

**Linked in**™

# IPv6@LinkedIn



Zaid Ali Kahn

Senior Director, Infrastructure Engineering  
LinkedIn

# Why does IPv6 matter @LinkedIn?



IPv4 is running out

- IANA exhausted free /8 pool Feb 2011
- APNIC final /8
- ARIN free pool depleted

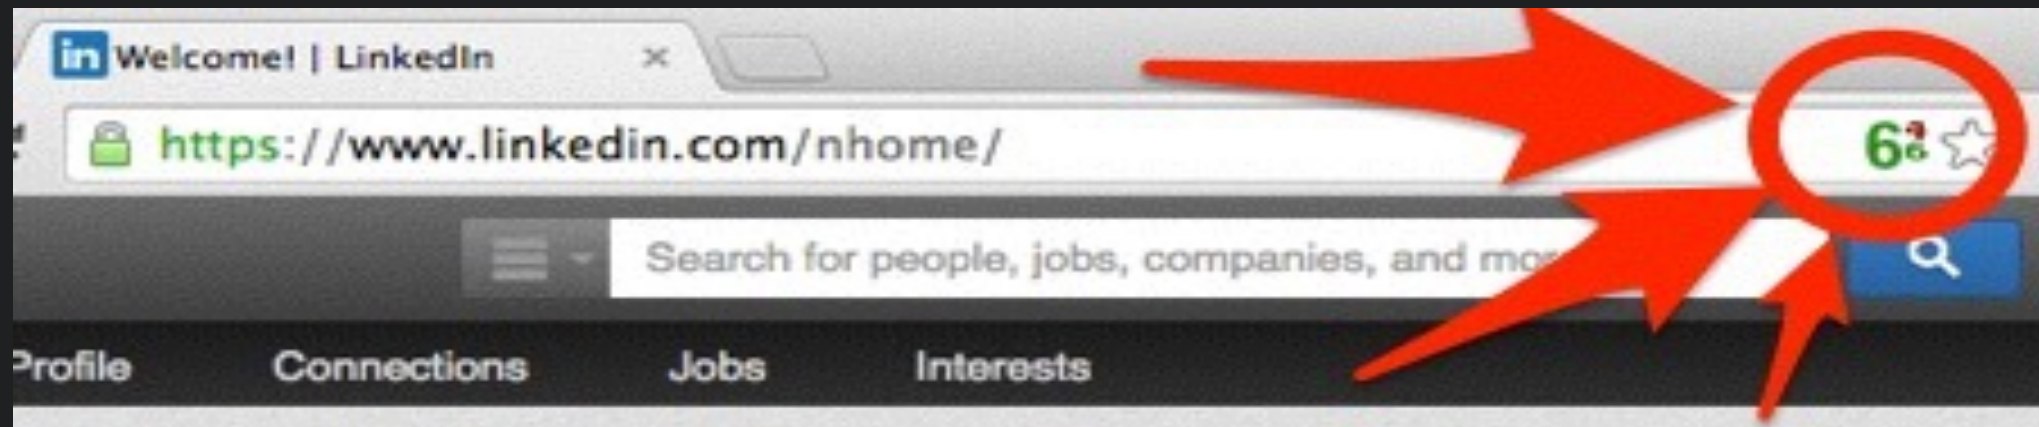
IPv6 performs better

- Strong data for mobile users
- Continued high growth with mobile carriers

Larger datacenters

- RFC1918 exhaustion

# LinkedIn IPv6 launch



Edge services Launched September 8<sup>th</sup> 2014

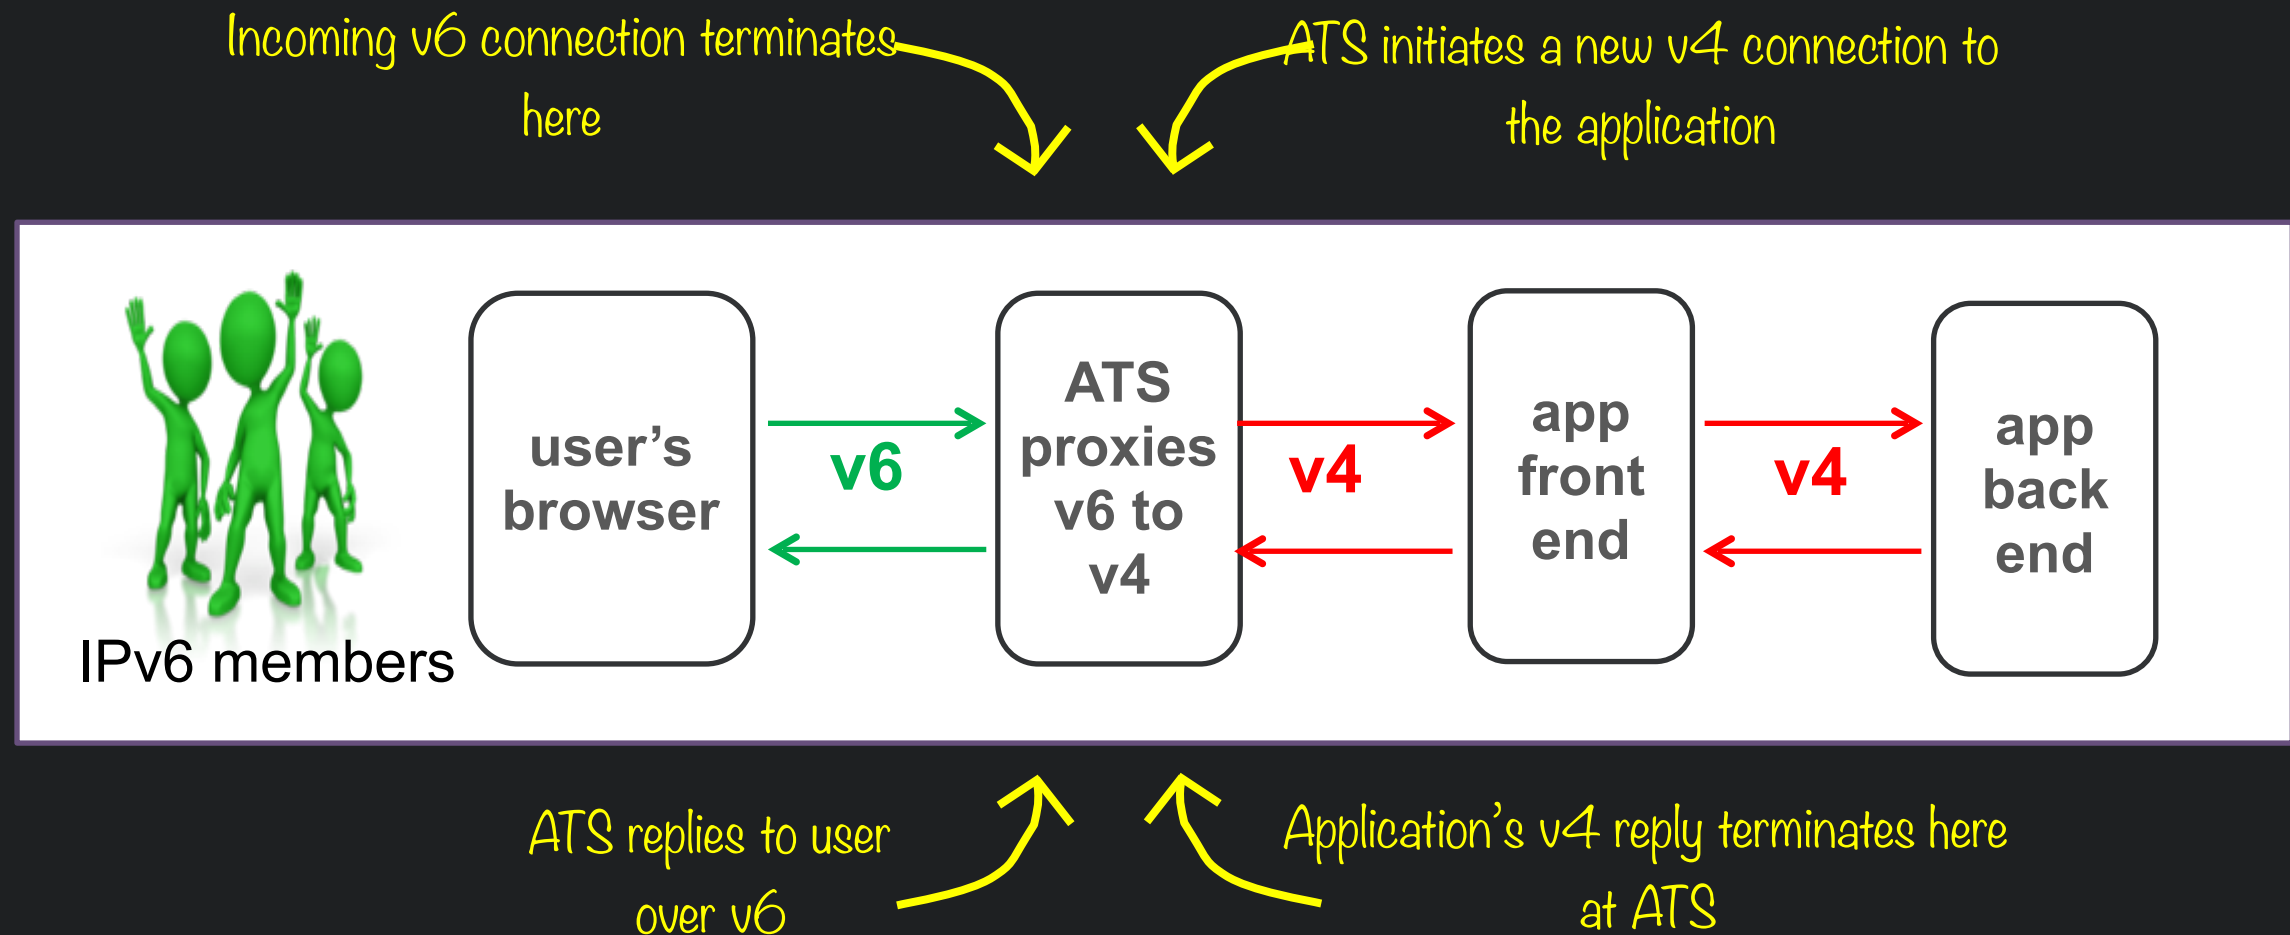
One year of work to transition

Multiple testing of site with real user traffic

Zero site outage

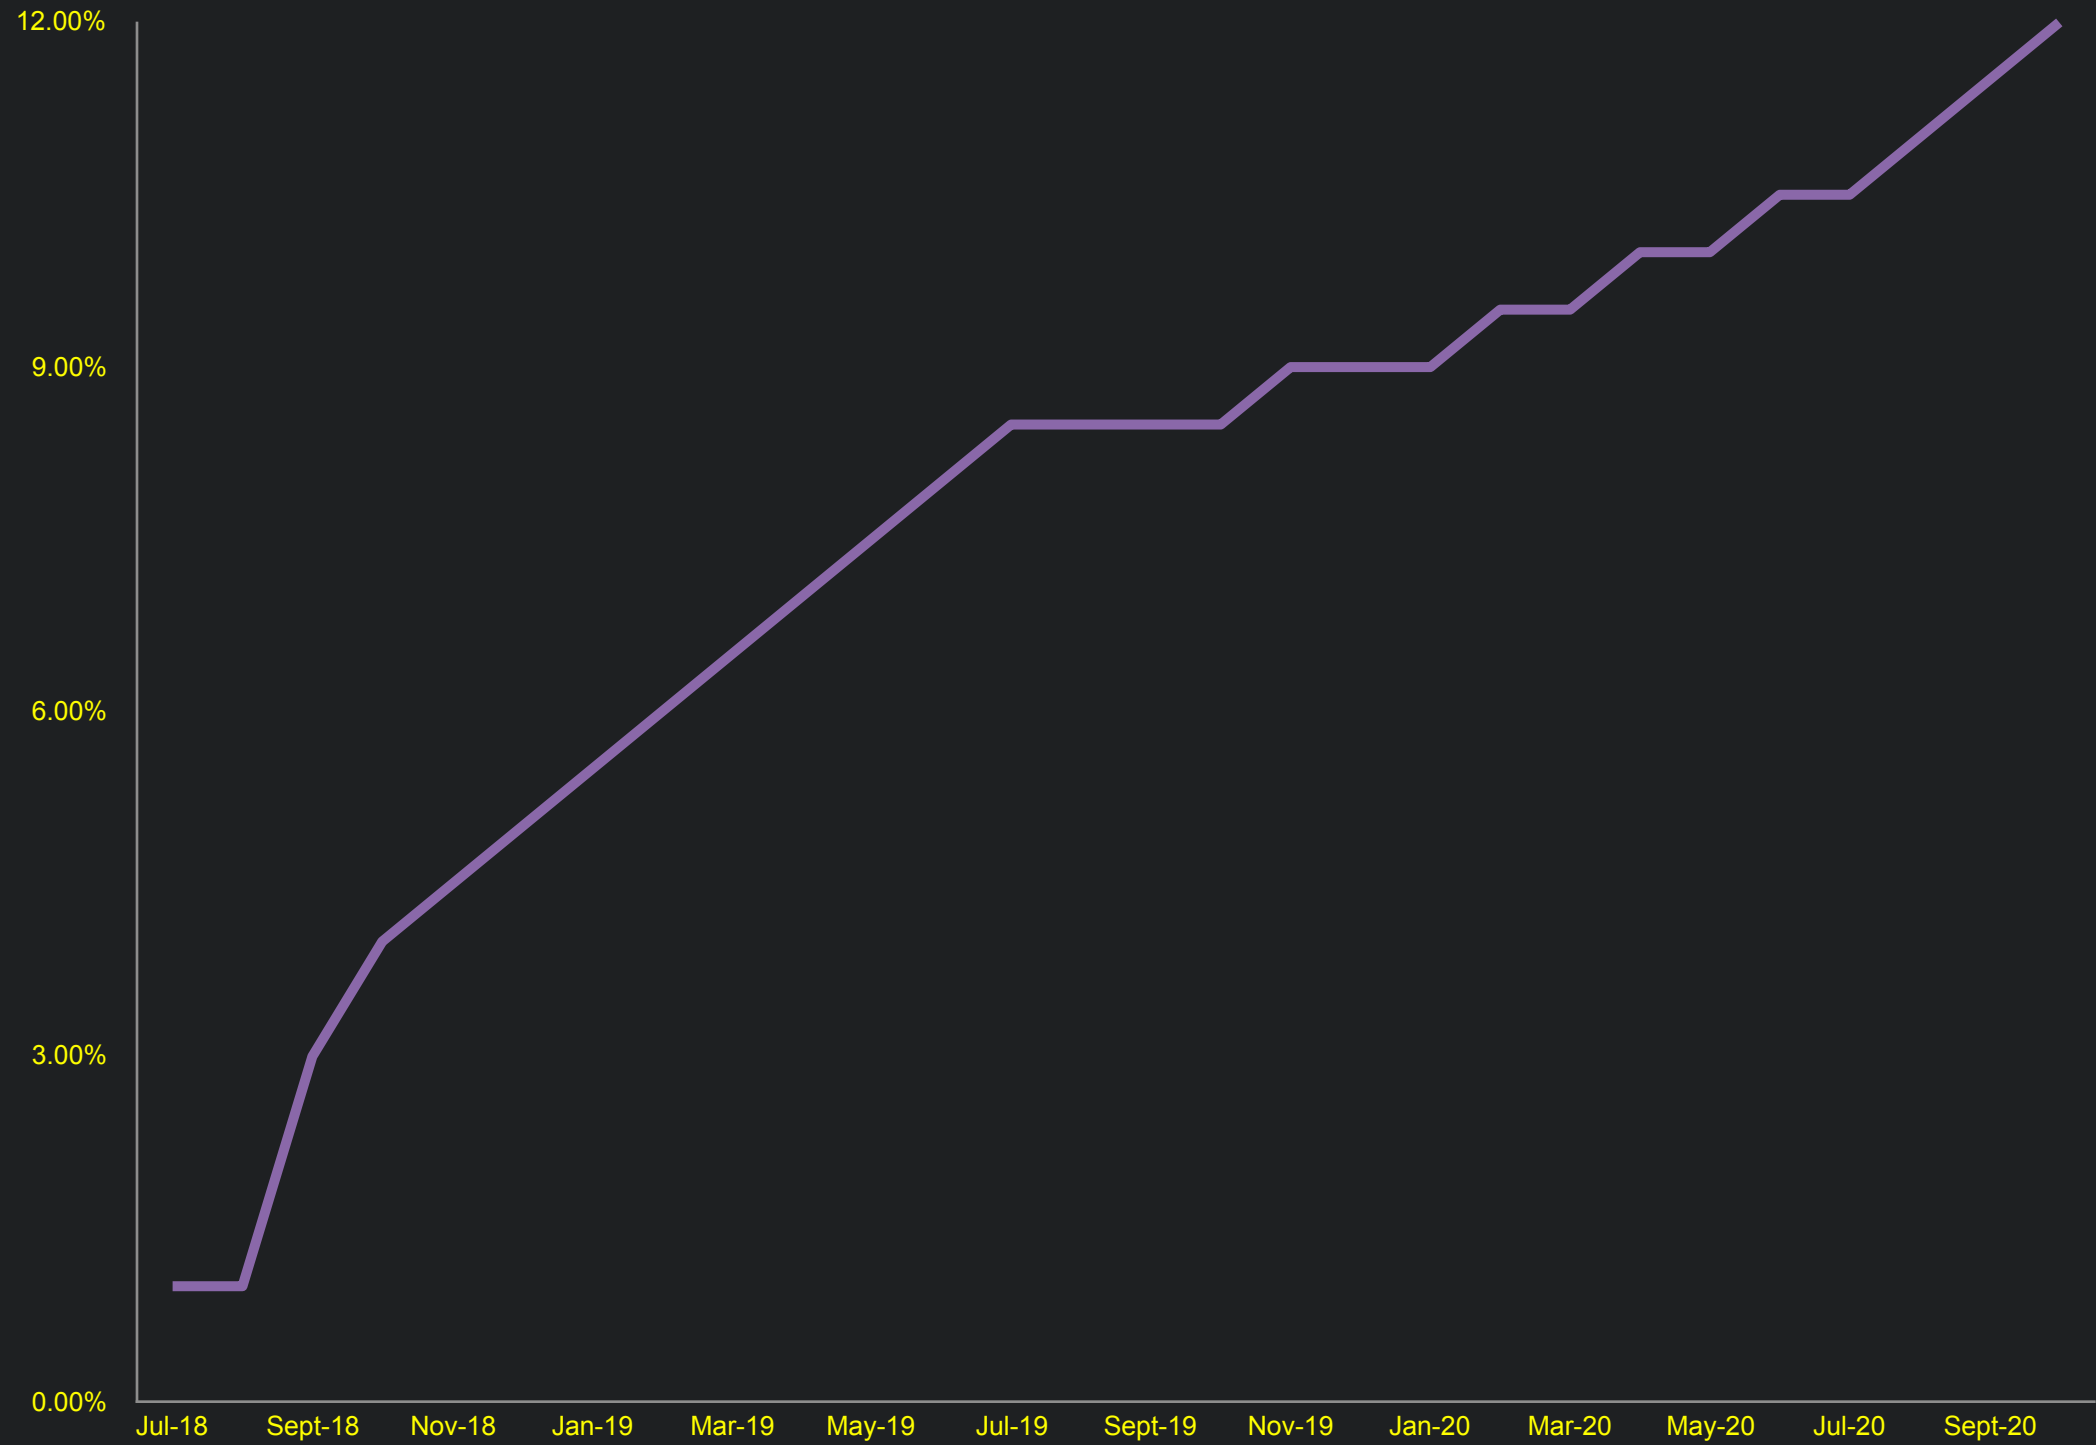
# IPv6-enable the edge

members can communicate over IPv6



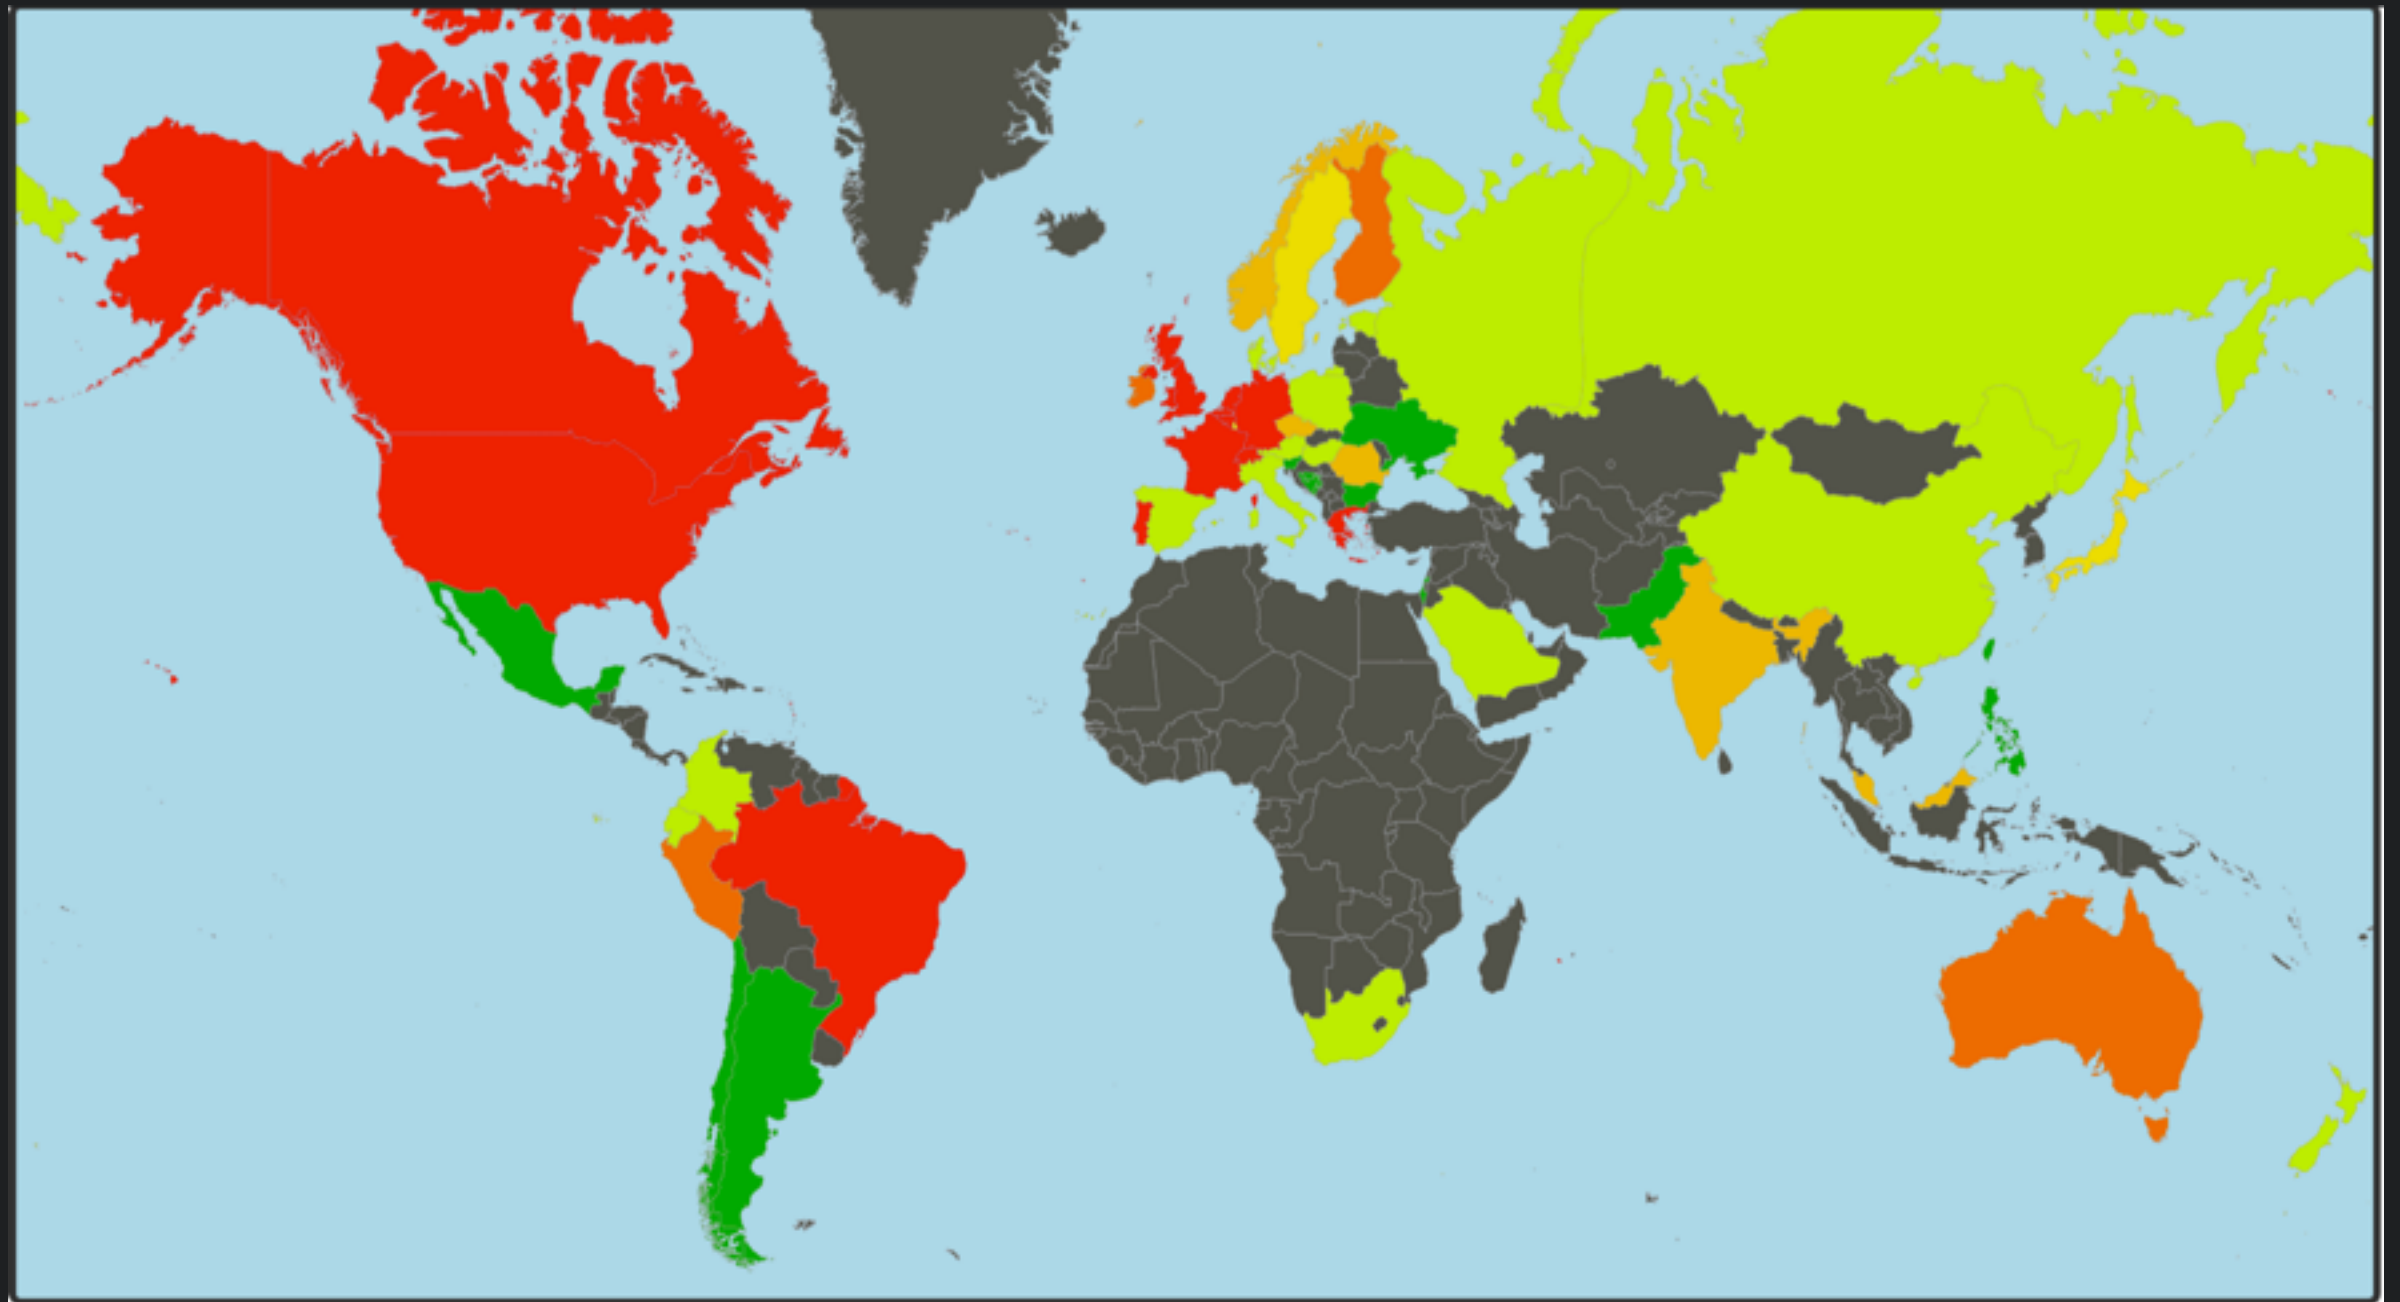


**Internal elements stay IPv4**

# IPv6 growth at LinkedIn



# IPv6 global heatmap



# RUM – Real User Monitoring

```
Elements Network Sources Timeline Profiles
<top frame>
> rumTiming = window.performance.timing;
  ▶ PerformanceTiming {loadEventEnd: 1405967695234,
    domContentLoadedEventEnd: 1405967695225,
  }
> rumTiming.connectEnd - rumTiming.connectStart
136
> rumTiming.responseEnd - rumTiming.responseStart
401
> |
```

Most browsers support RUM  
Navigation timing API

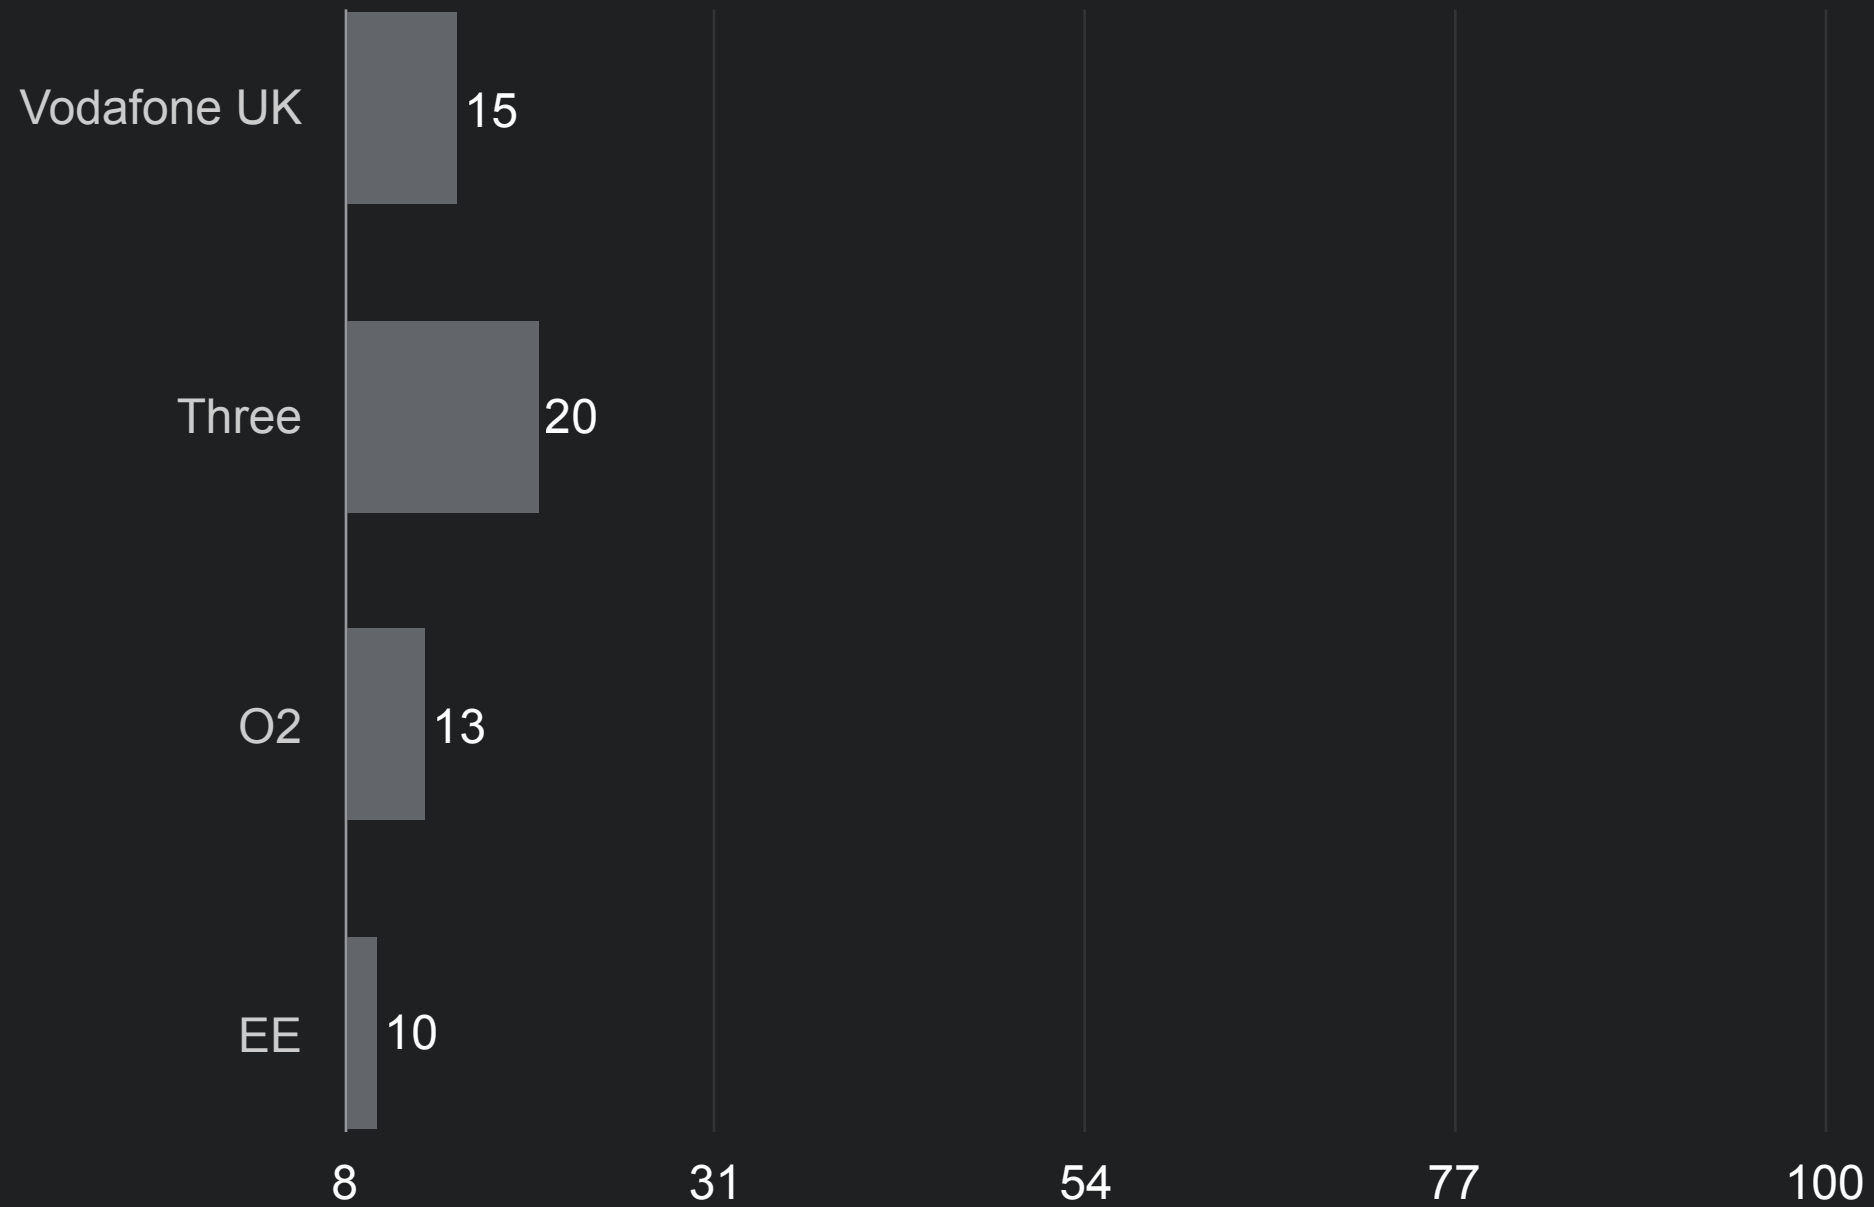
- Browsers record performance timers for each payload.

Mobile end-end control


- Client/Server control gives full visibility
- Light weight beacon




# Mobile Performance (UK)



# Mobile Performance (USA)



# IPv6/IPv4 timeout

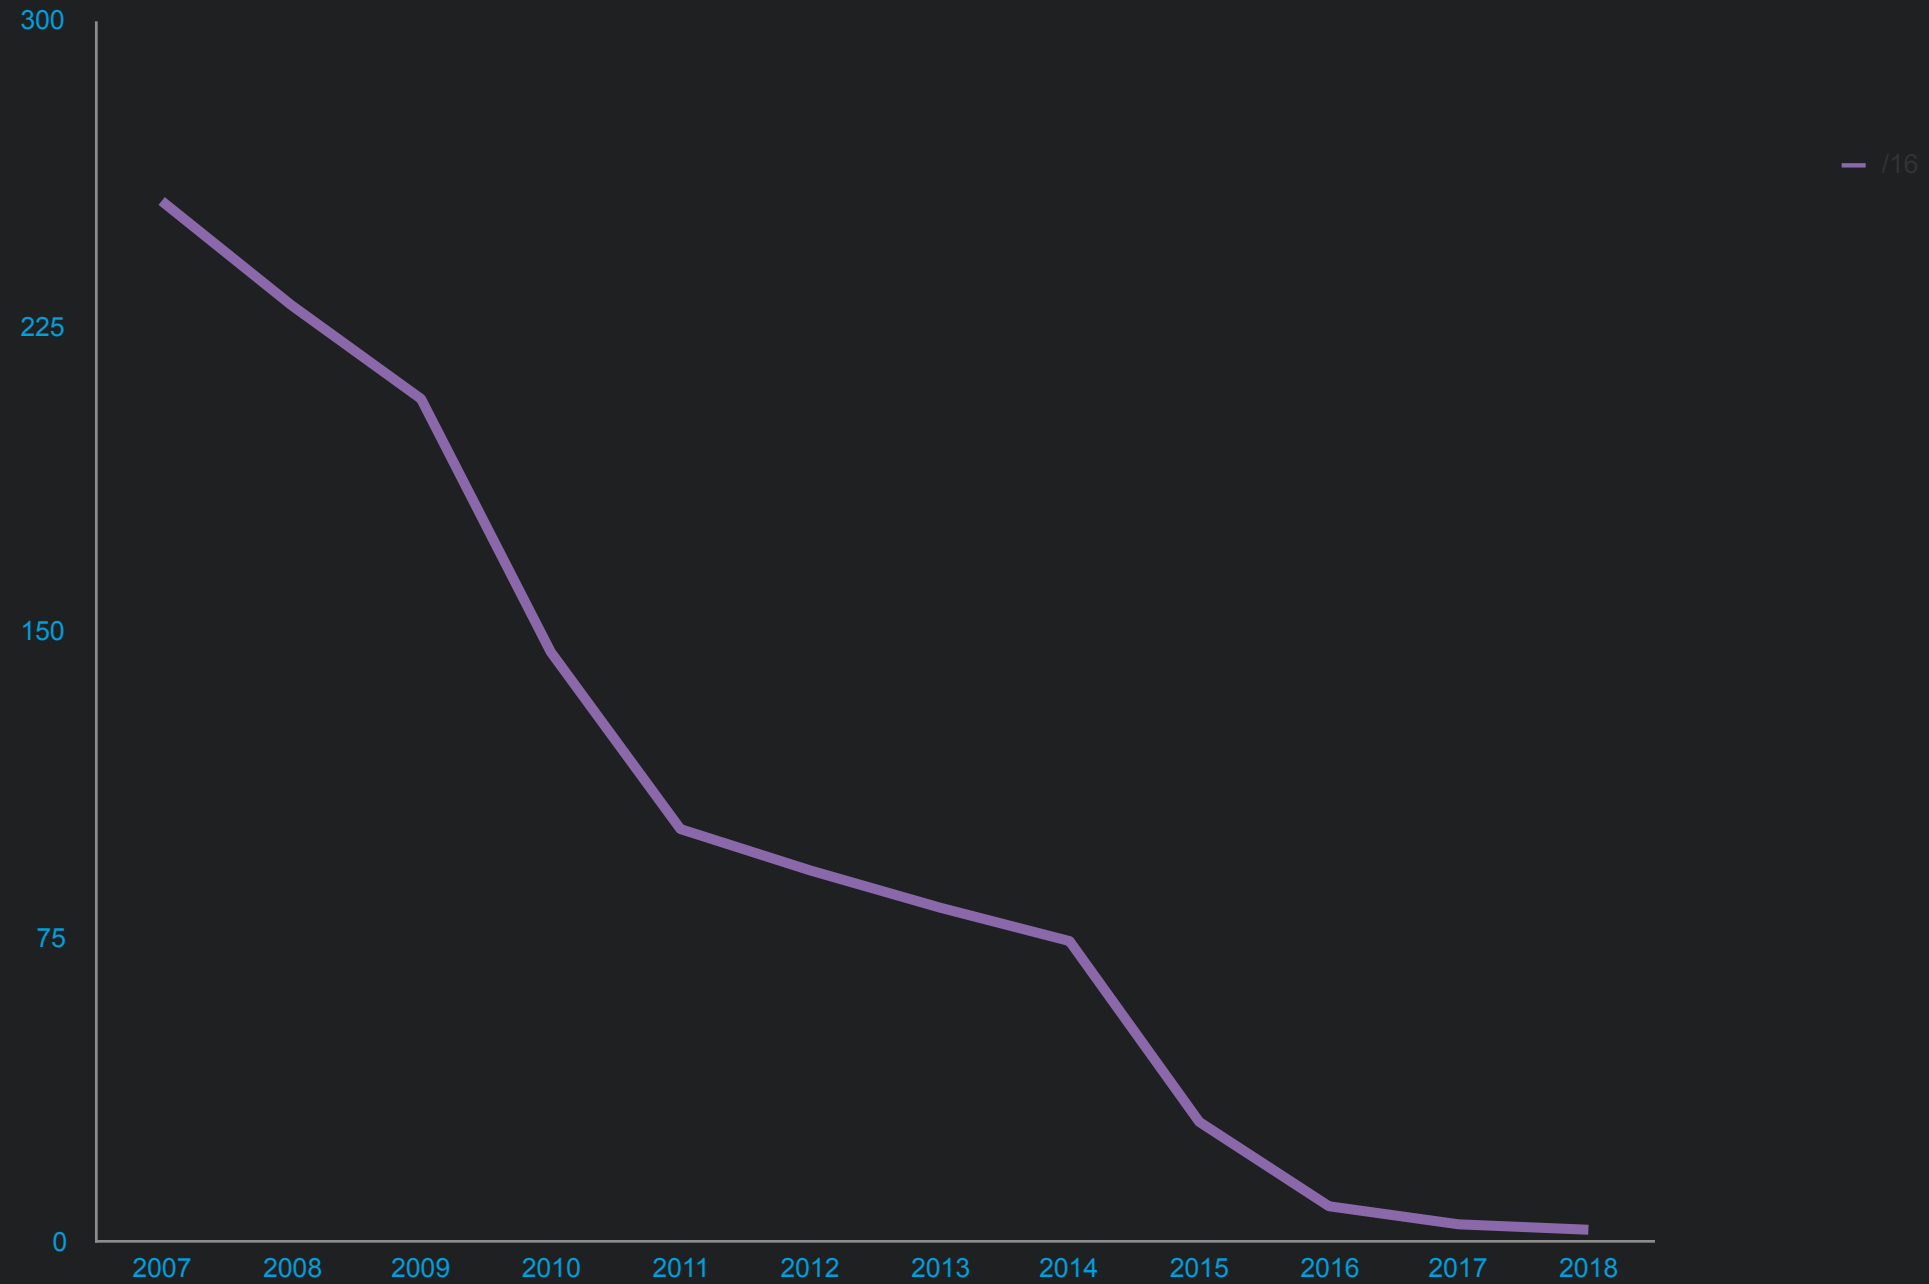


**Higher network request time means higher TCP timeouts yielding in higher page download load times.**

# IPv6 inside the datacenter



# Datacenters RFC1918 /8 exhaustion



# Plan/Key Steps

- Dual stack latest datacenter in Oregon (2016)
- Dual stack remaining datacenters (2016-2017)
- Application stack migrates to support IPv6
- New datacenters in 2017 IPv6 only
- Turn off IPv4 by 2018

# What you have to think about

- Verifying all existing devices and futures devices fully support IPv6 and able to scale
- Having conversations with application developers and SRE teams.
- Making a case with data.
- Subnetting and numbering policies
- DHCPv6/SLAAC
- Infrastructure automation tools
- DC RDU, Power management
- DNS, Syslog, Auth, NTP
- Packet too big (PTB) – ensure messages are generated.

# Things you should not do

- Any kind of NAT. Live in the future!
- When designing address pools don't map to IPv4. Think IPv6!
- Don't do this:

```
2620:abcd:efef::192.168.1.1
```

## Instead

```
find AAAA
if not found {
  get A
  find IPv6 paired_subnet_id of A
  Take last 2 bytes of A, convert to hex, pseudo AAAA is IPv6 paired_subnet_id with last 2 bytes of A
}
```

```
$ host dc1-gggg.env1.linkedin.com
```

```
dc1-gggg.env1.linkedin.com has address 10.136.150.167
```

```
host dc1-gggg.env1.linkedin.com has IPv6 address 2620:f549:14:4124::96a7
```



# Future work

- IETF draft: draft-ali-ipv6rtr-reqs-00
- “Requirements for IPv6 Routers”
  - YANG [RFC7223] [RFC7224] [RFC7277] [RFC7317]
  - Netconf [RFC6241] / Restconf [I-D.ietf-netconf-restconf]
  - Zero Touch Provisioning [I-D.ietf-dhc-rfc3315bis]
  - SLAAC [RFC7217] [RFC7527]
  - Telemetry
  - Router handling of IPv6 addresses
  - MTU/Jumbo Frames
- Active discussions, please join and contribute!



