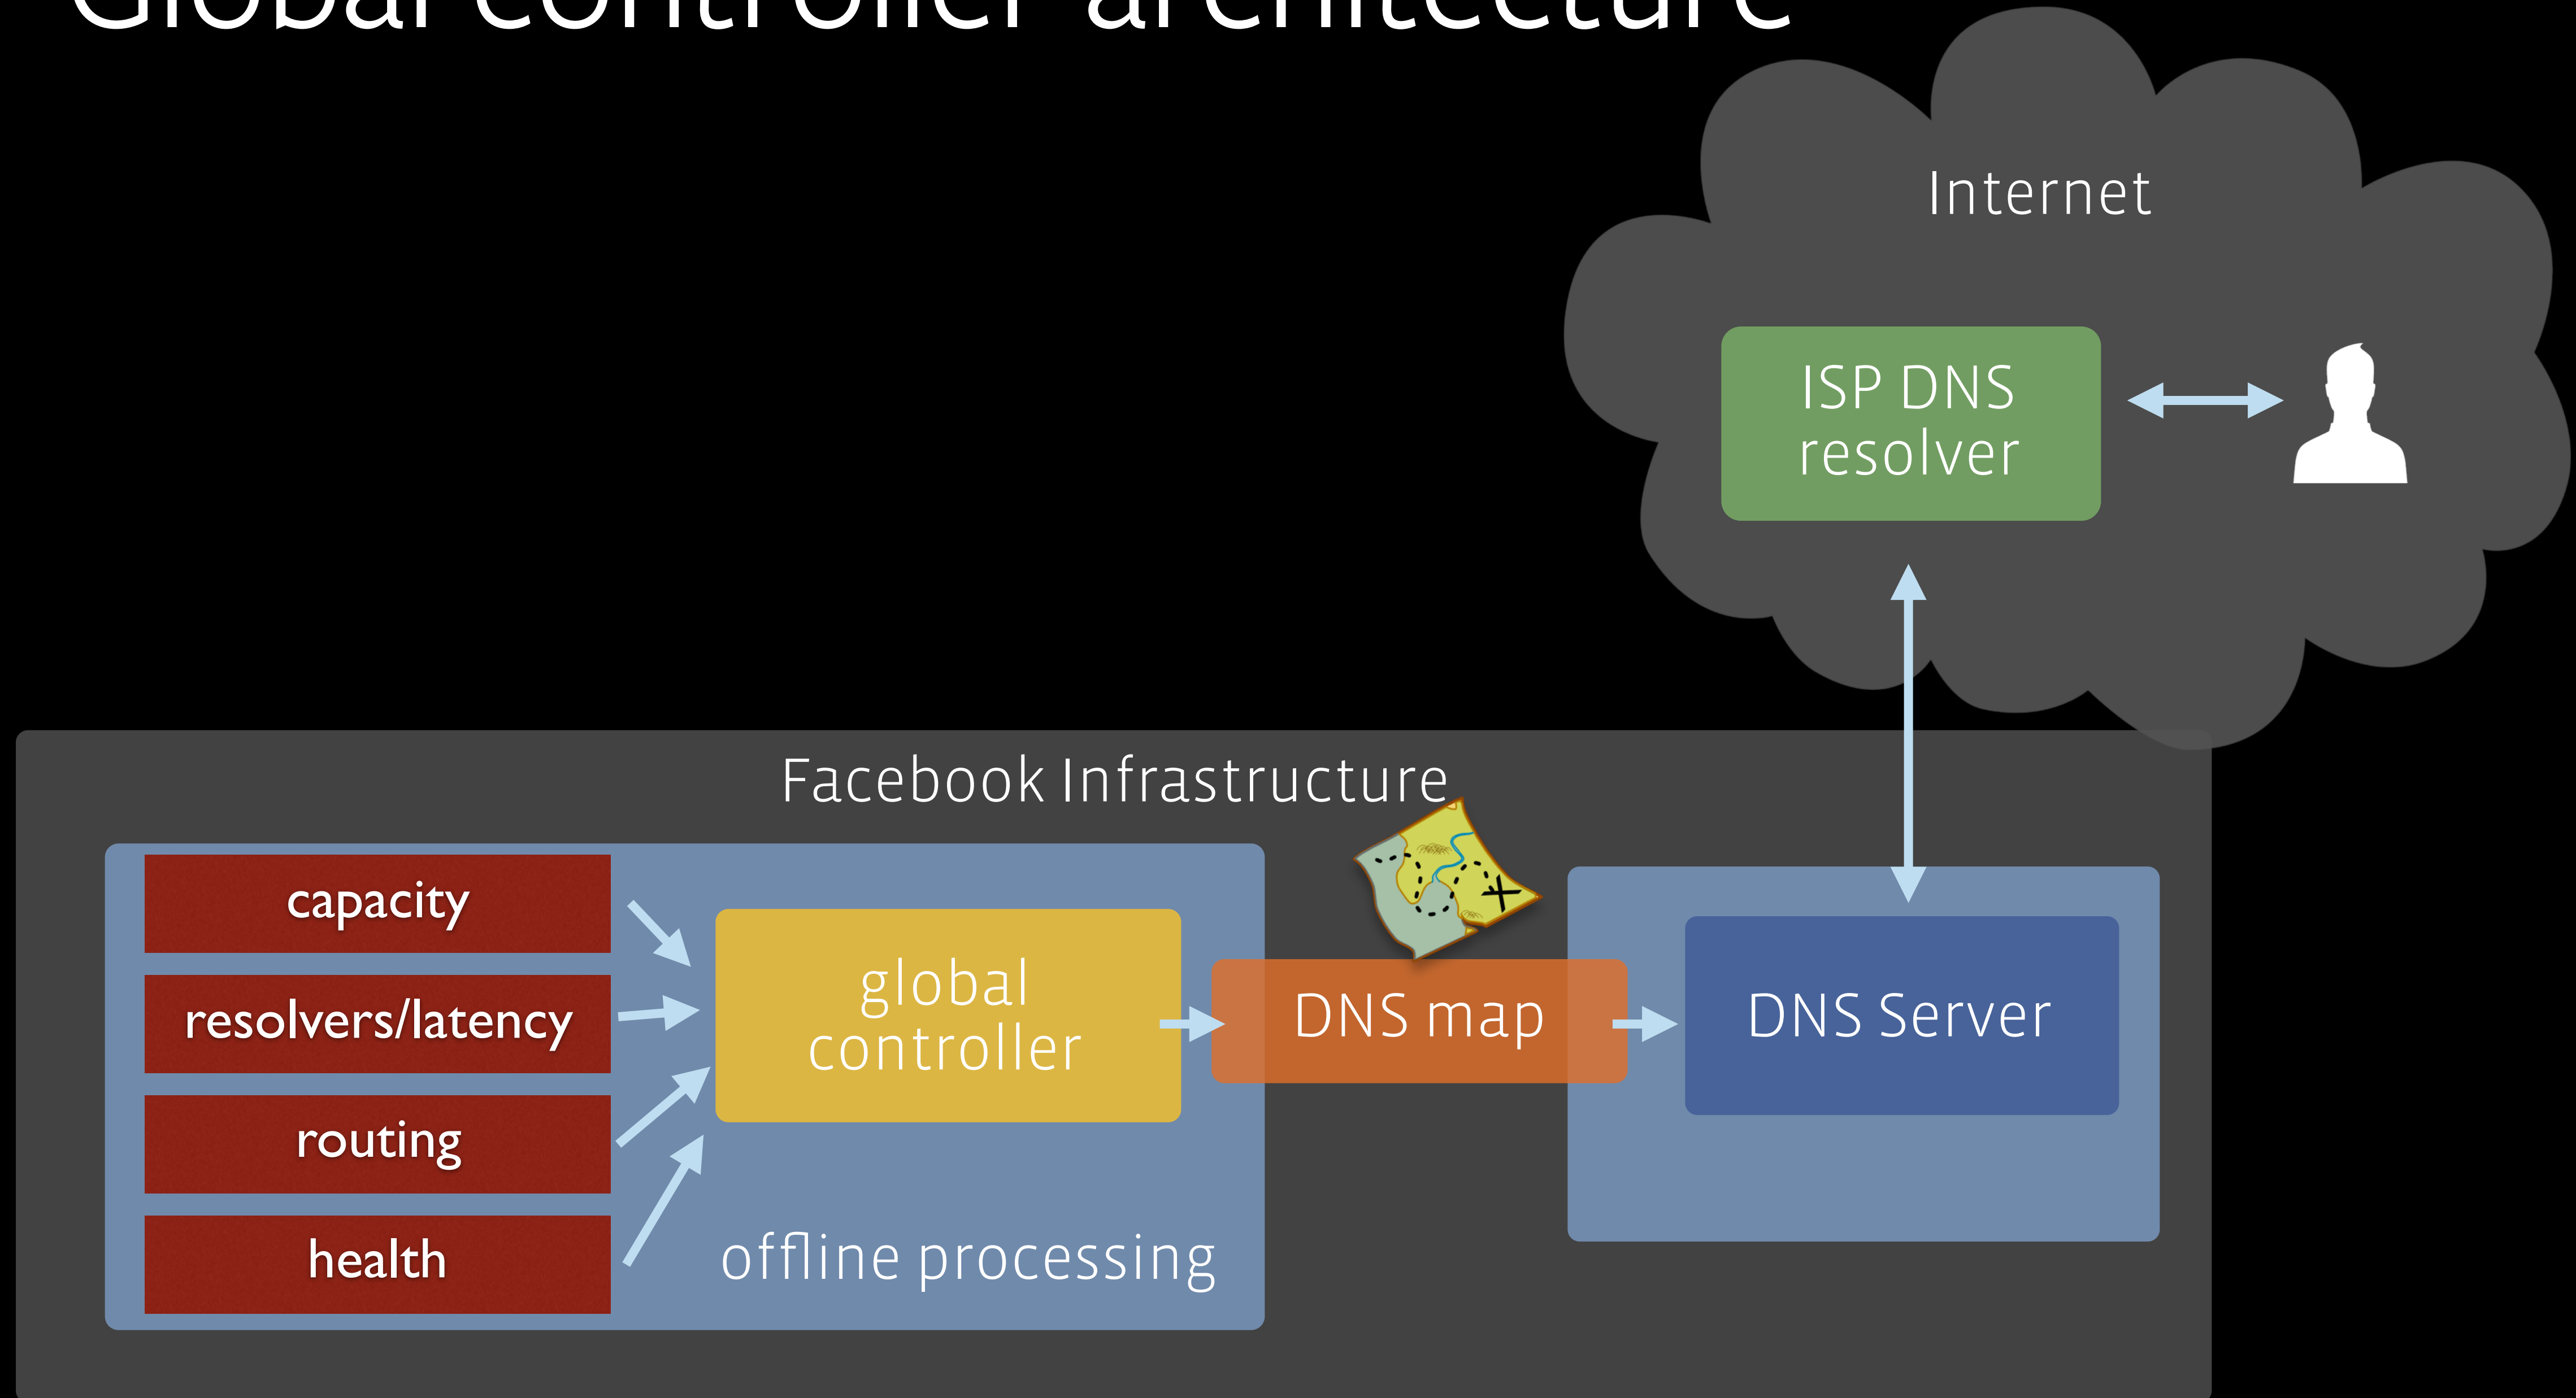


???

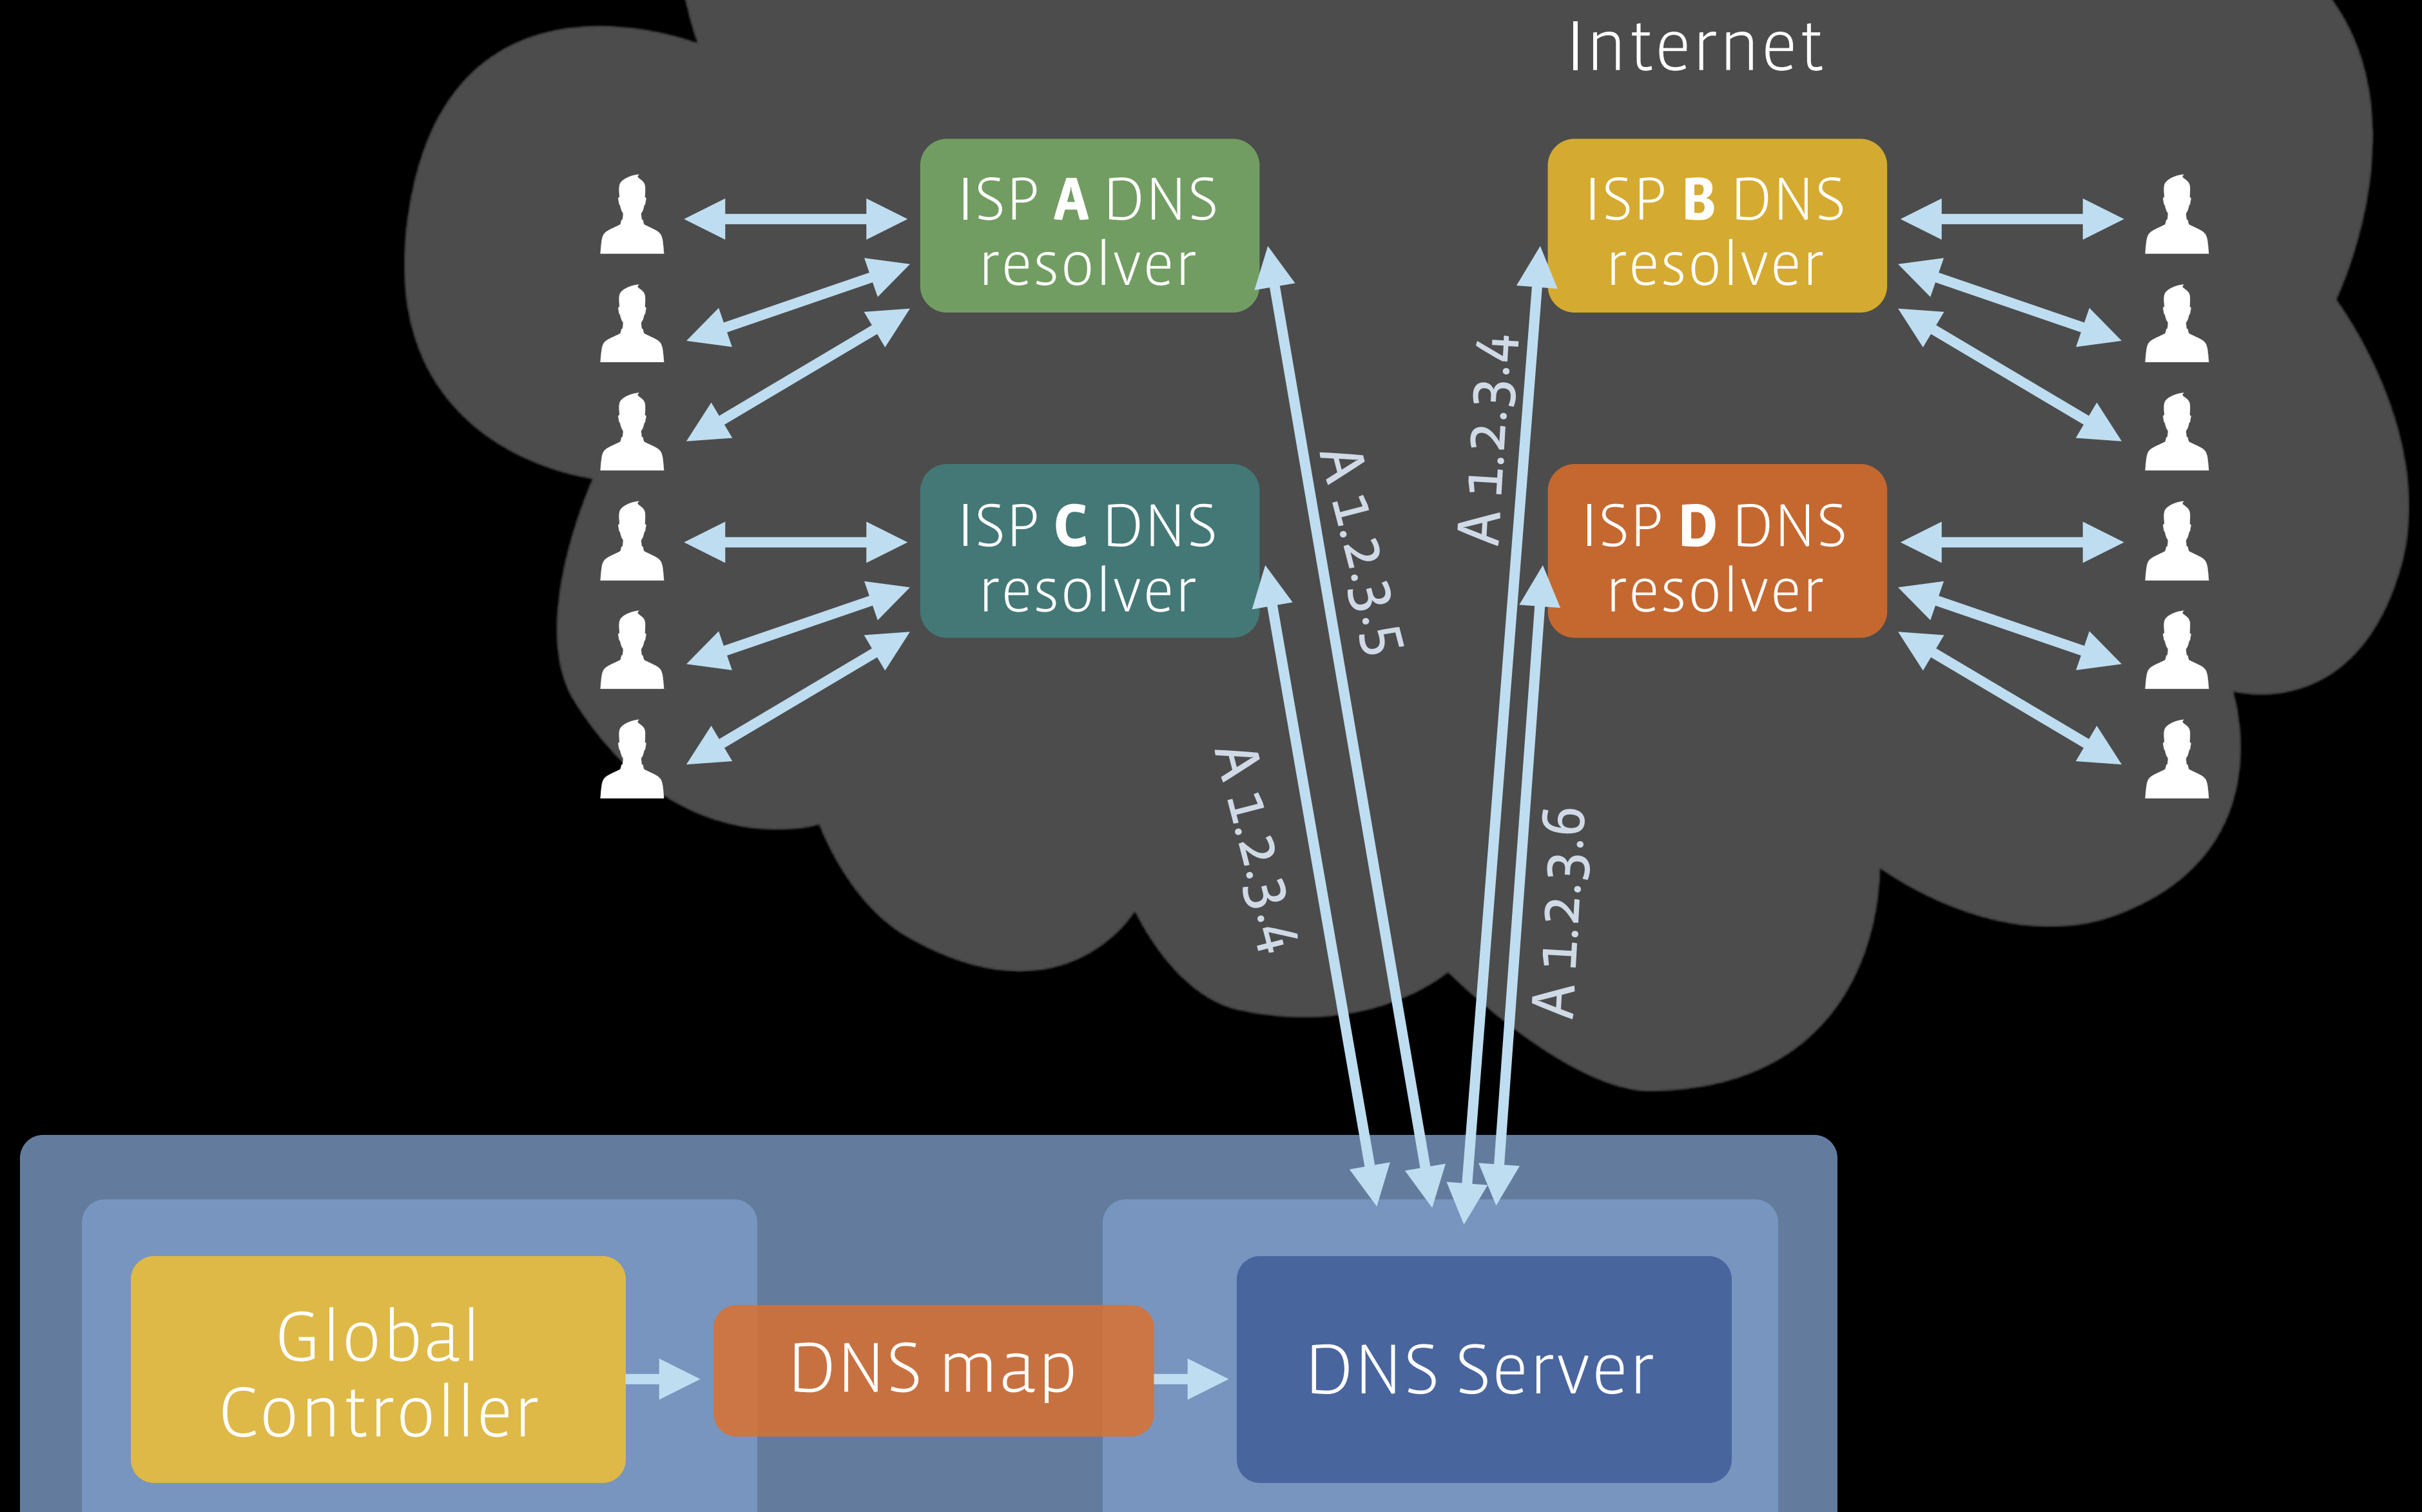
Sonar: Measuring “closeness”



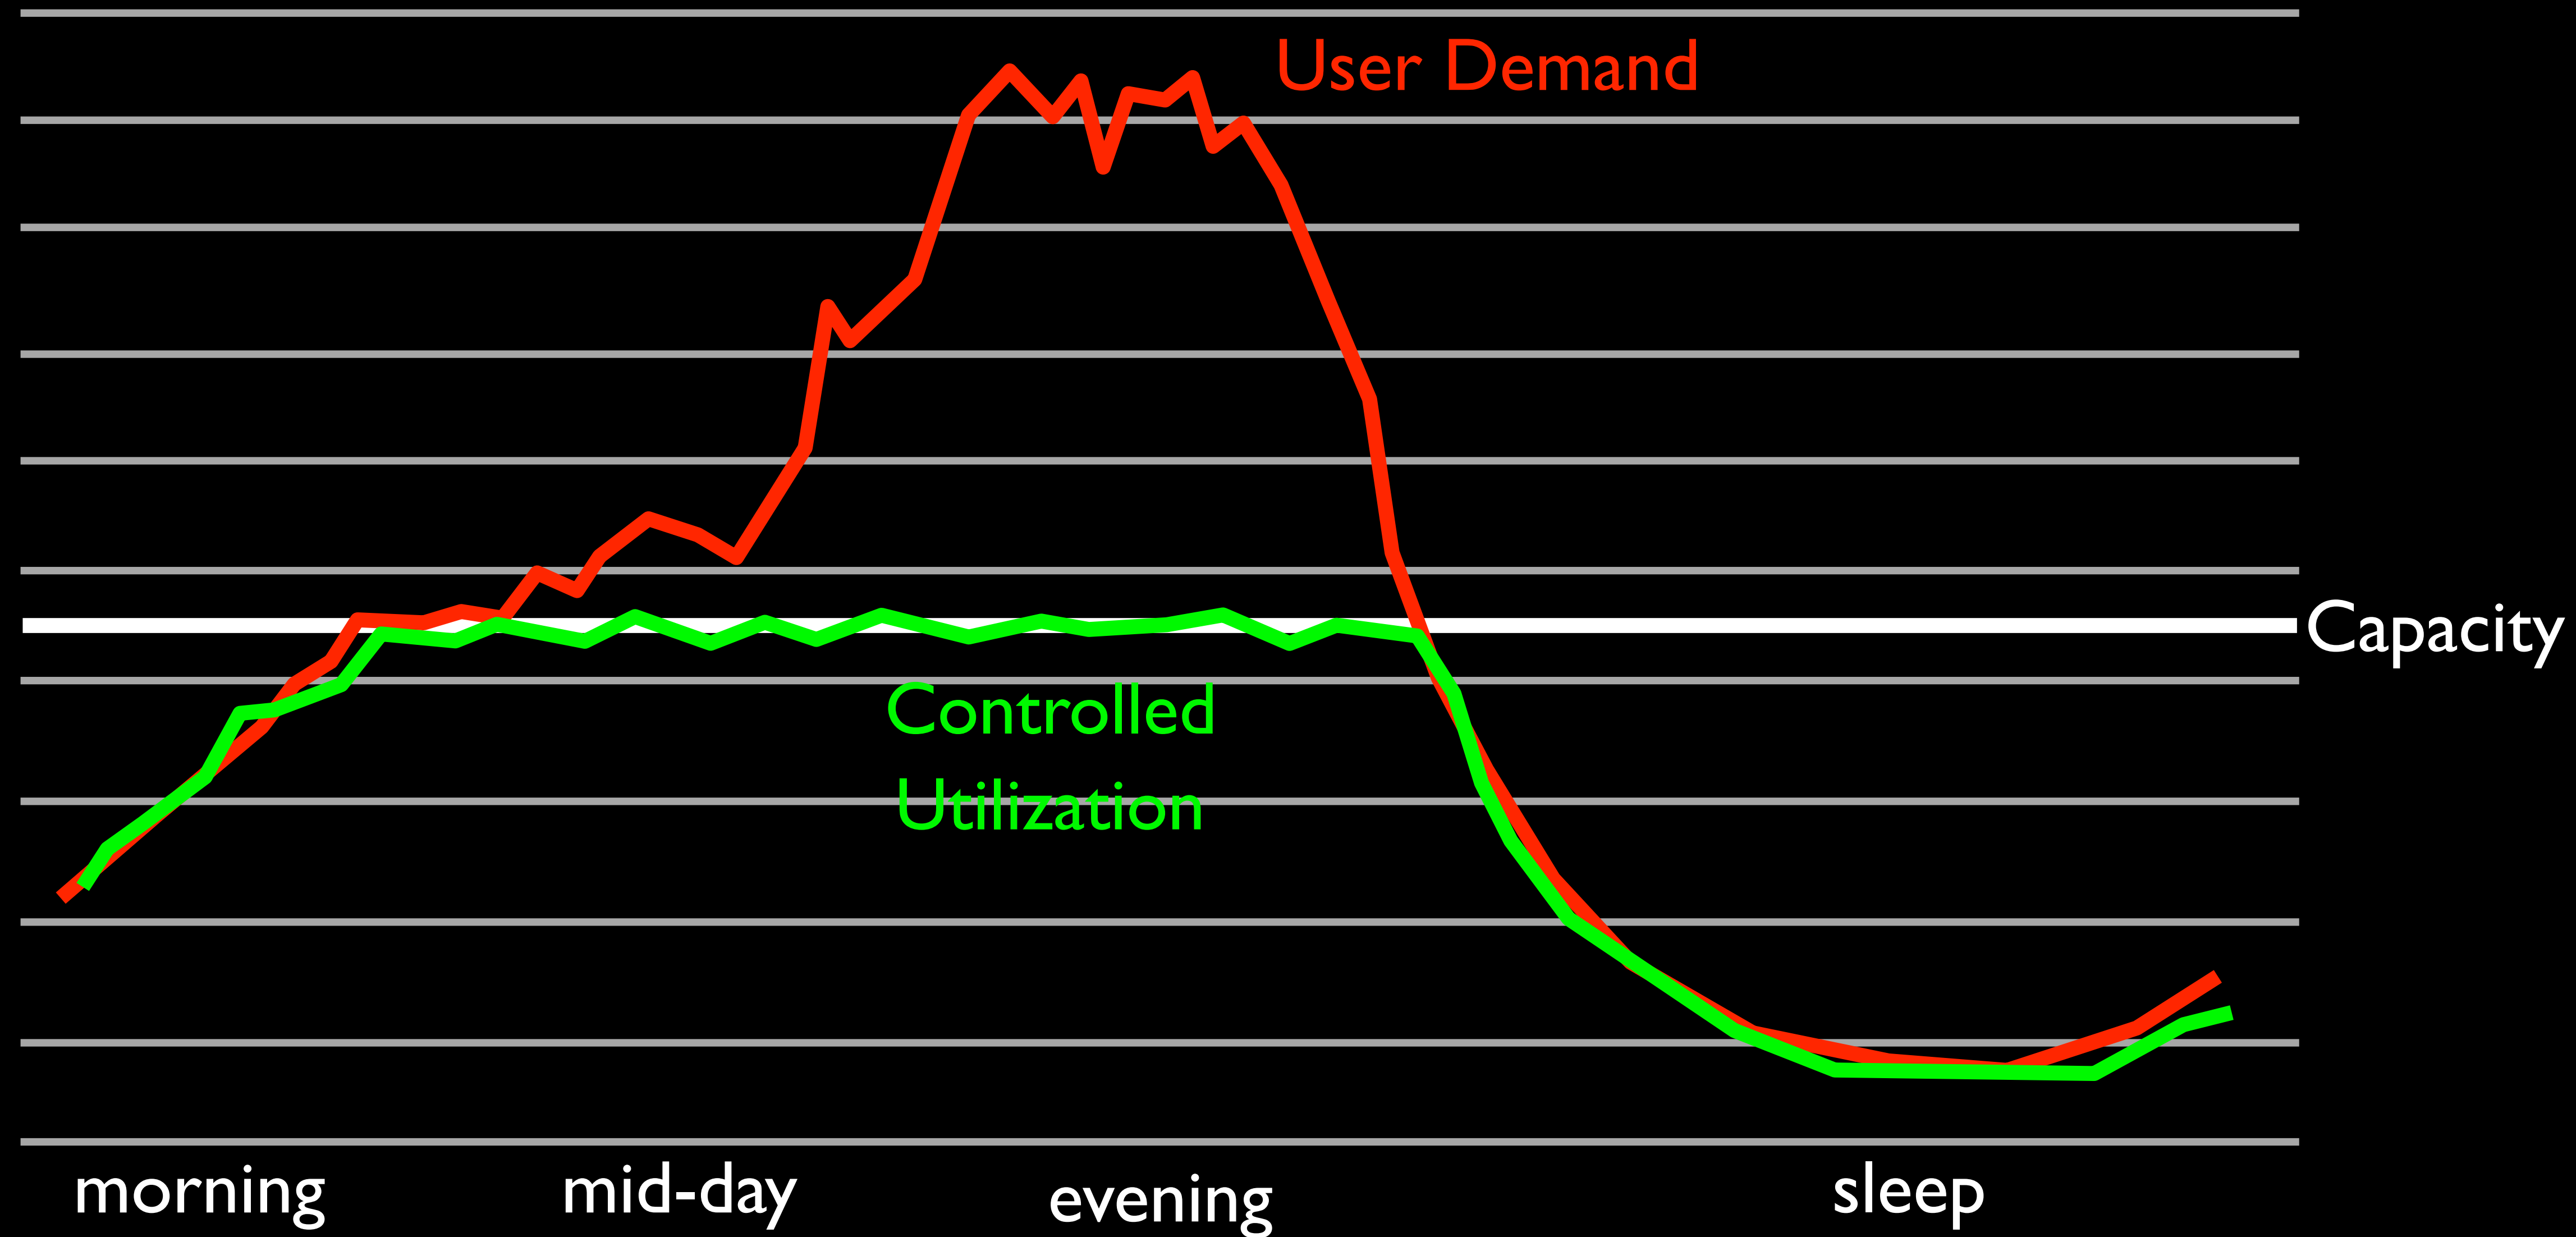
Global controller architecture



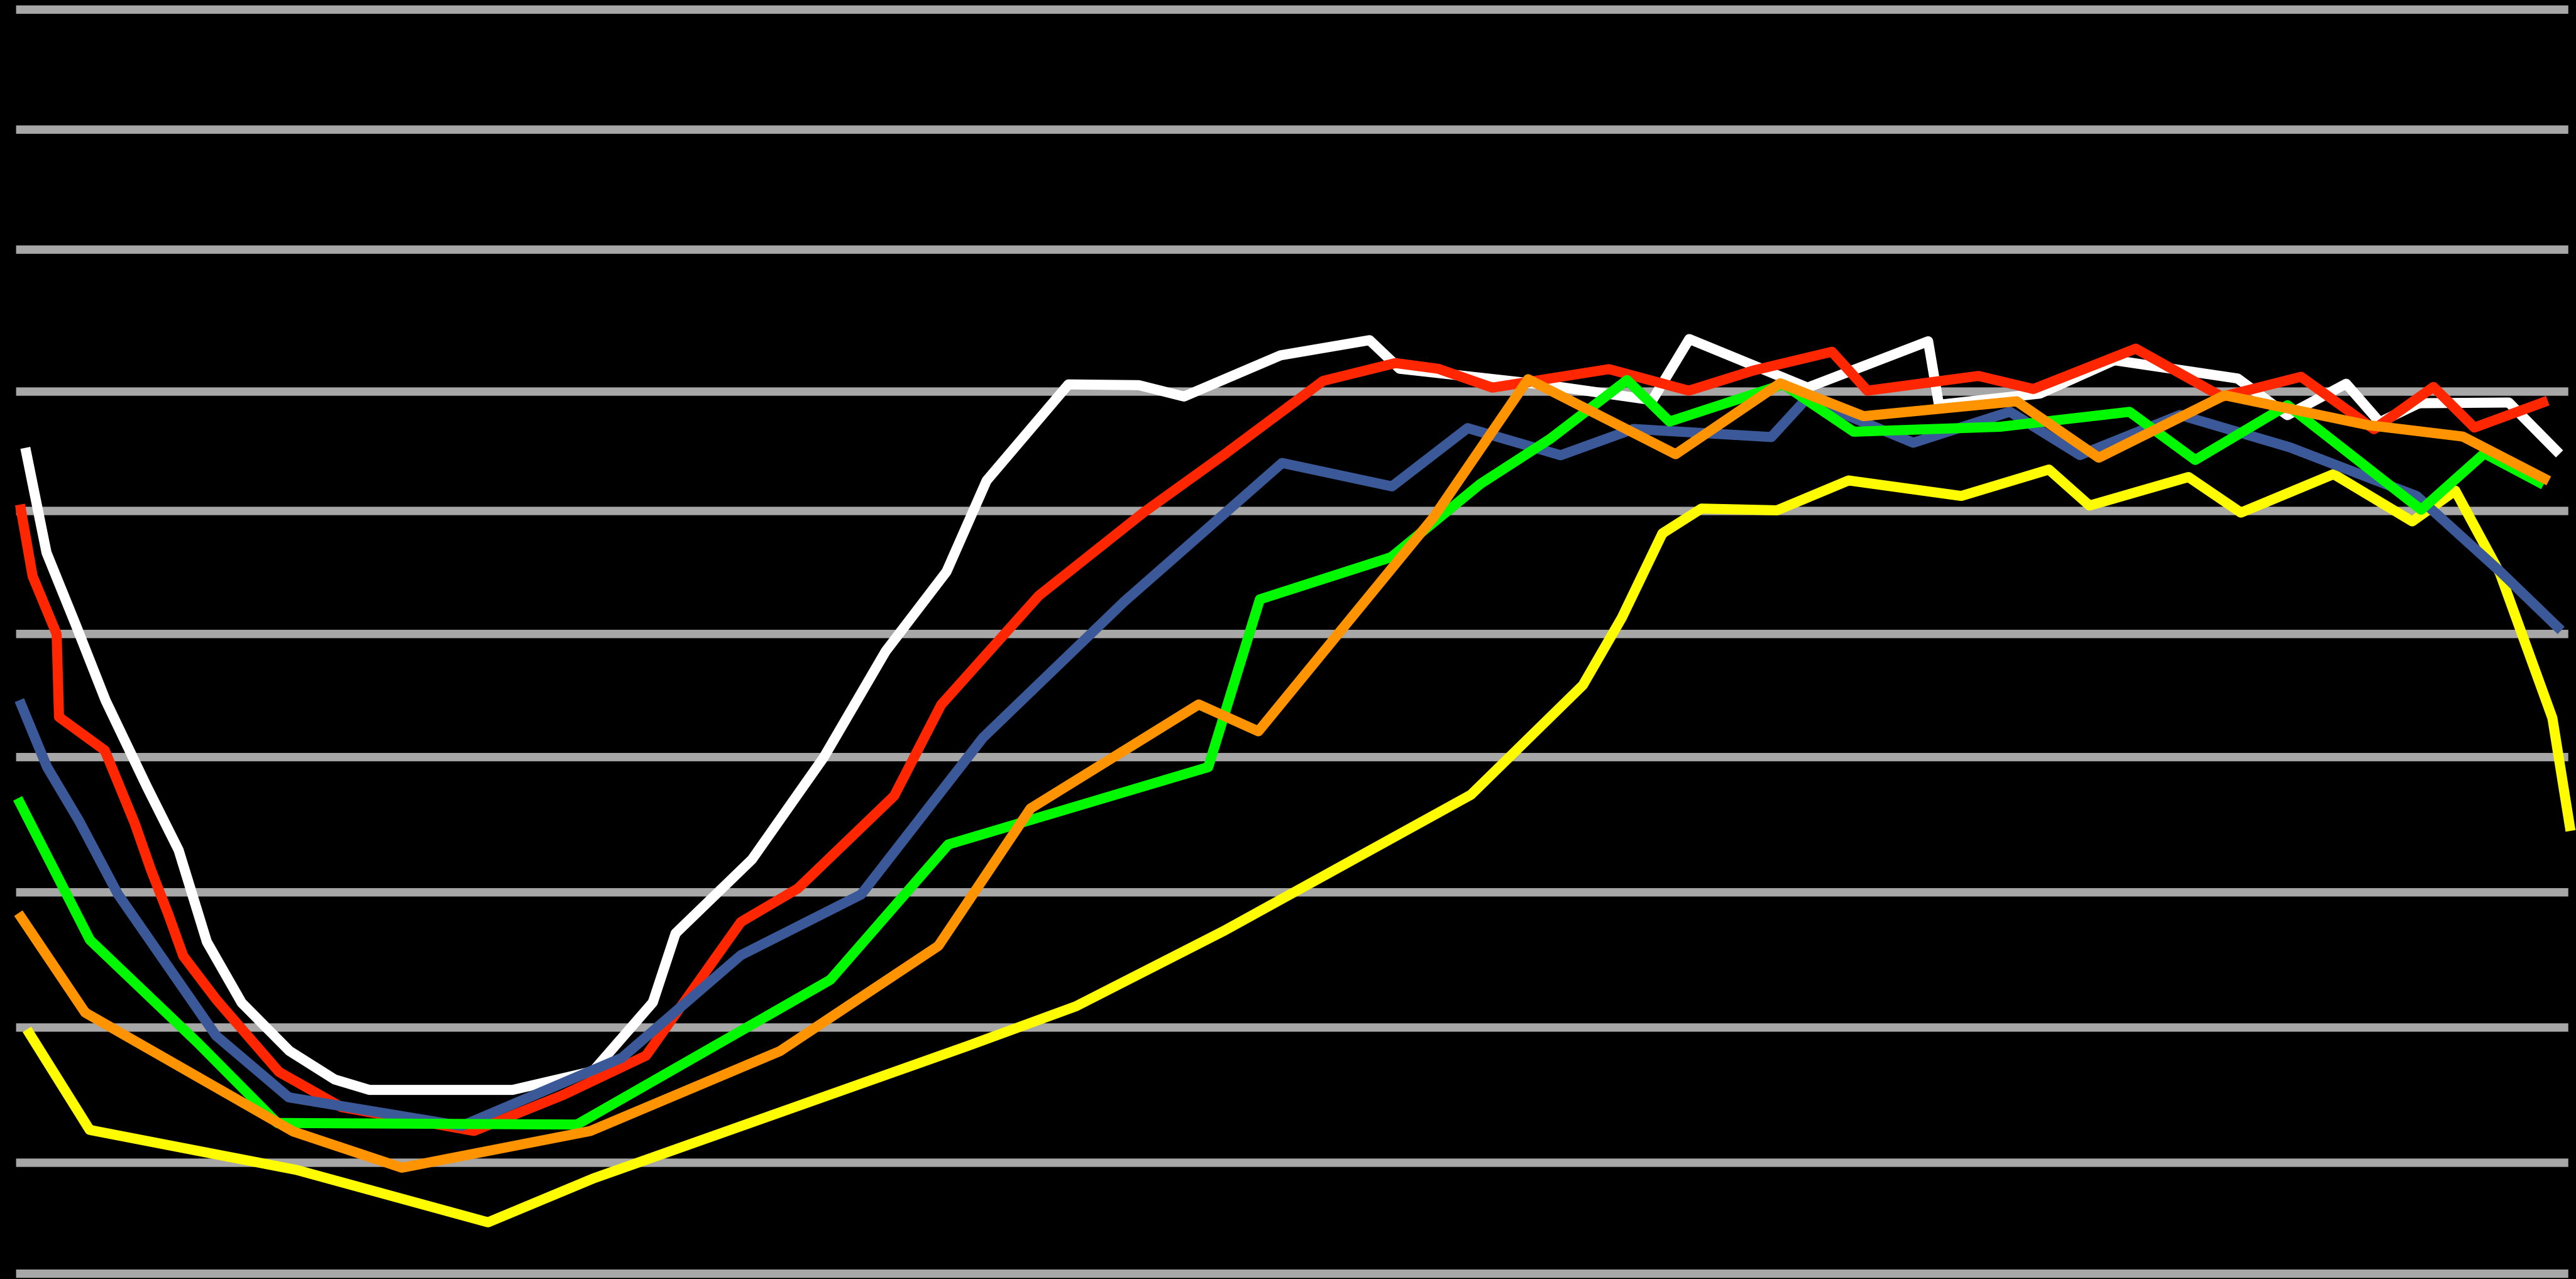
Global controller architecture




Global controller in action






Regional load shedding

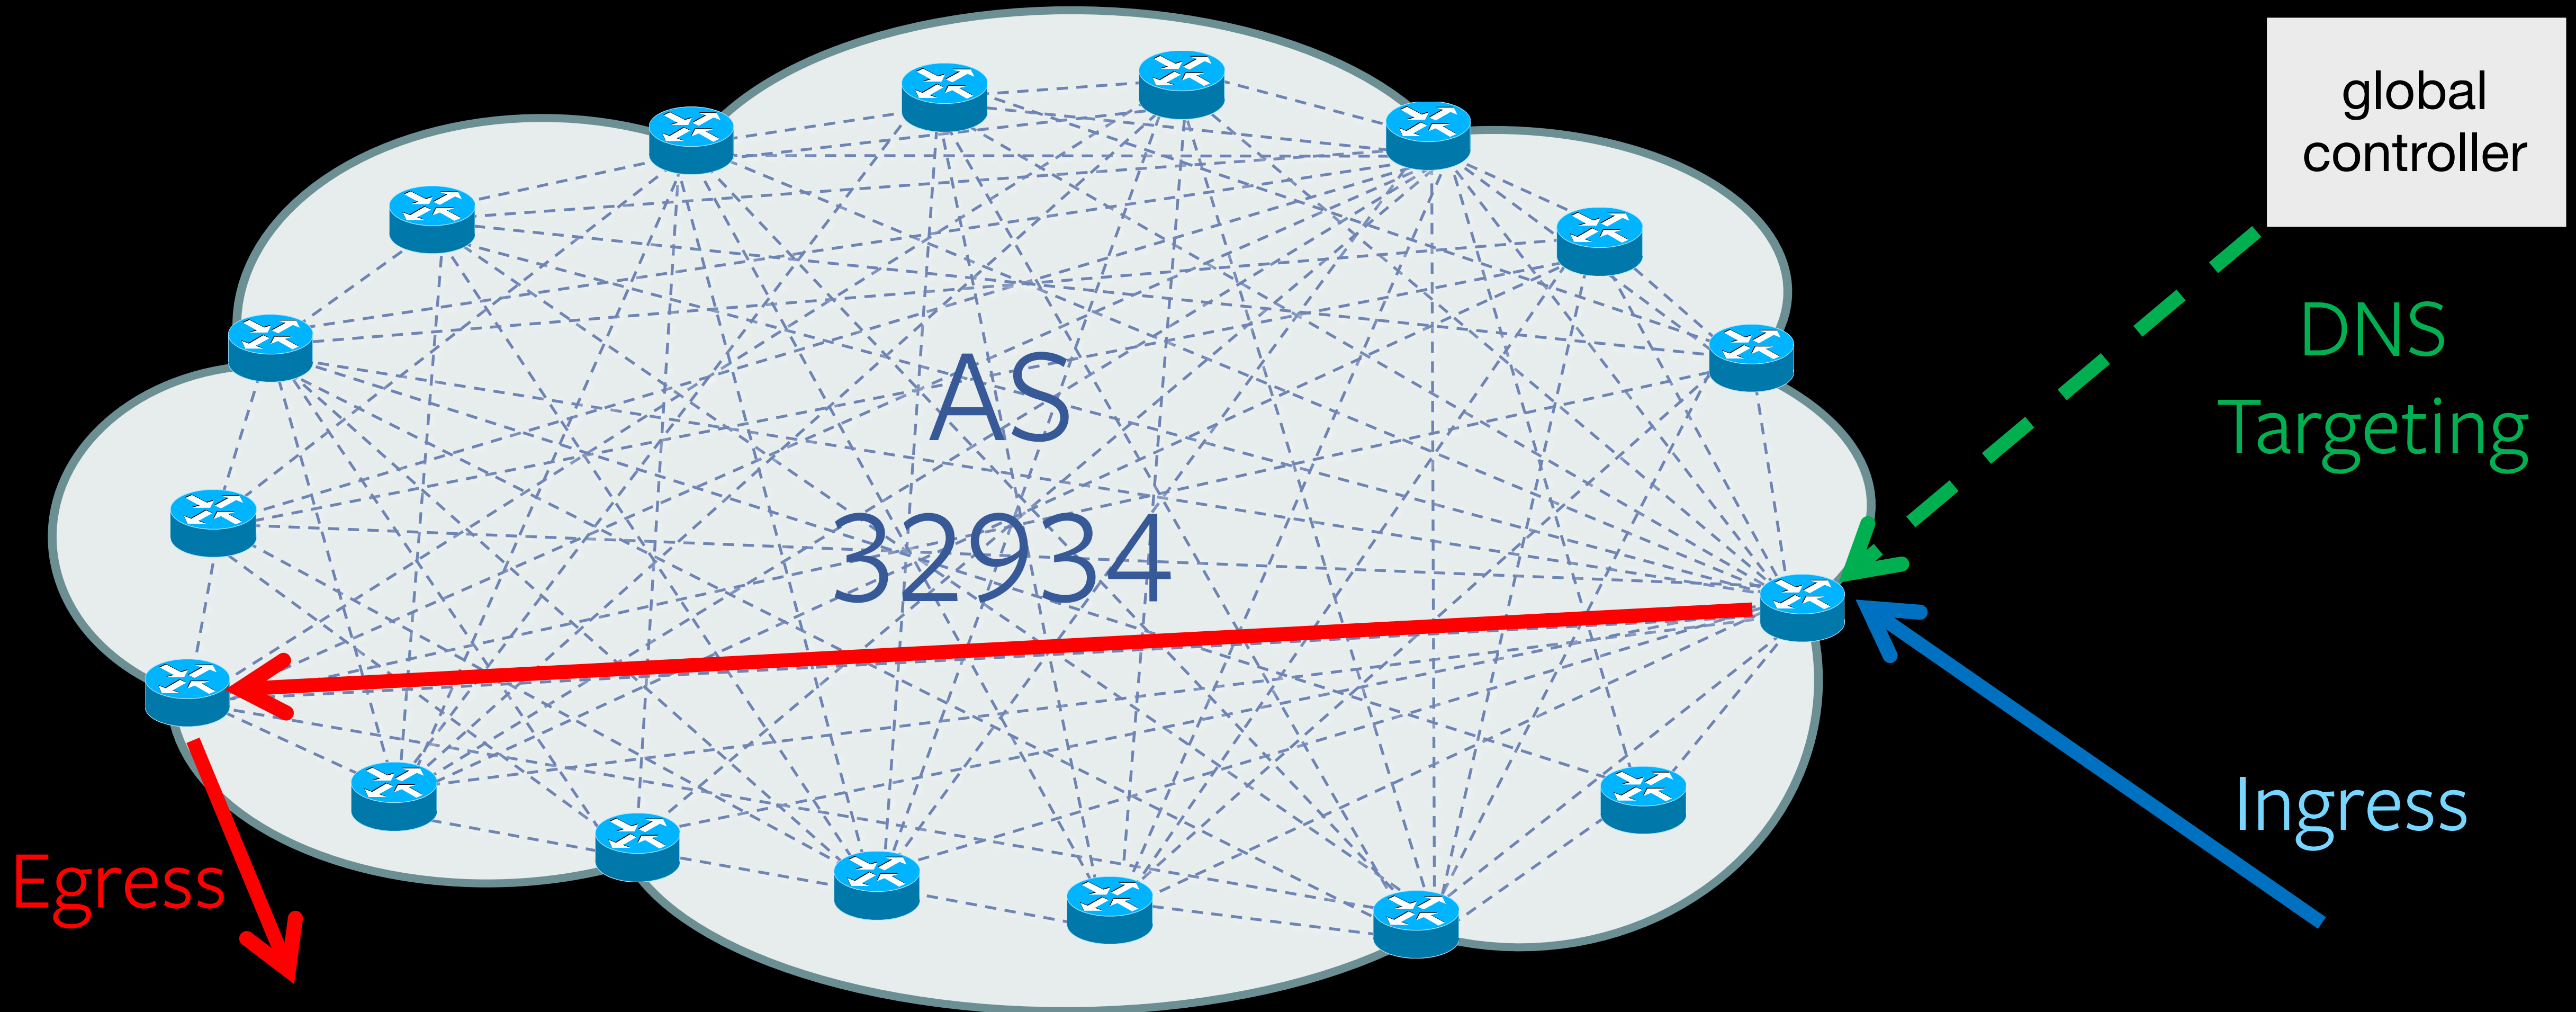


Global load shedding

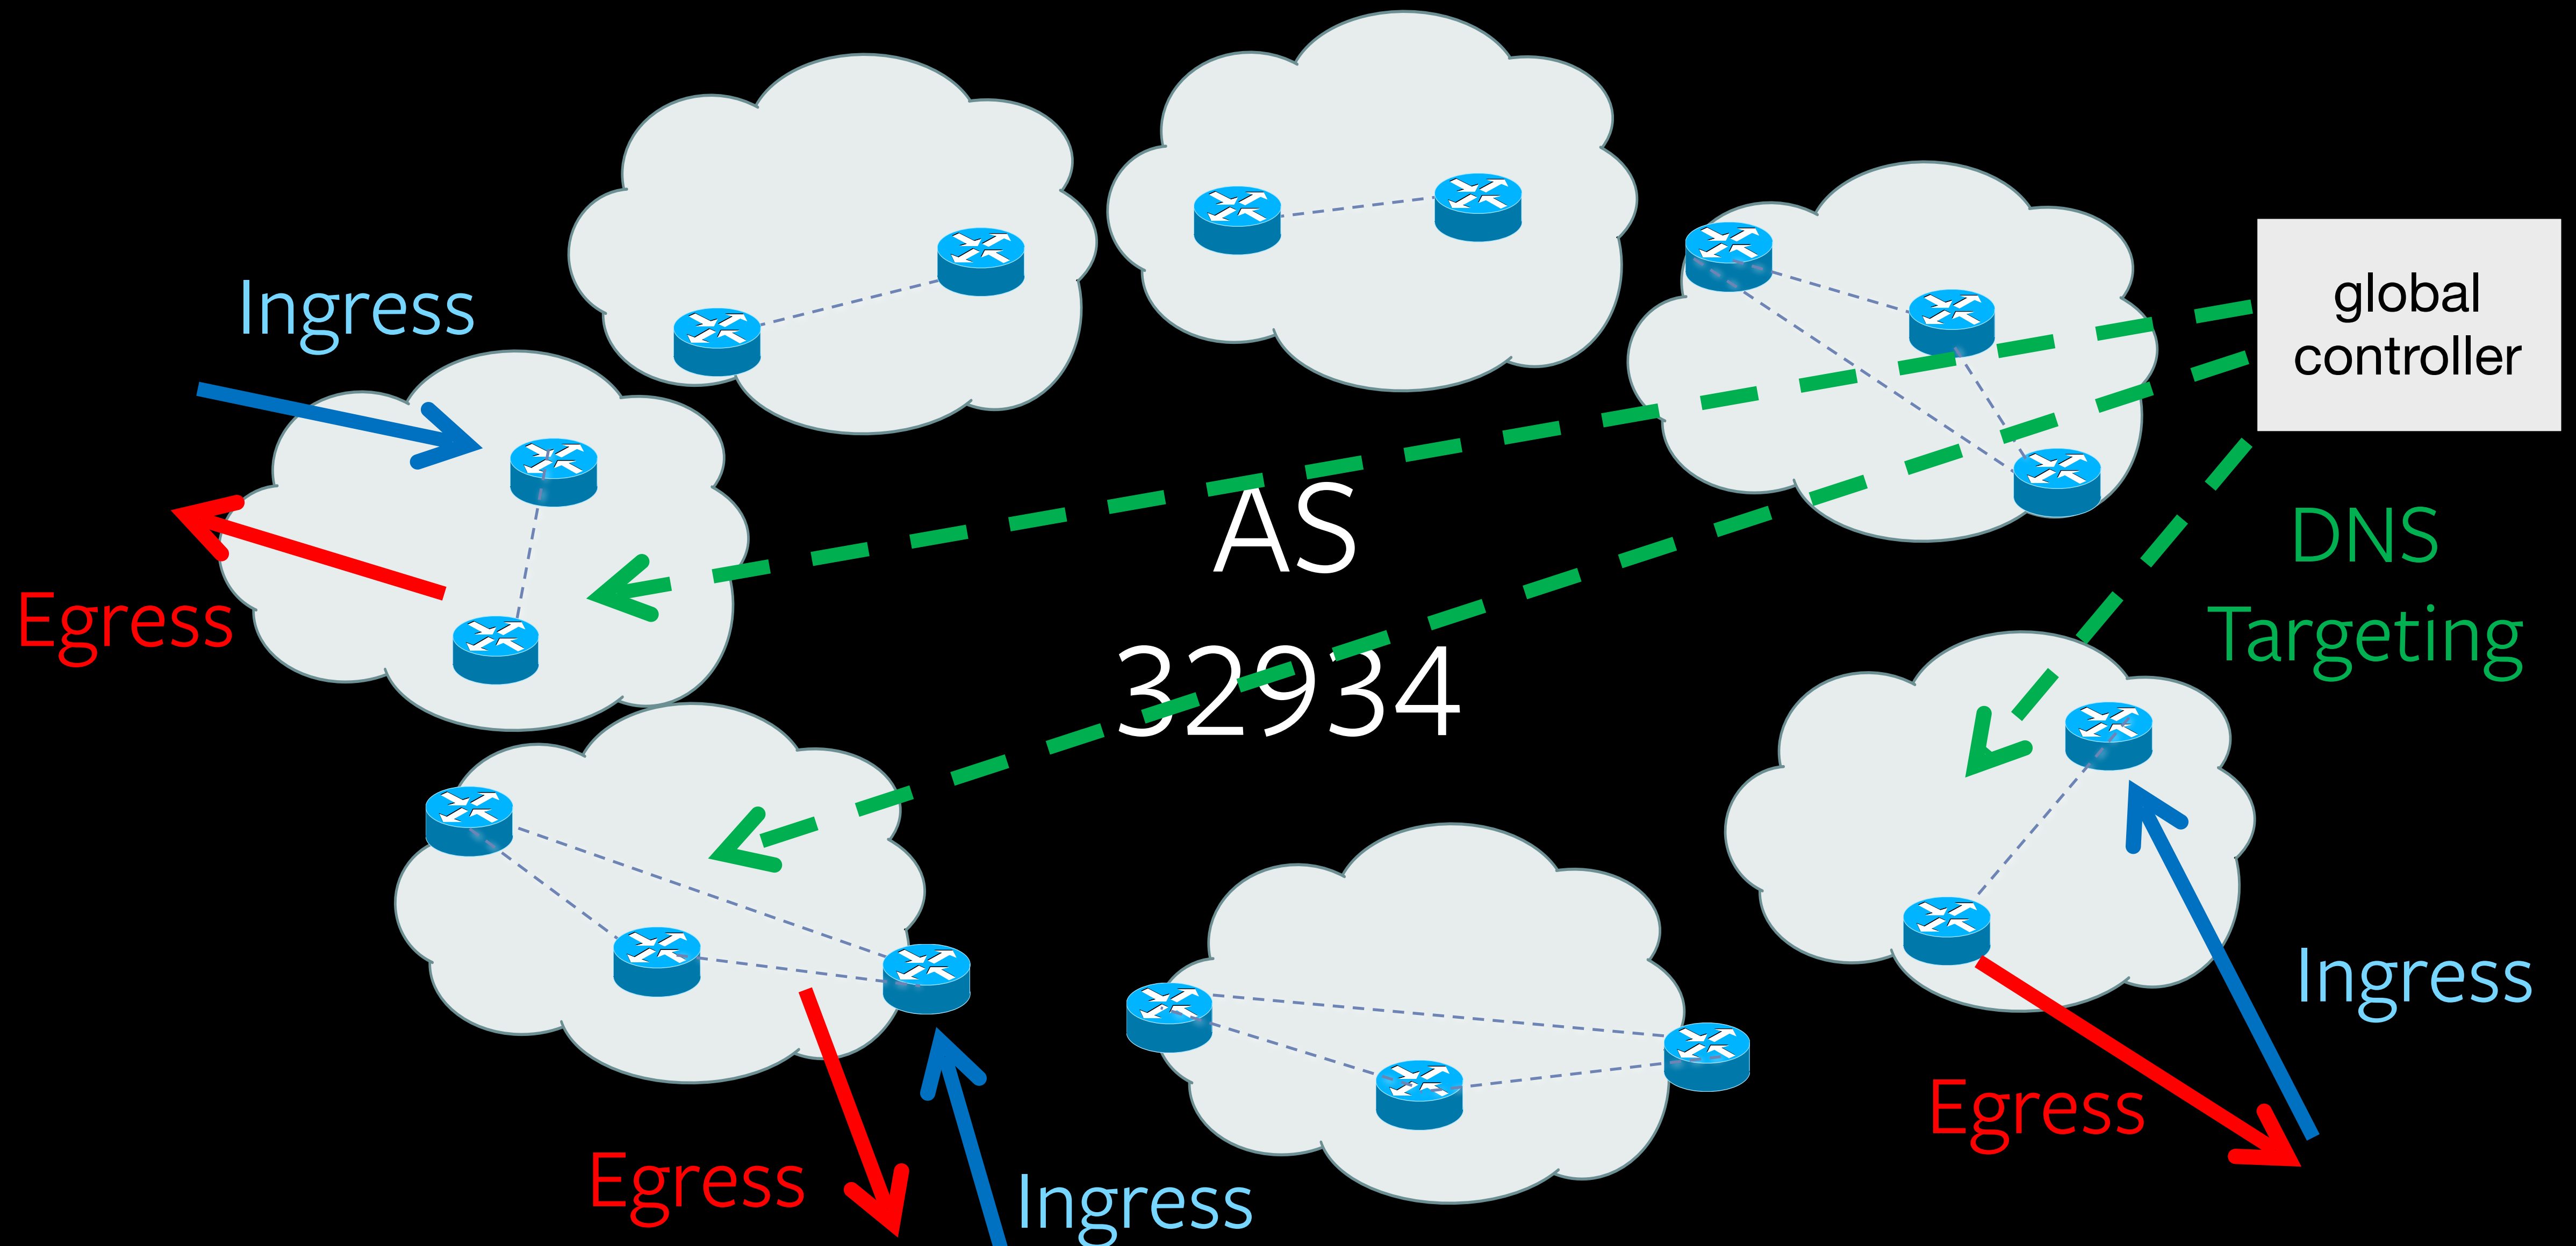




But what about BGP & targeting?



Metro-level BGP peering islands



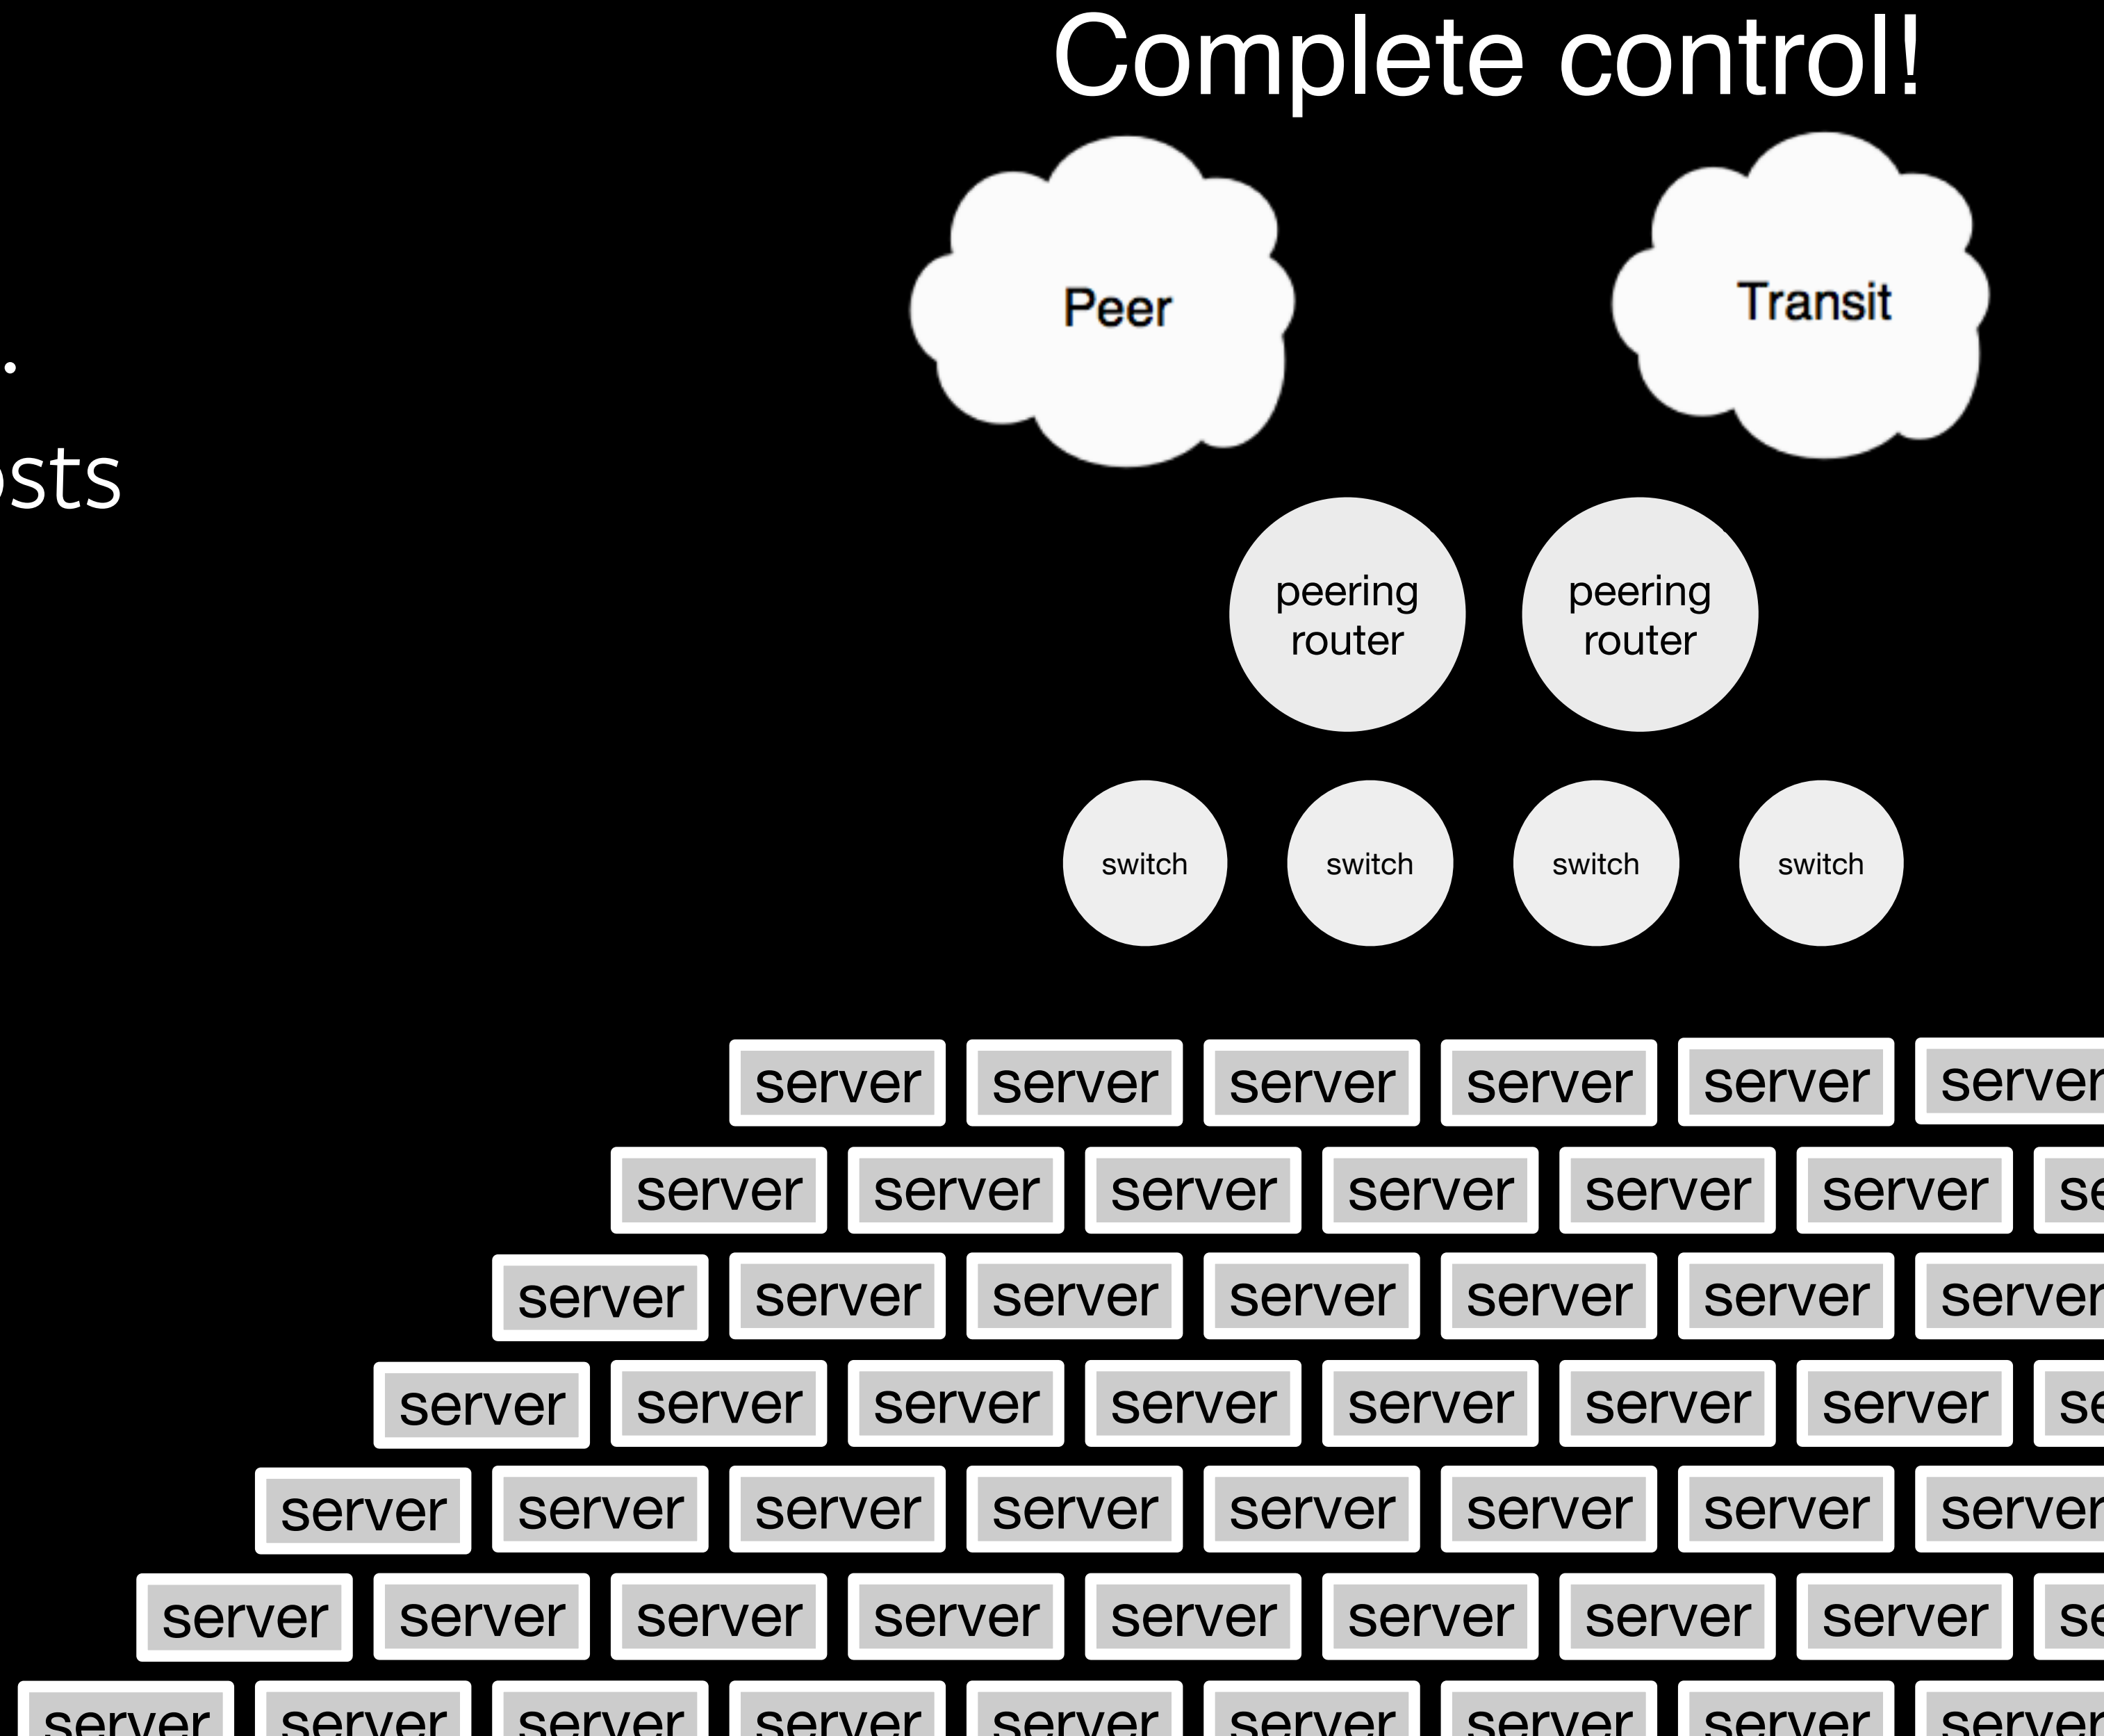
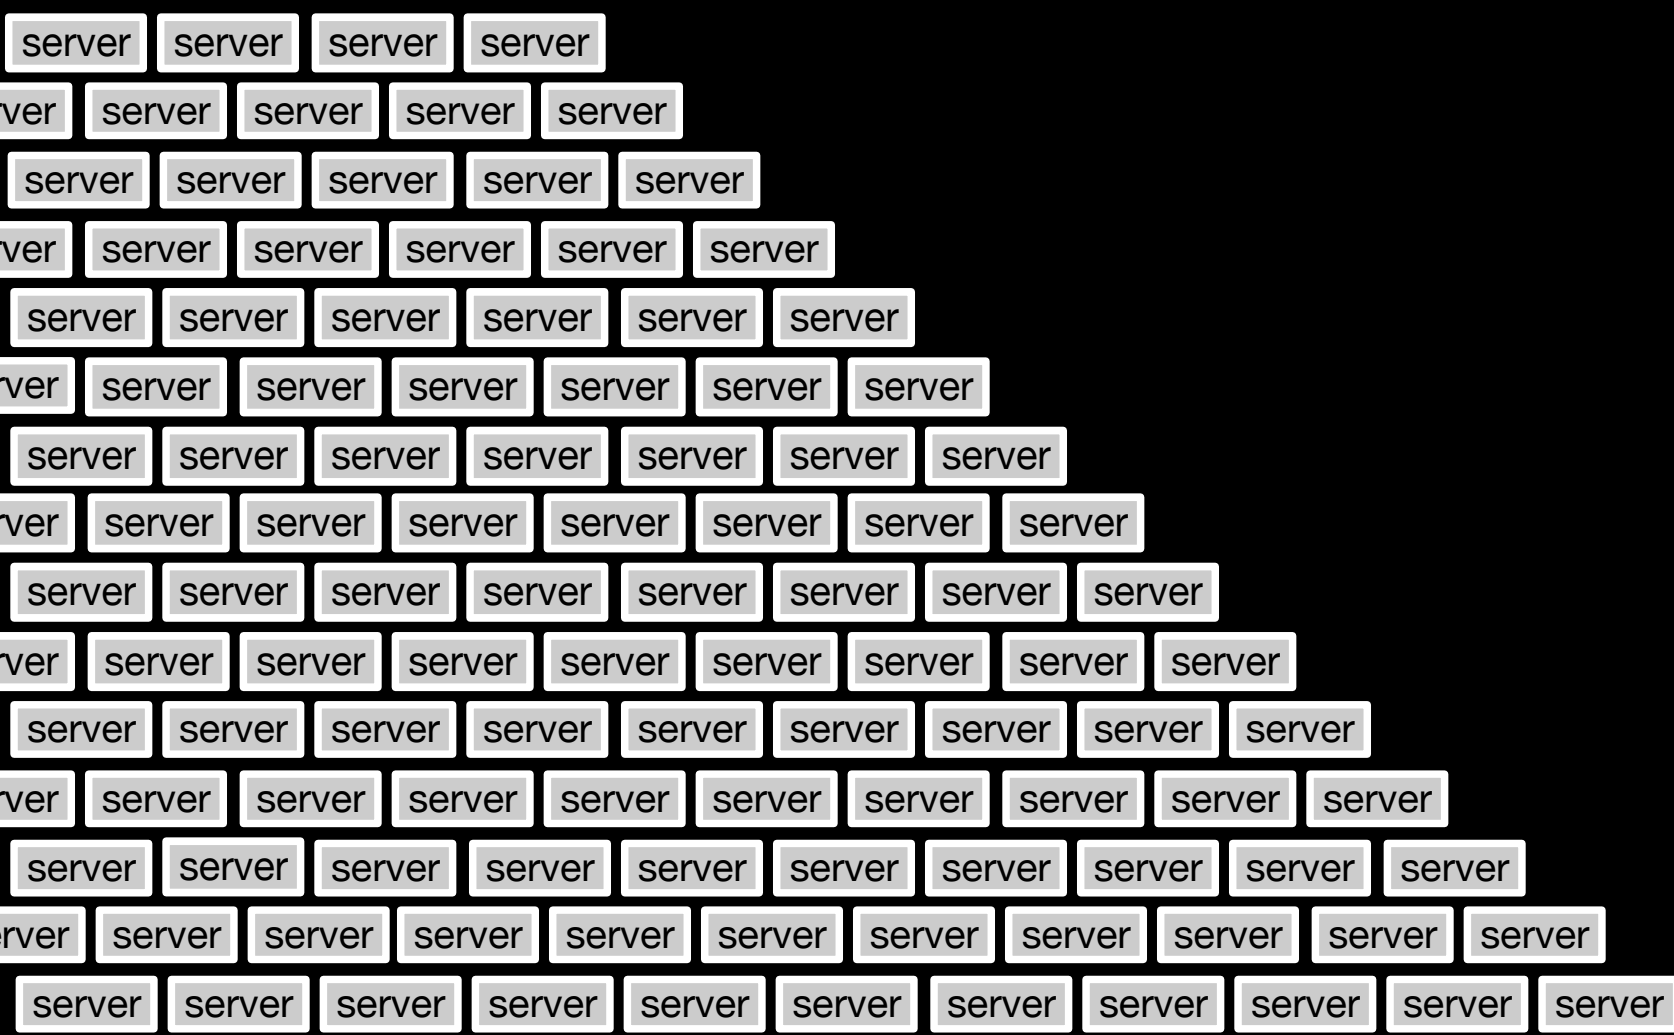
Even on an island, BGP is limited

BGP cares about...	We care about...
<ol style="list-style-type: none">1. longest prefix2. local preference3. AS path4. MED	<ol style="list-style-type: none">1. capacity2. packet loss3. latency4. <u>service performance</u>

Local network controllers

Host routing!

- Looks great in slides
- Not so simple in practice...
 - Synchronization across hosts
 - *A lot* of hosts
 - Non-local traffic sources
 - And which encapsulation?

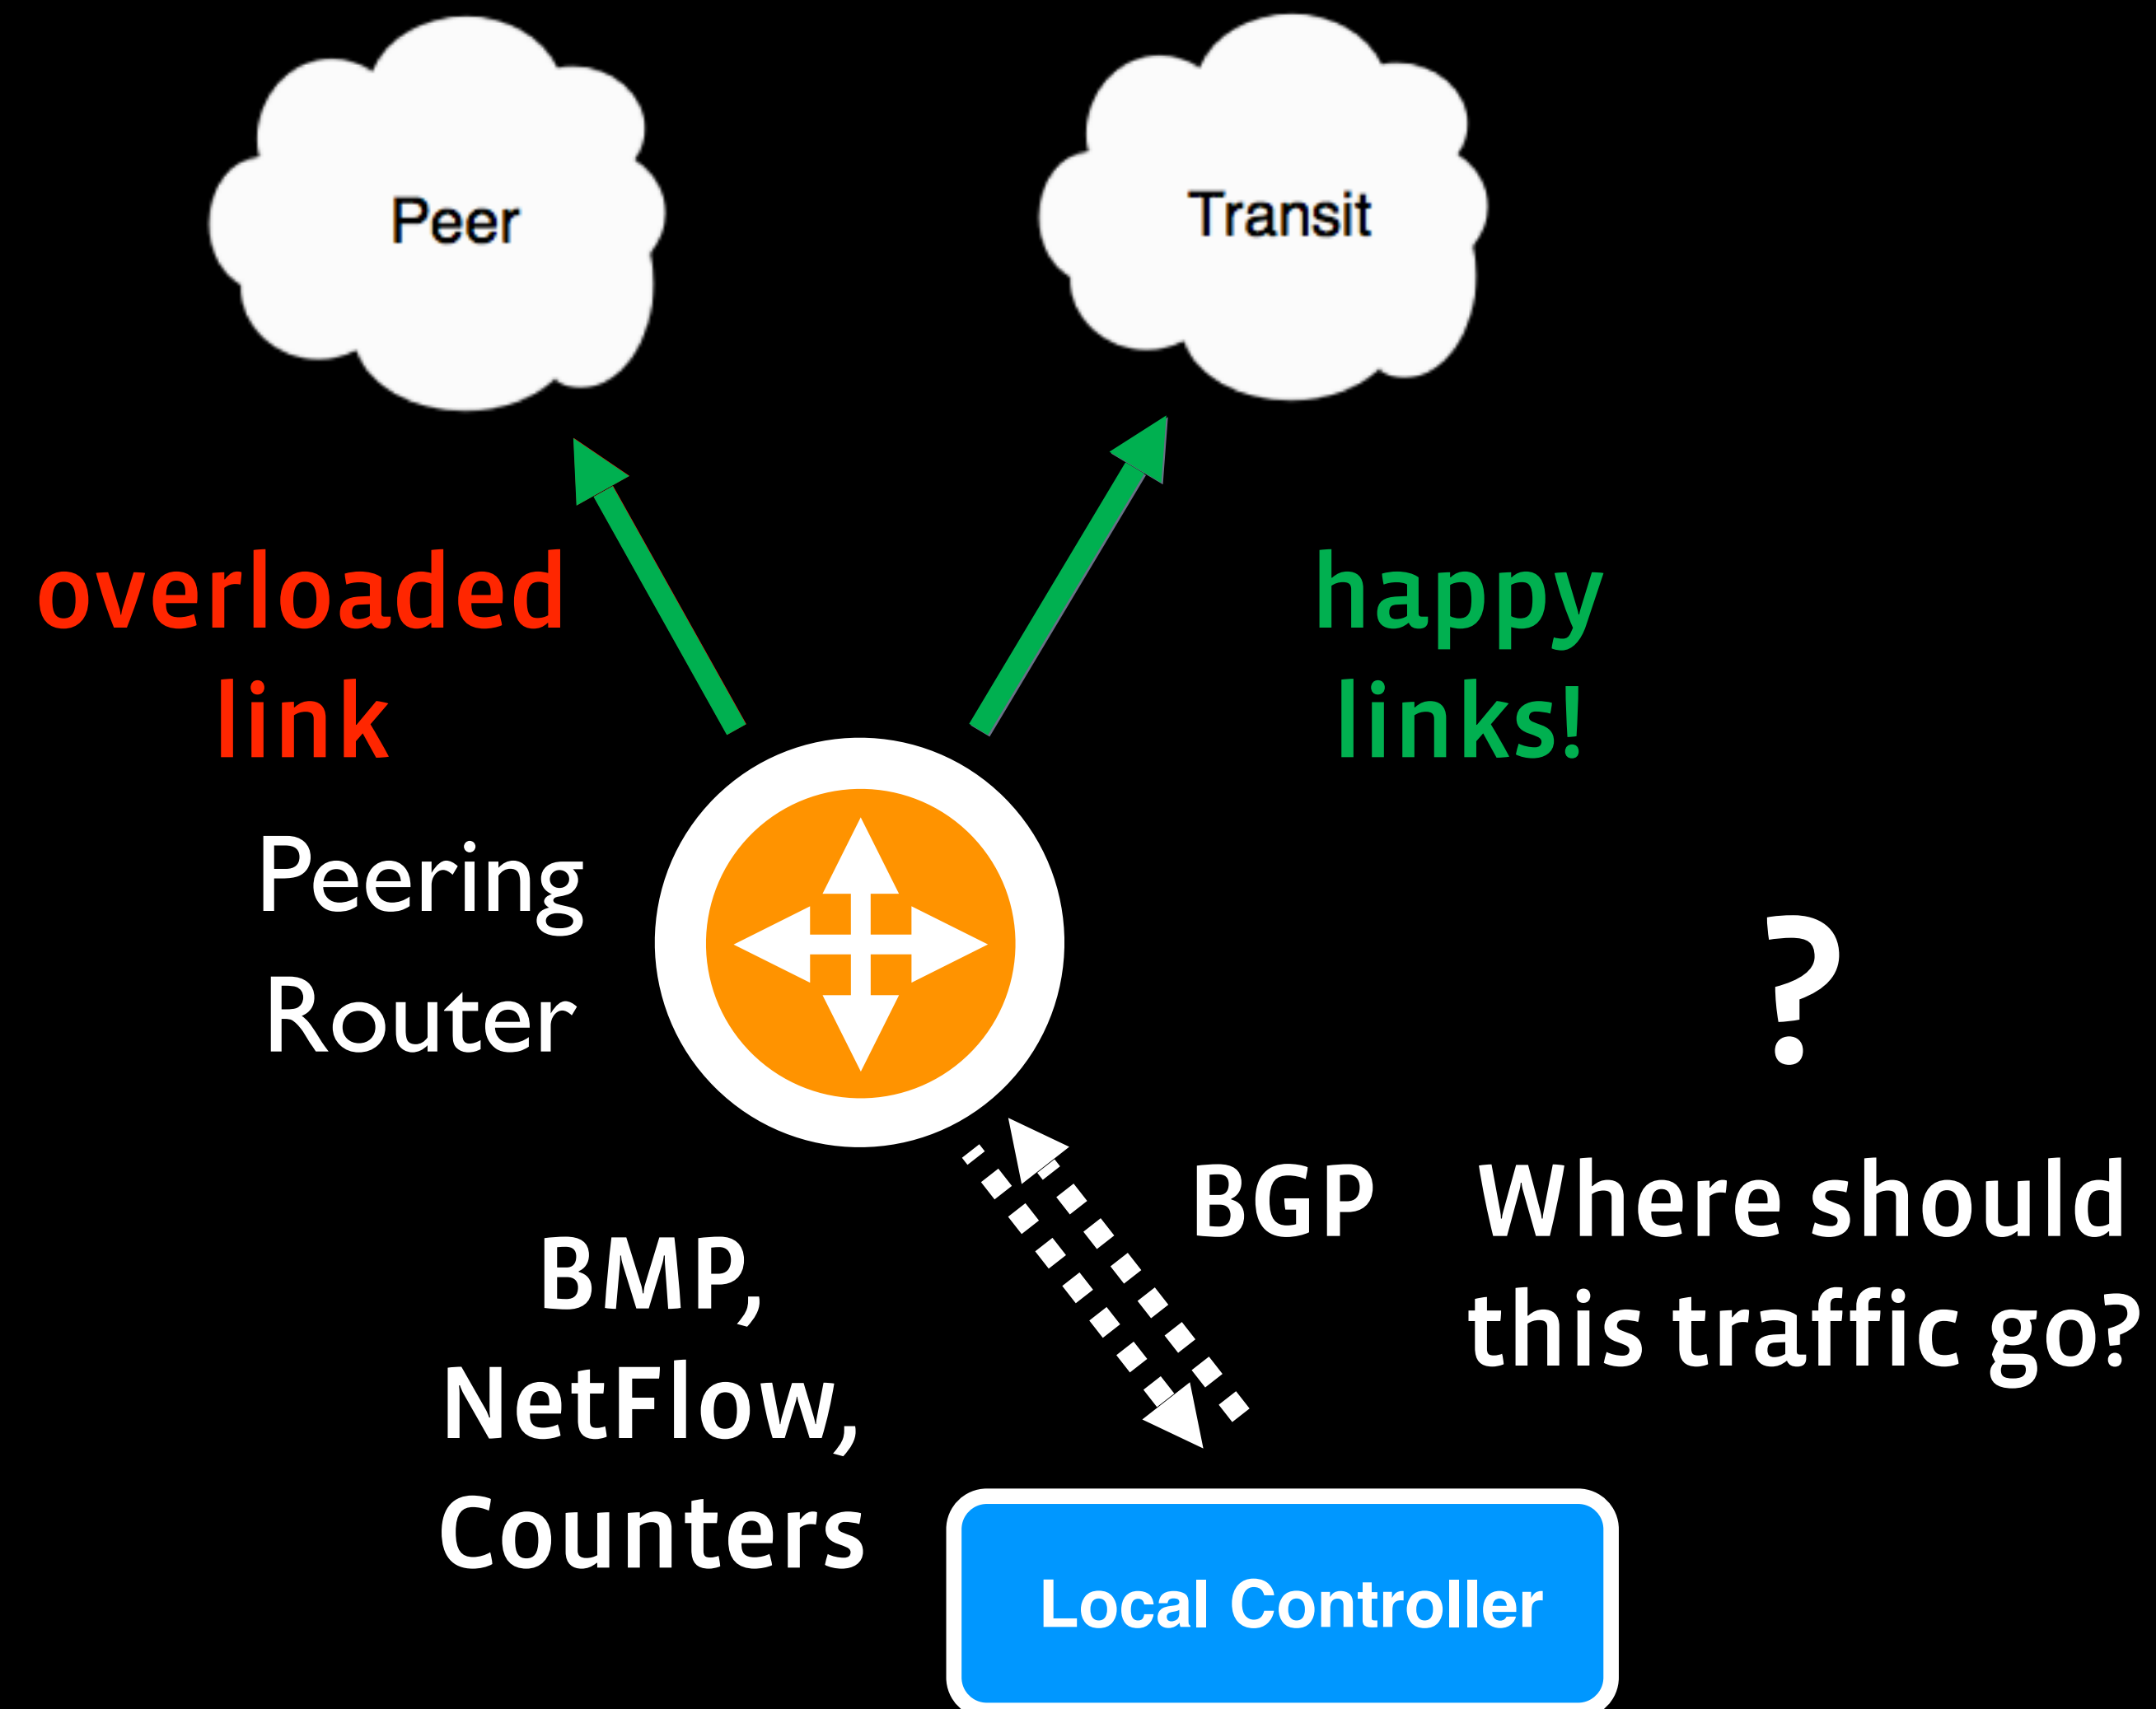


How do the servers encapsulate!?

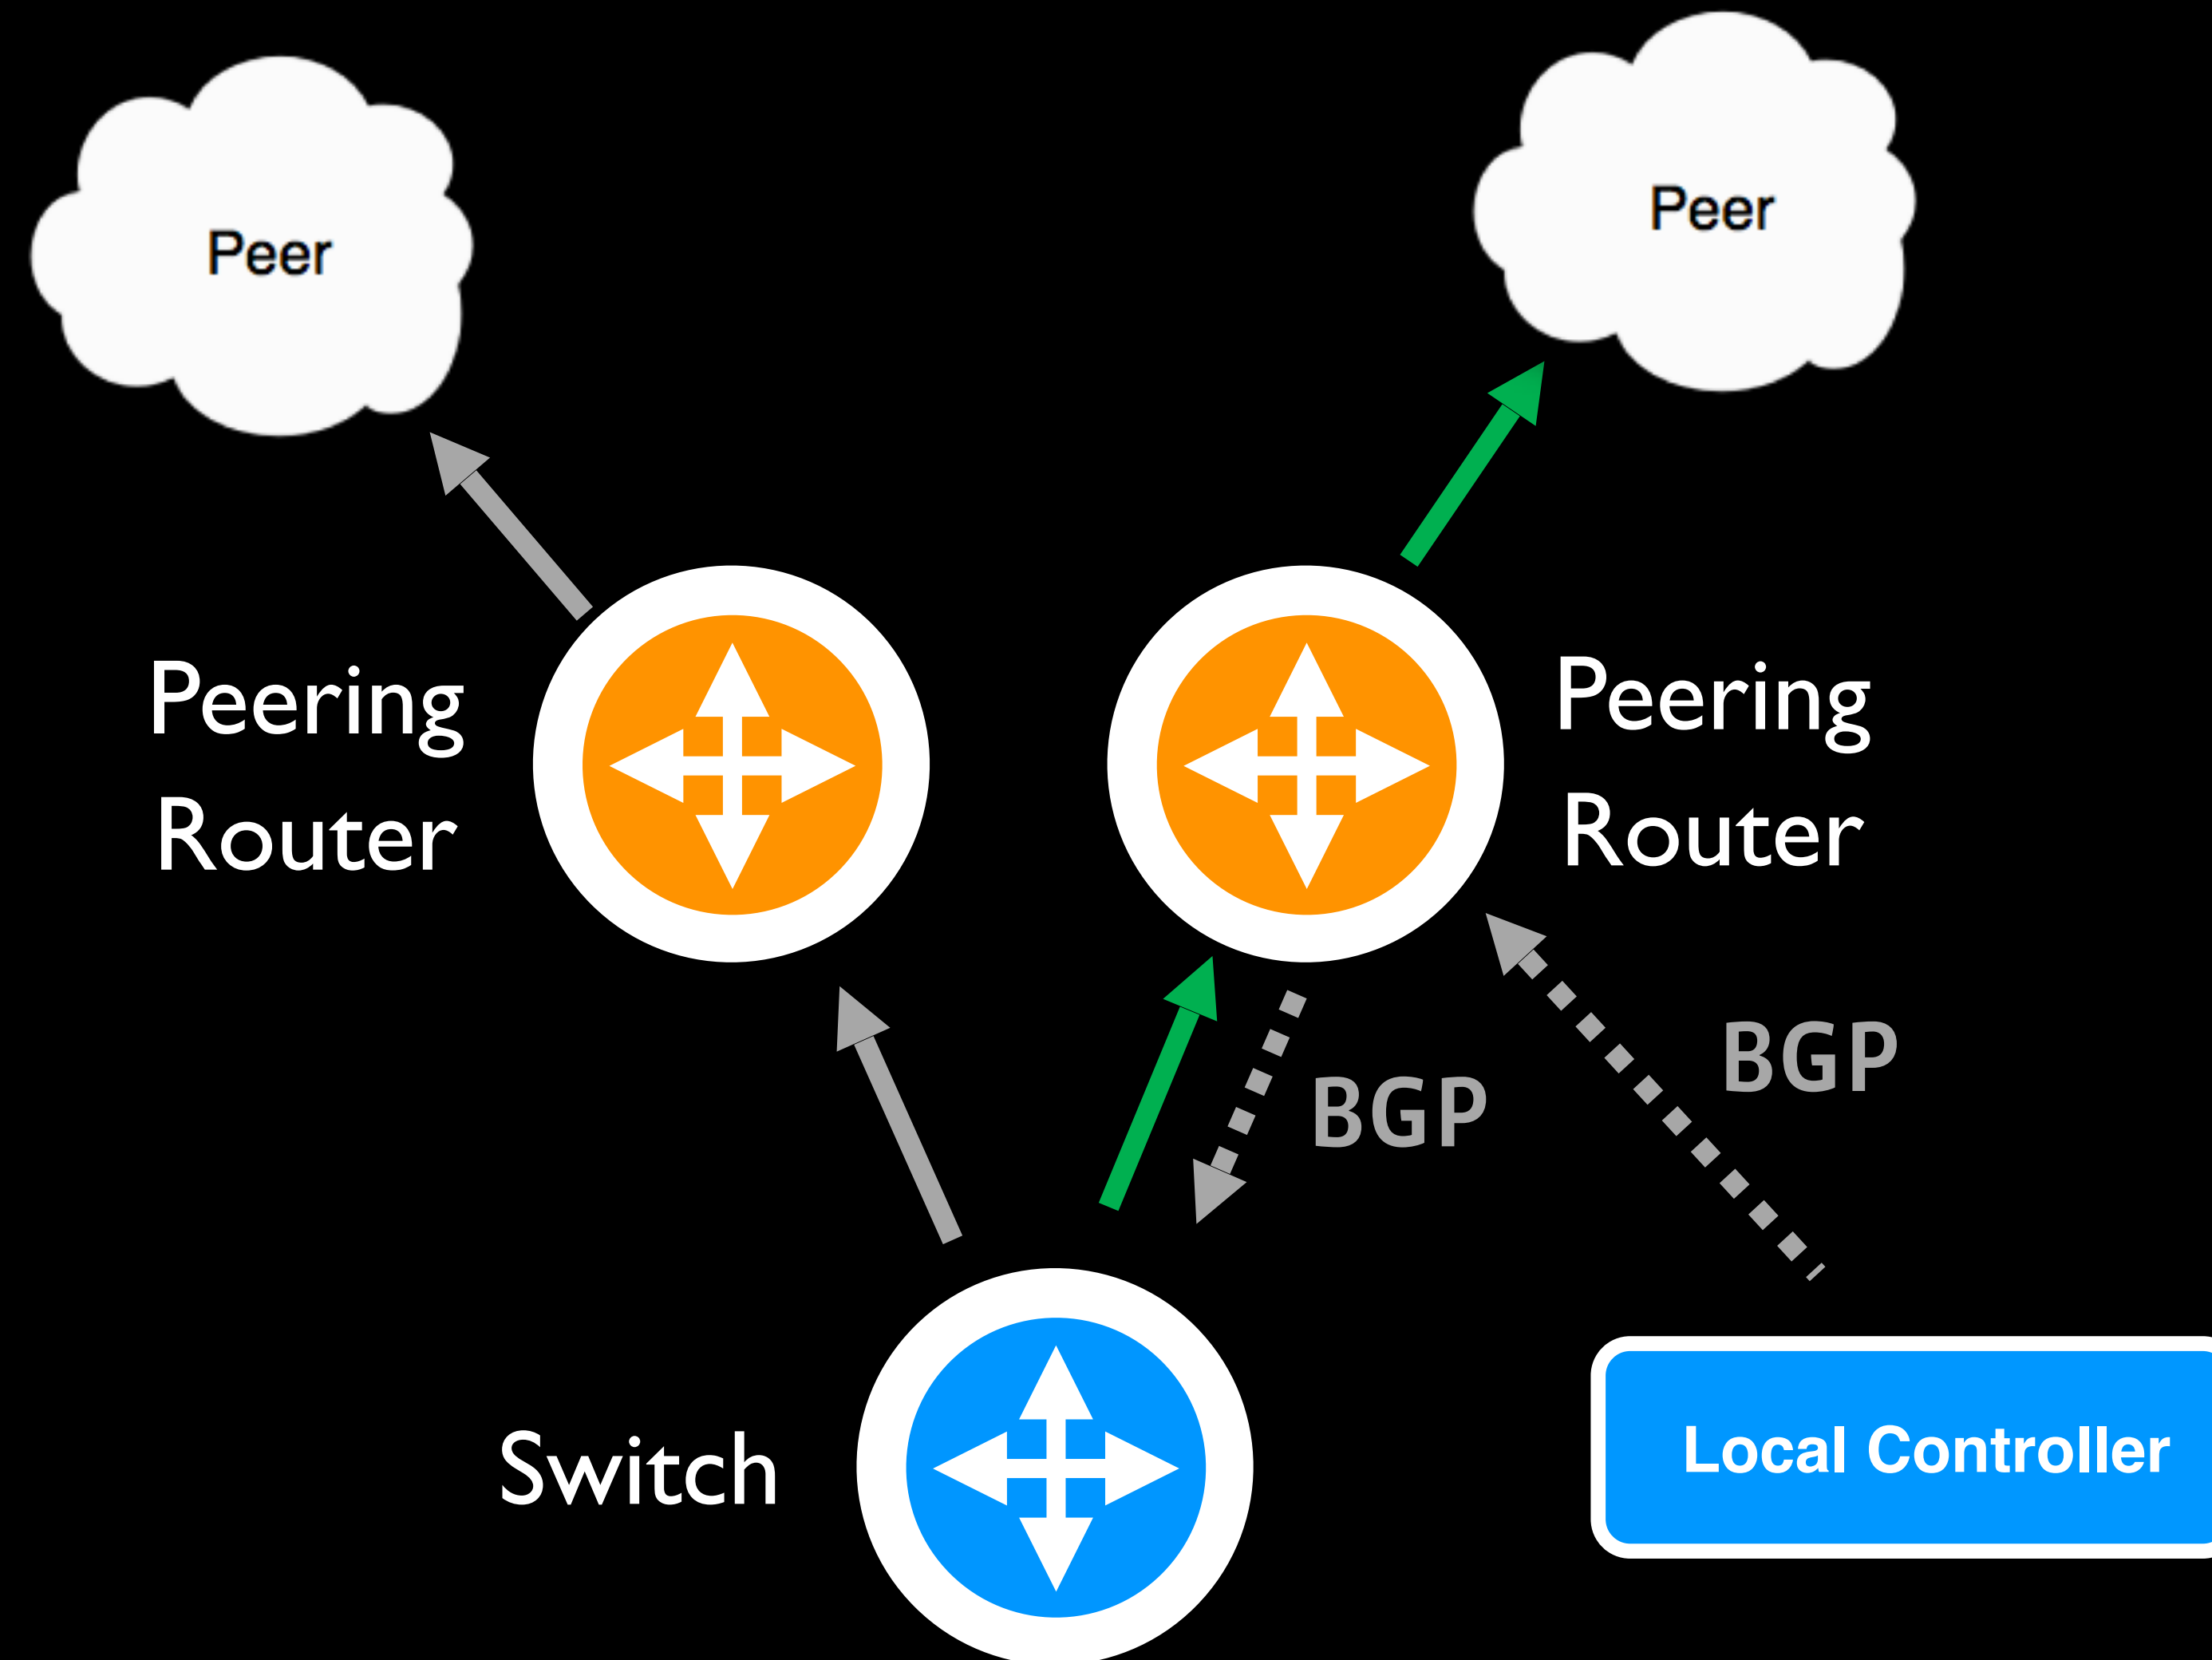
- MPLS – Kernel support problems, router features
- PBR - limits platform diversity, router configuration state
 - DSCP – very coarse range
 - GRE key – repeated router bugs

What we learned... keep it simple

Evolving beyond BGP ...with BGP

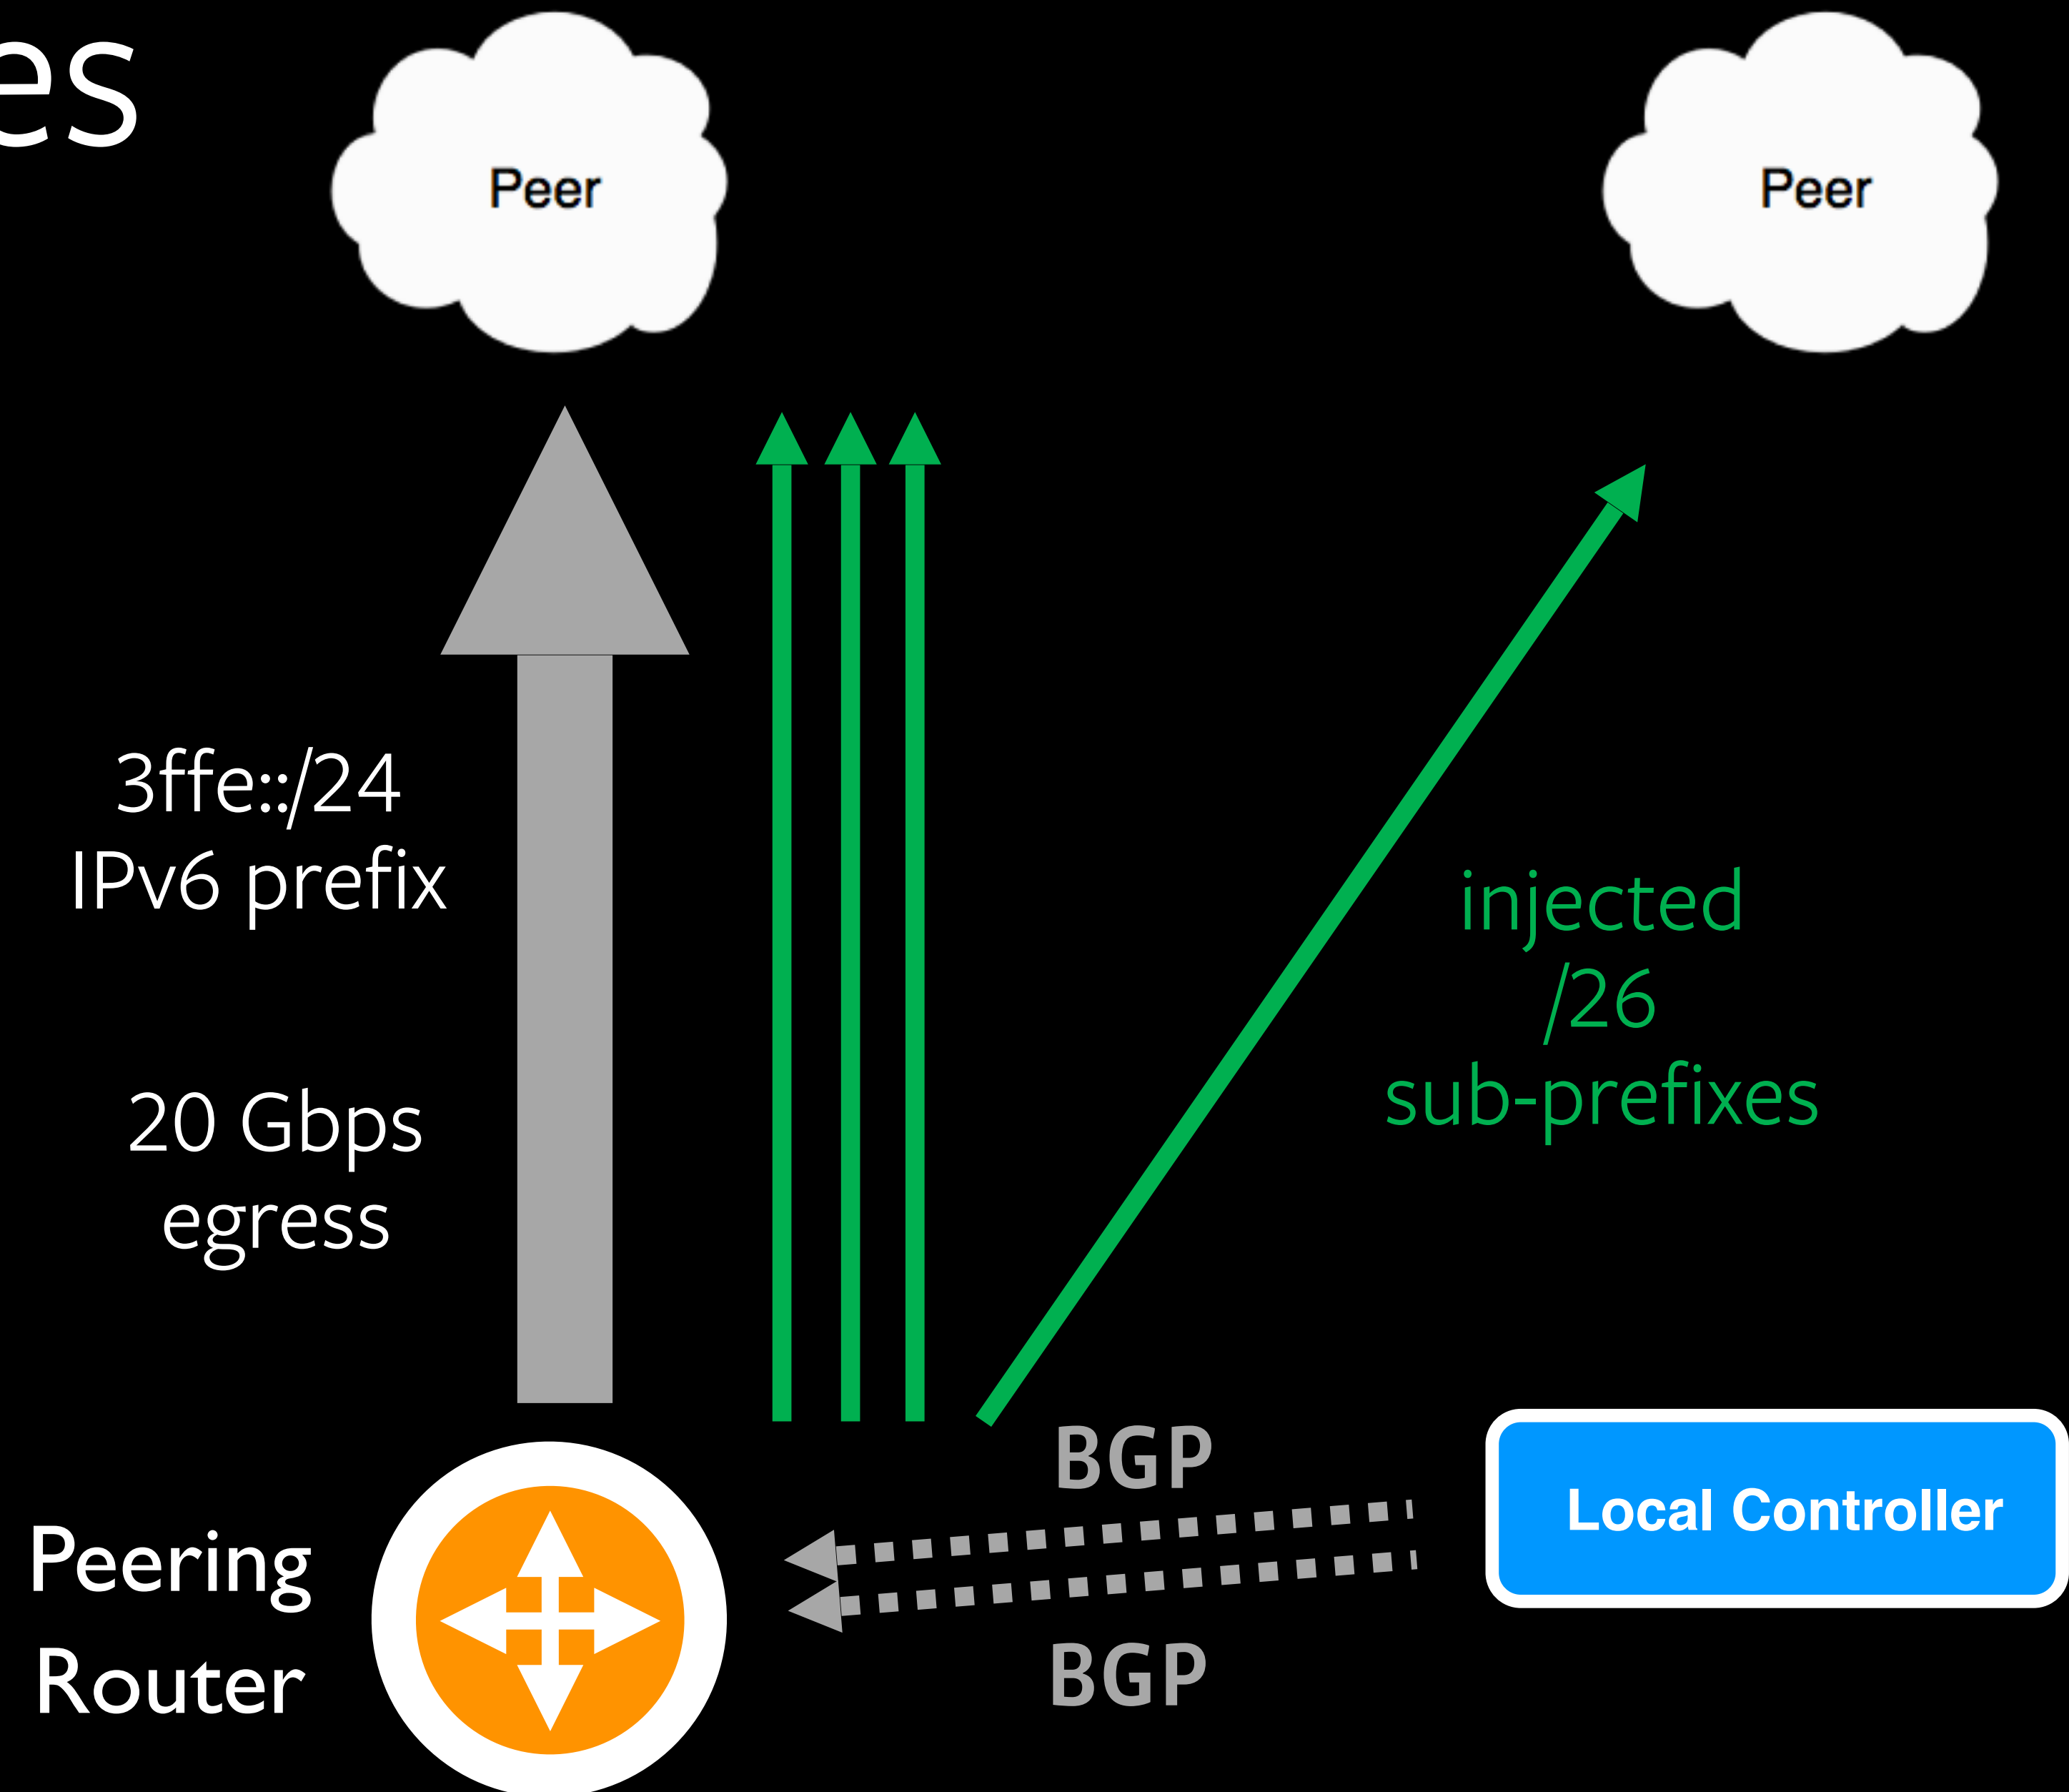


Metro Traffic Engineering




Awesome! So it's all solved?

Huge Prefixes




Oscillation between controllers

Local and Global Controller interaction



Clear roles so they don't oscillate




Oscillation from imprecise data

Lack of Precise Bandwidth Estimation


20 Gbps



10 Gbps



w/o precision




Oscillation from imprecise data

Lack of Precise Bandwidth Estimation


20 Gbps



10 Gbps




w/o precision




w/ precision

All together...




All together...



...and the future

Increasing:

- Modularity
- Programmability
- Performance



Summary

Questions?



Photo Credits

<http://www.flickr.com/photos/27587002@N07/5170590074>

<http://www.flickr.com/photos/yaketyyakyak/7001664846>

<http://www.flickr.com/photos/hinnosaar/3778903507>

<http://www.flickr.com/photos/eamoncurry/8698726494>

<http://www.flickr.com/photos/43158397@N02/4514113429>

<http://www.flickr.com/photos/nobusue/6876280595>

<http://www.flickr.com/photos/29487672@N07/14760573314>

<http://www.flickr.com/photos/joyosity/3595242078>

<http://www.flickr.com/photos/kyntharyn74/3262089319>

<http://www.flickr.com/photos/rexipe/826987087>

<http://www.flickr.com/photos/lablasco/6815671096>