Buffer sizing and Video QoE Measurements at Netflix

Bruce Spang, Brady Walsh, Te-Yuan Huang,
Tom Rusnock, Joe Lawrence, Nick McKeown

February 10, 2020

What are we talking about?

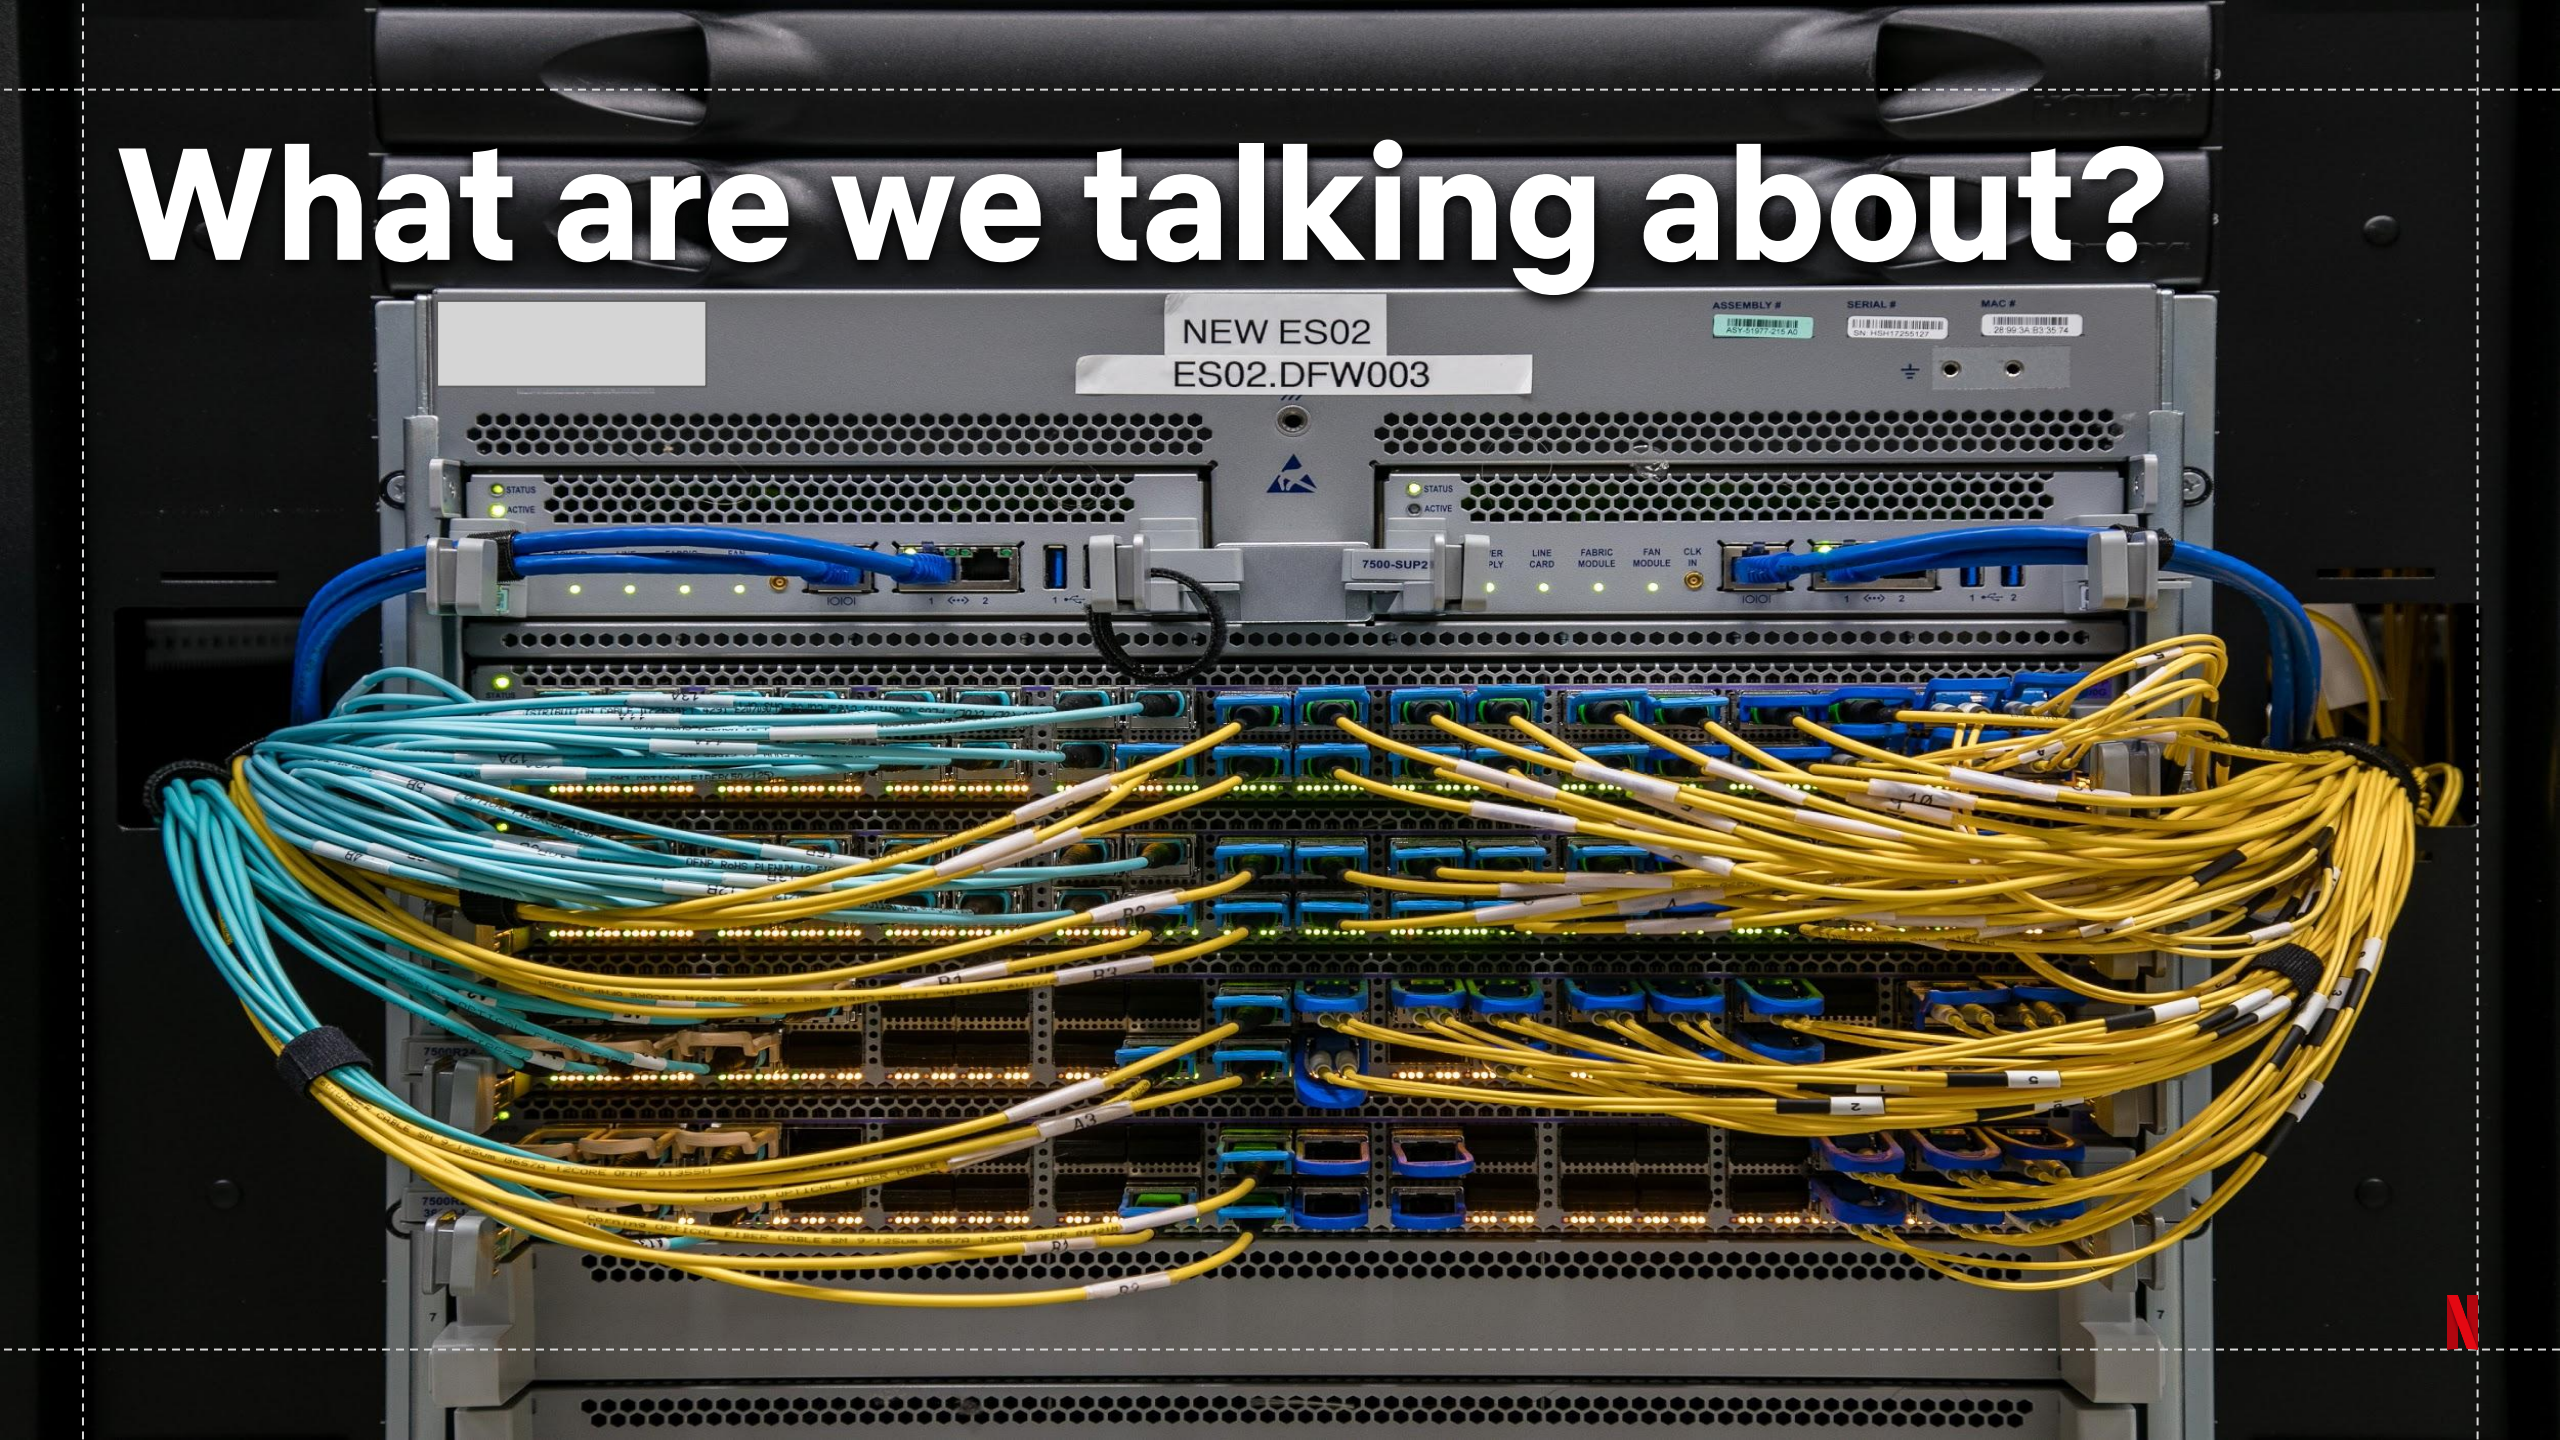

What are we talking about?

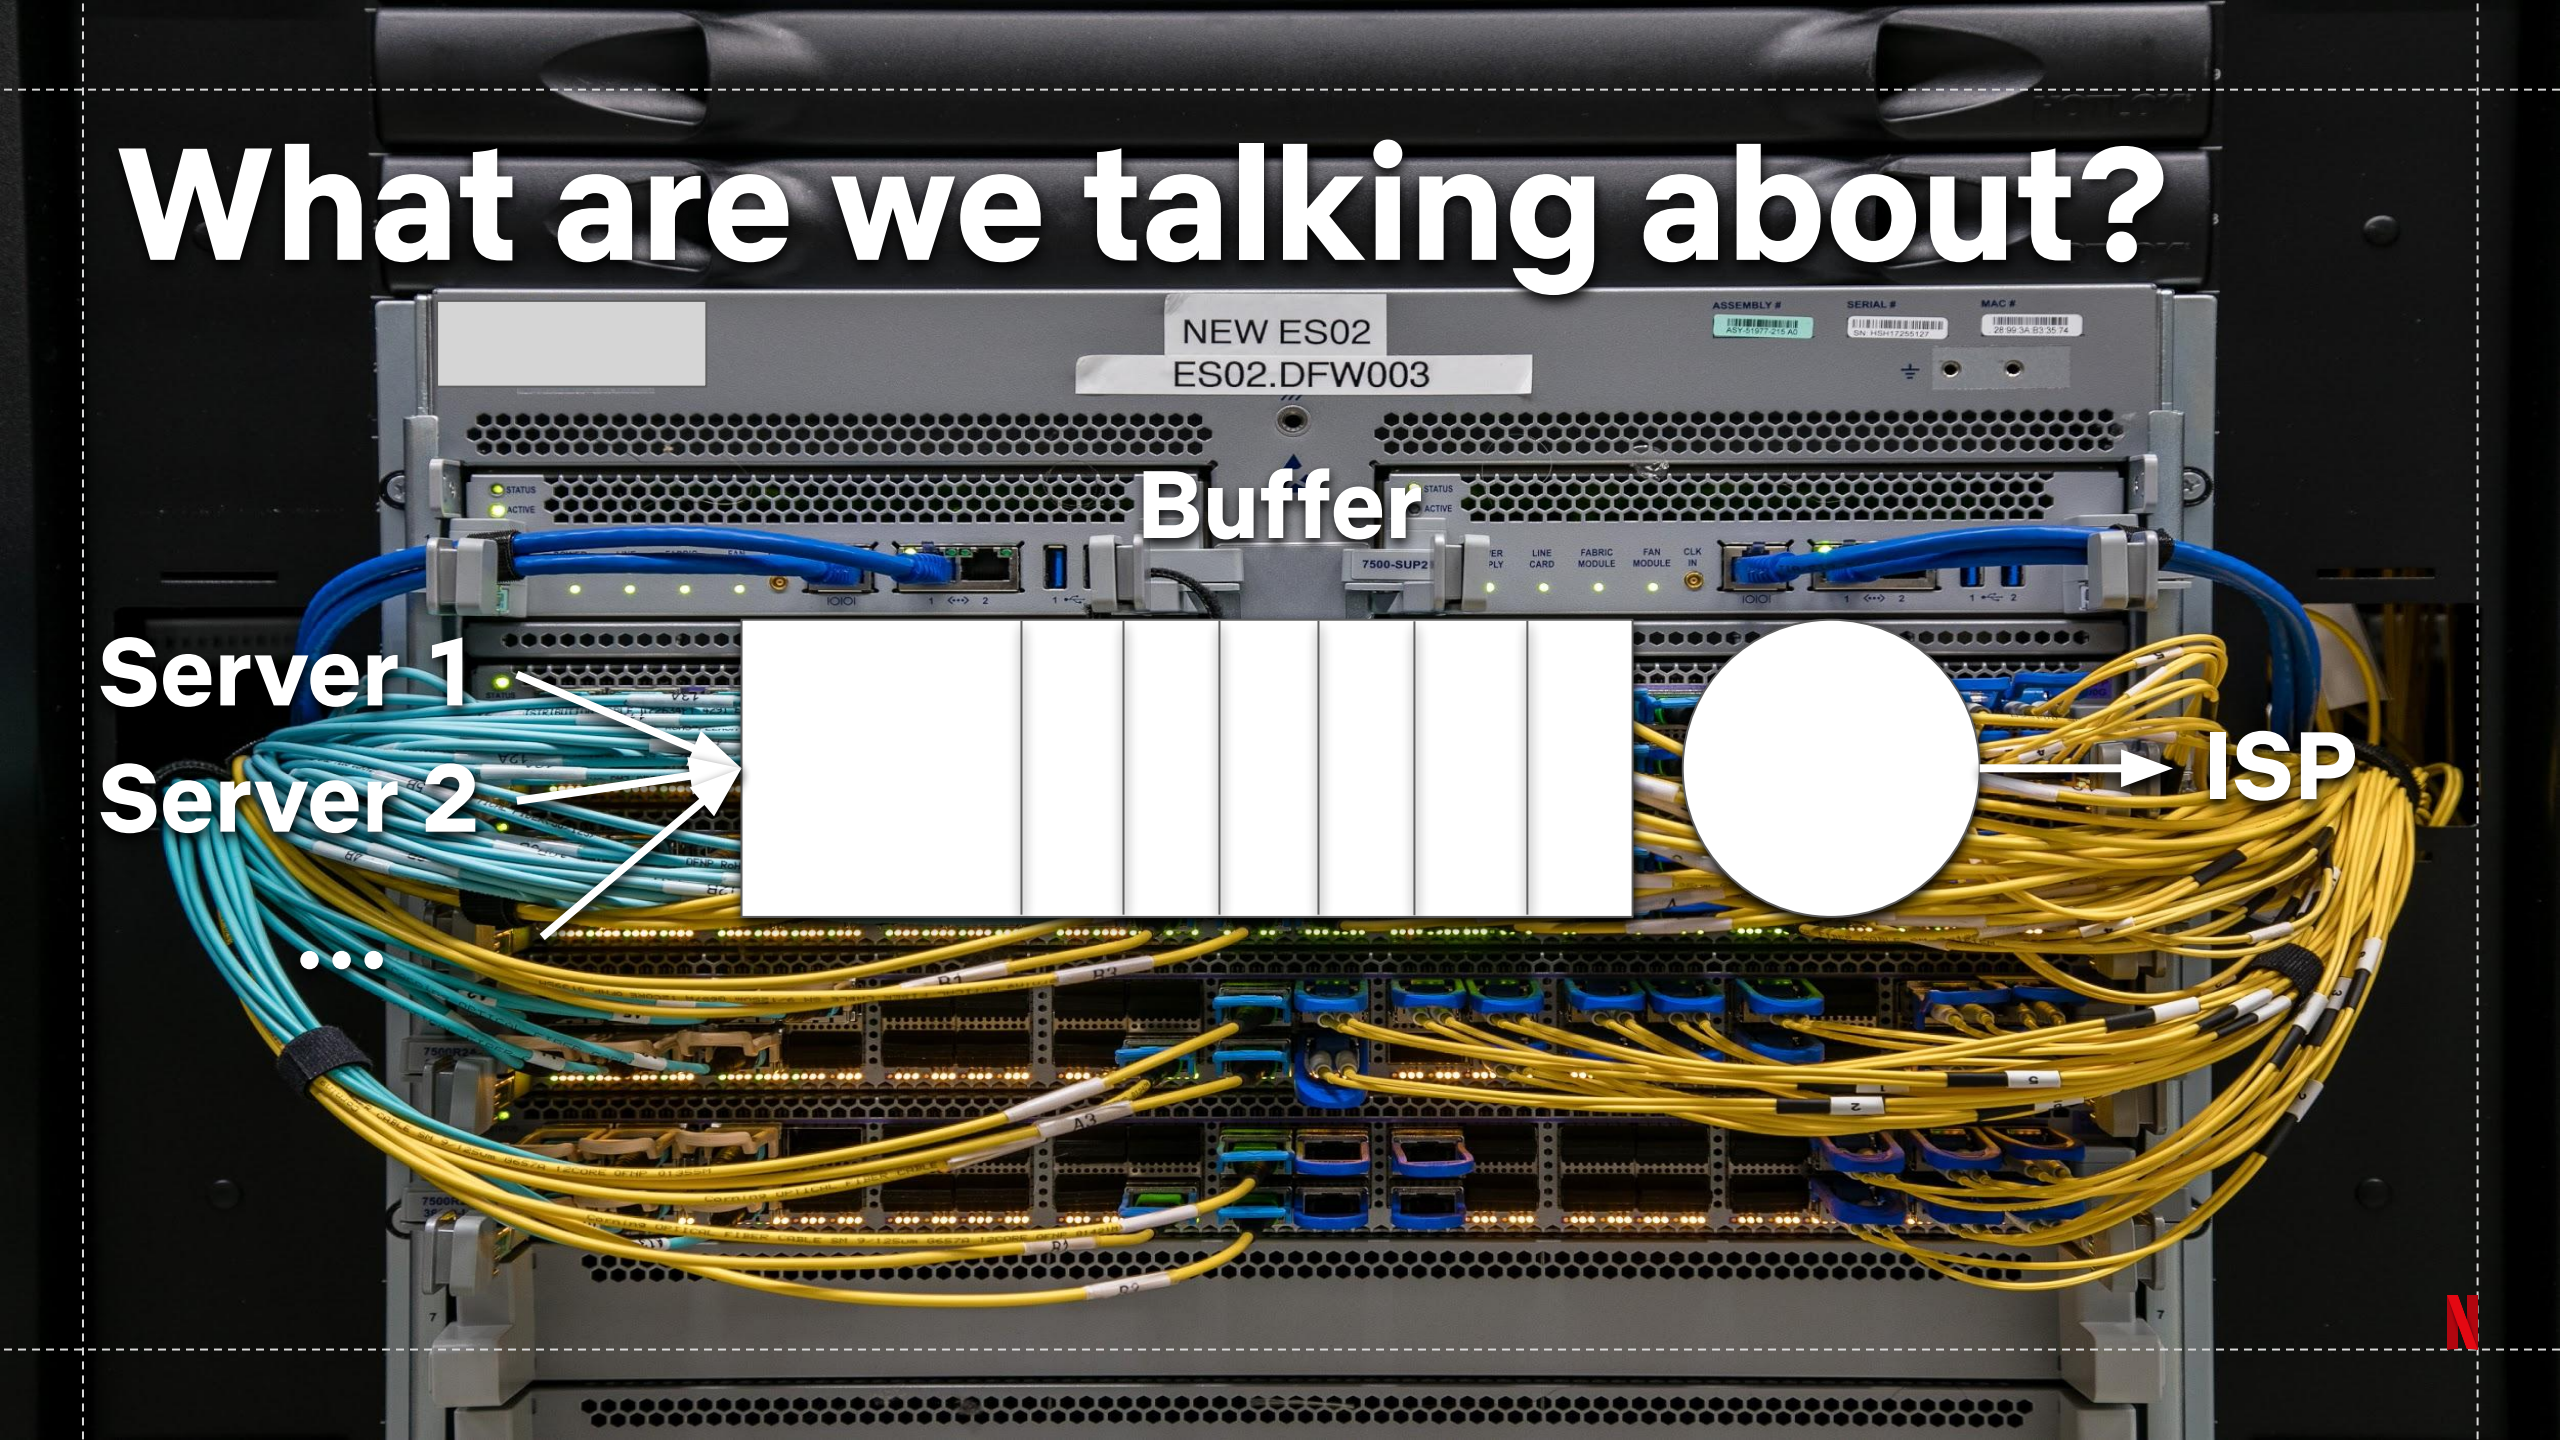

How big should a buffer be?

Too big: packets wait for too long

Too small: too many packets thrown away

“A buffer should be at least one BDP” [Villamizar, Song 1994]

“A buffer should be at least one BDP” [Villamizar, Song 1994]

BDP=Bandwidth x Delay
# of packets in a link for full utilization

“A buffer should be at least one BDP” [Villamizar, Song 1994]

BDP=Bandwidth x Delay
# of packets in a link for full utilization

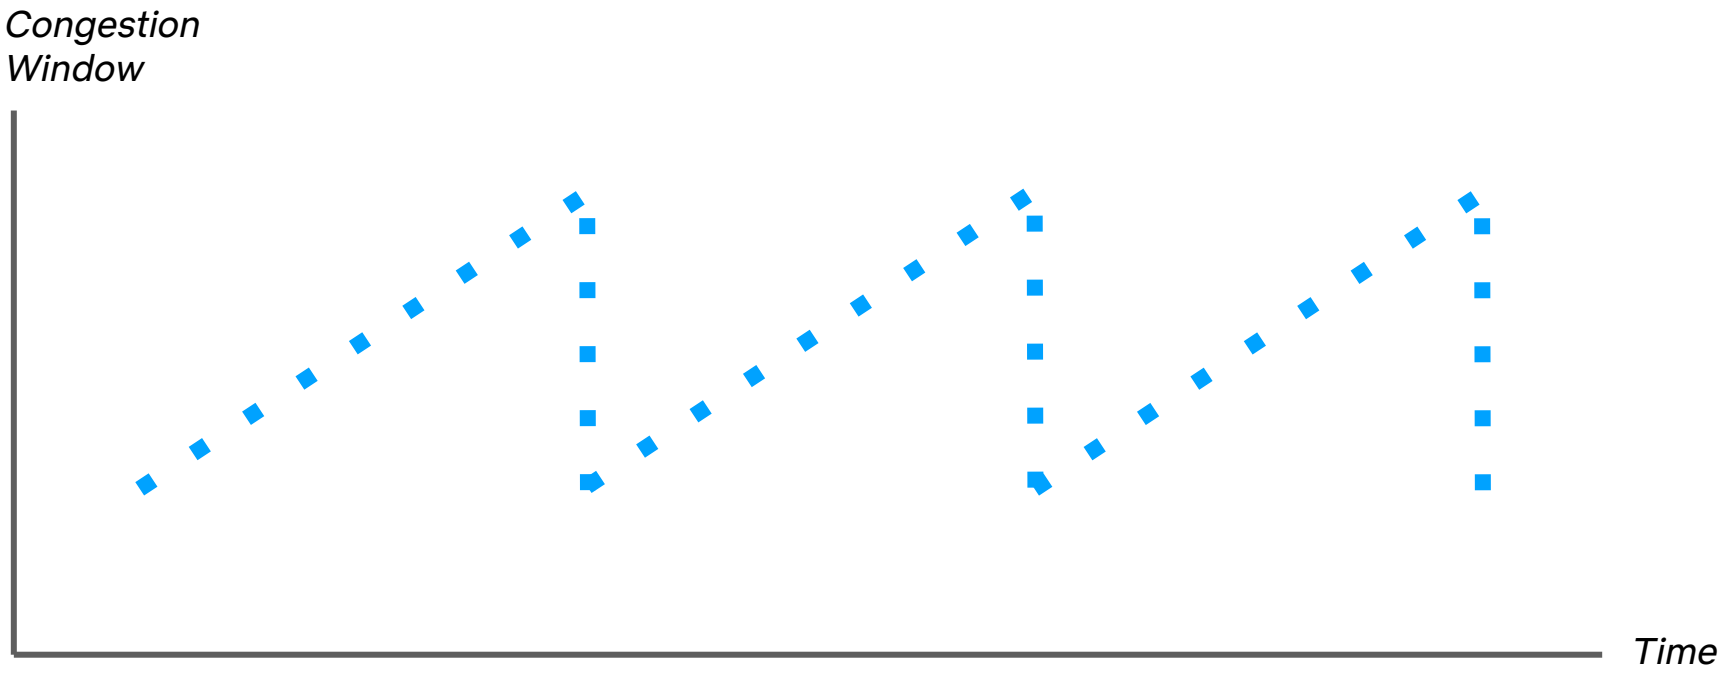

“A buffer should be at least one BDP” [Villamizar, Song 1994]

BDP=Bandwidth x Delay
# of packets in a link for full utilization

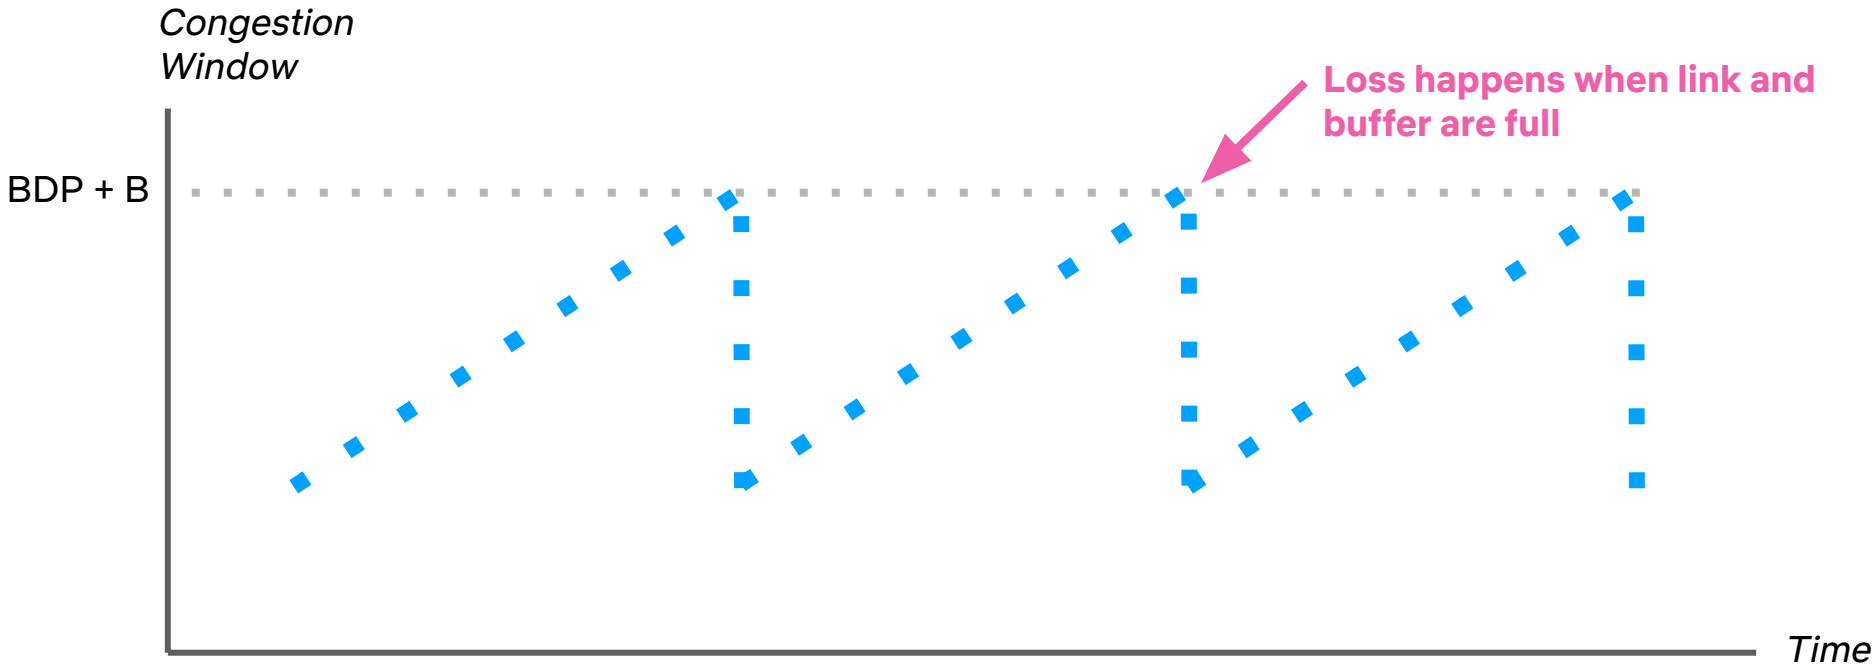

“A buffer should be at least one BDP” [Villamizar, Song 1994]

BDP=Bandwidth x Delay
# of packets in a link for full utilization

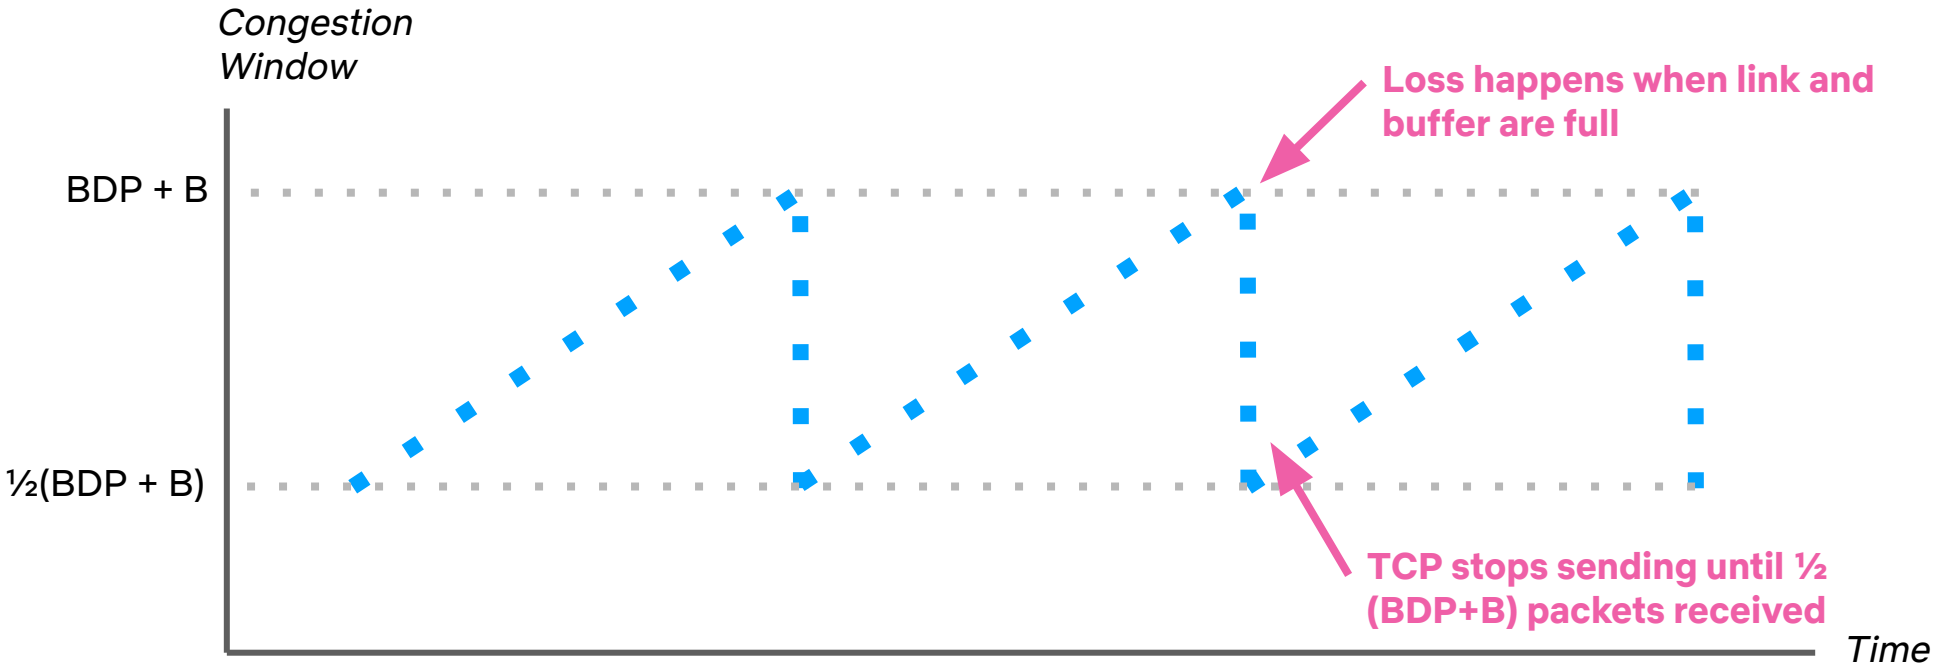

“A buffer should be at least one BDP” [Villamizar, Song 1994]

BDP=Bandwidth x Delay
# of packets in a link for full utilization

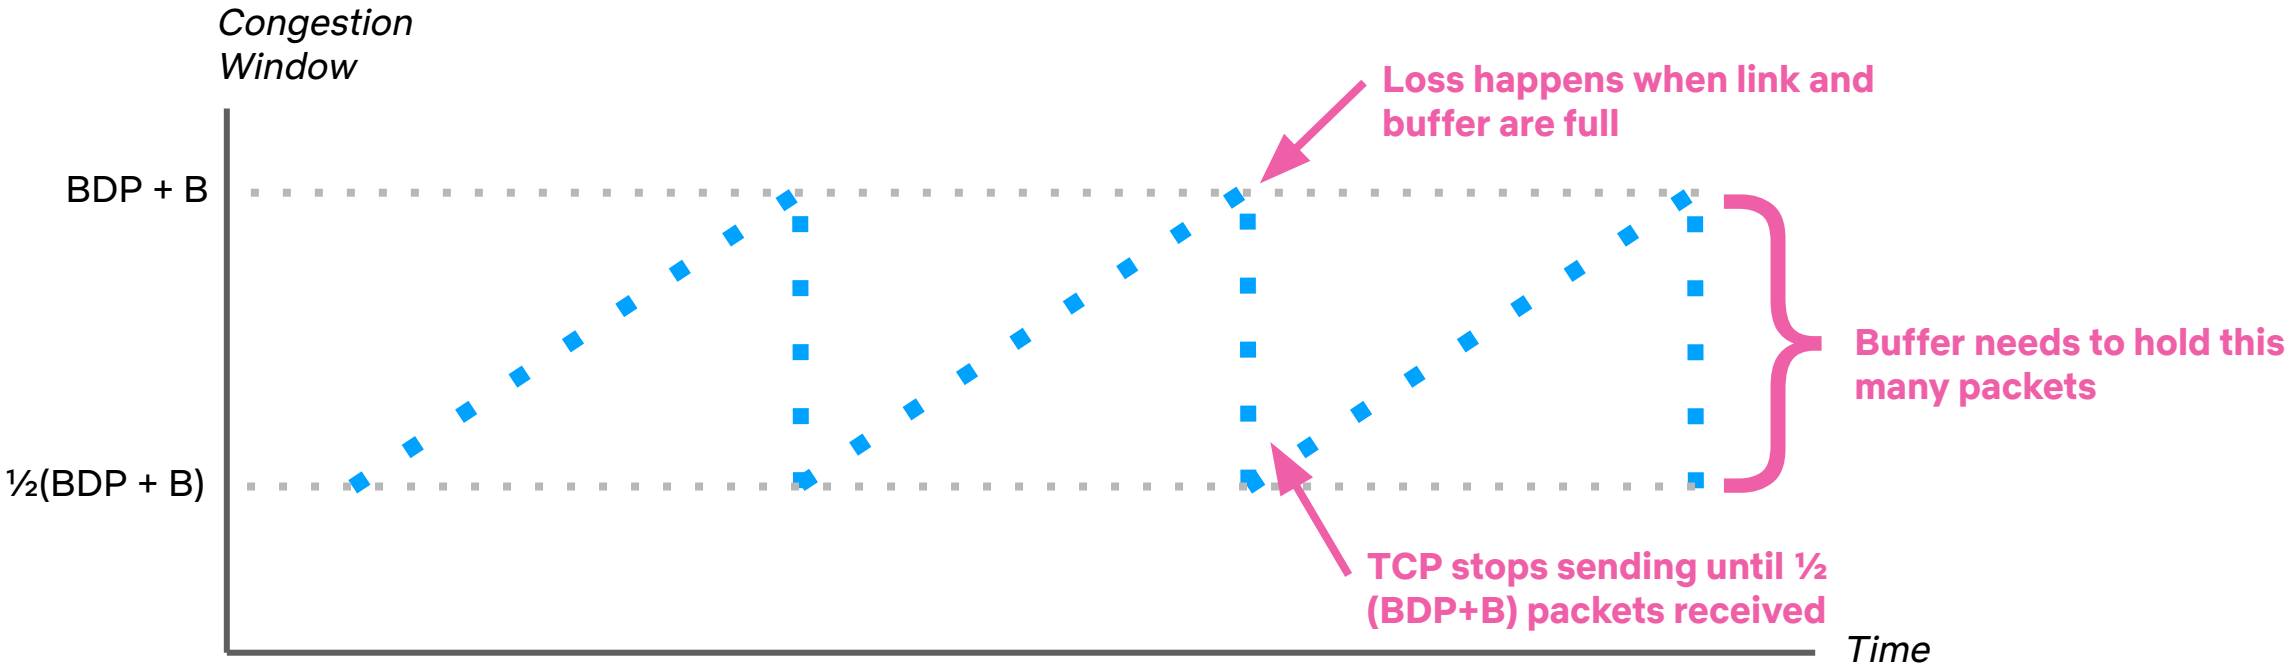

How big should a buffer be?

BDP: Villamizar and Song 1994

BDP/√n: Appenzeller, McKeown, Keslassy 2004

O(n): Dhamdhere, Jiang, Dovrolis 2005

O(1): Enachescu, Ganjali, Goel, McKeown, Roughgarden 2006

Which is correct?

It’s complicated

1. TCP New Reno (mostly) behaves as expected
2. Video performance varies
3. Real routers complicate this story


Our Experiment

Catalog servers

Uses spinning disks, cheaply stores entire catalog

Offload servers

Use SSDs to serve top ~30% of content faster

These three racks are called a **stack**

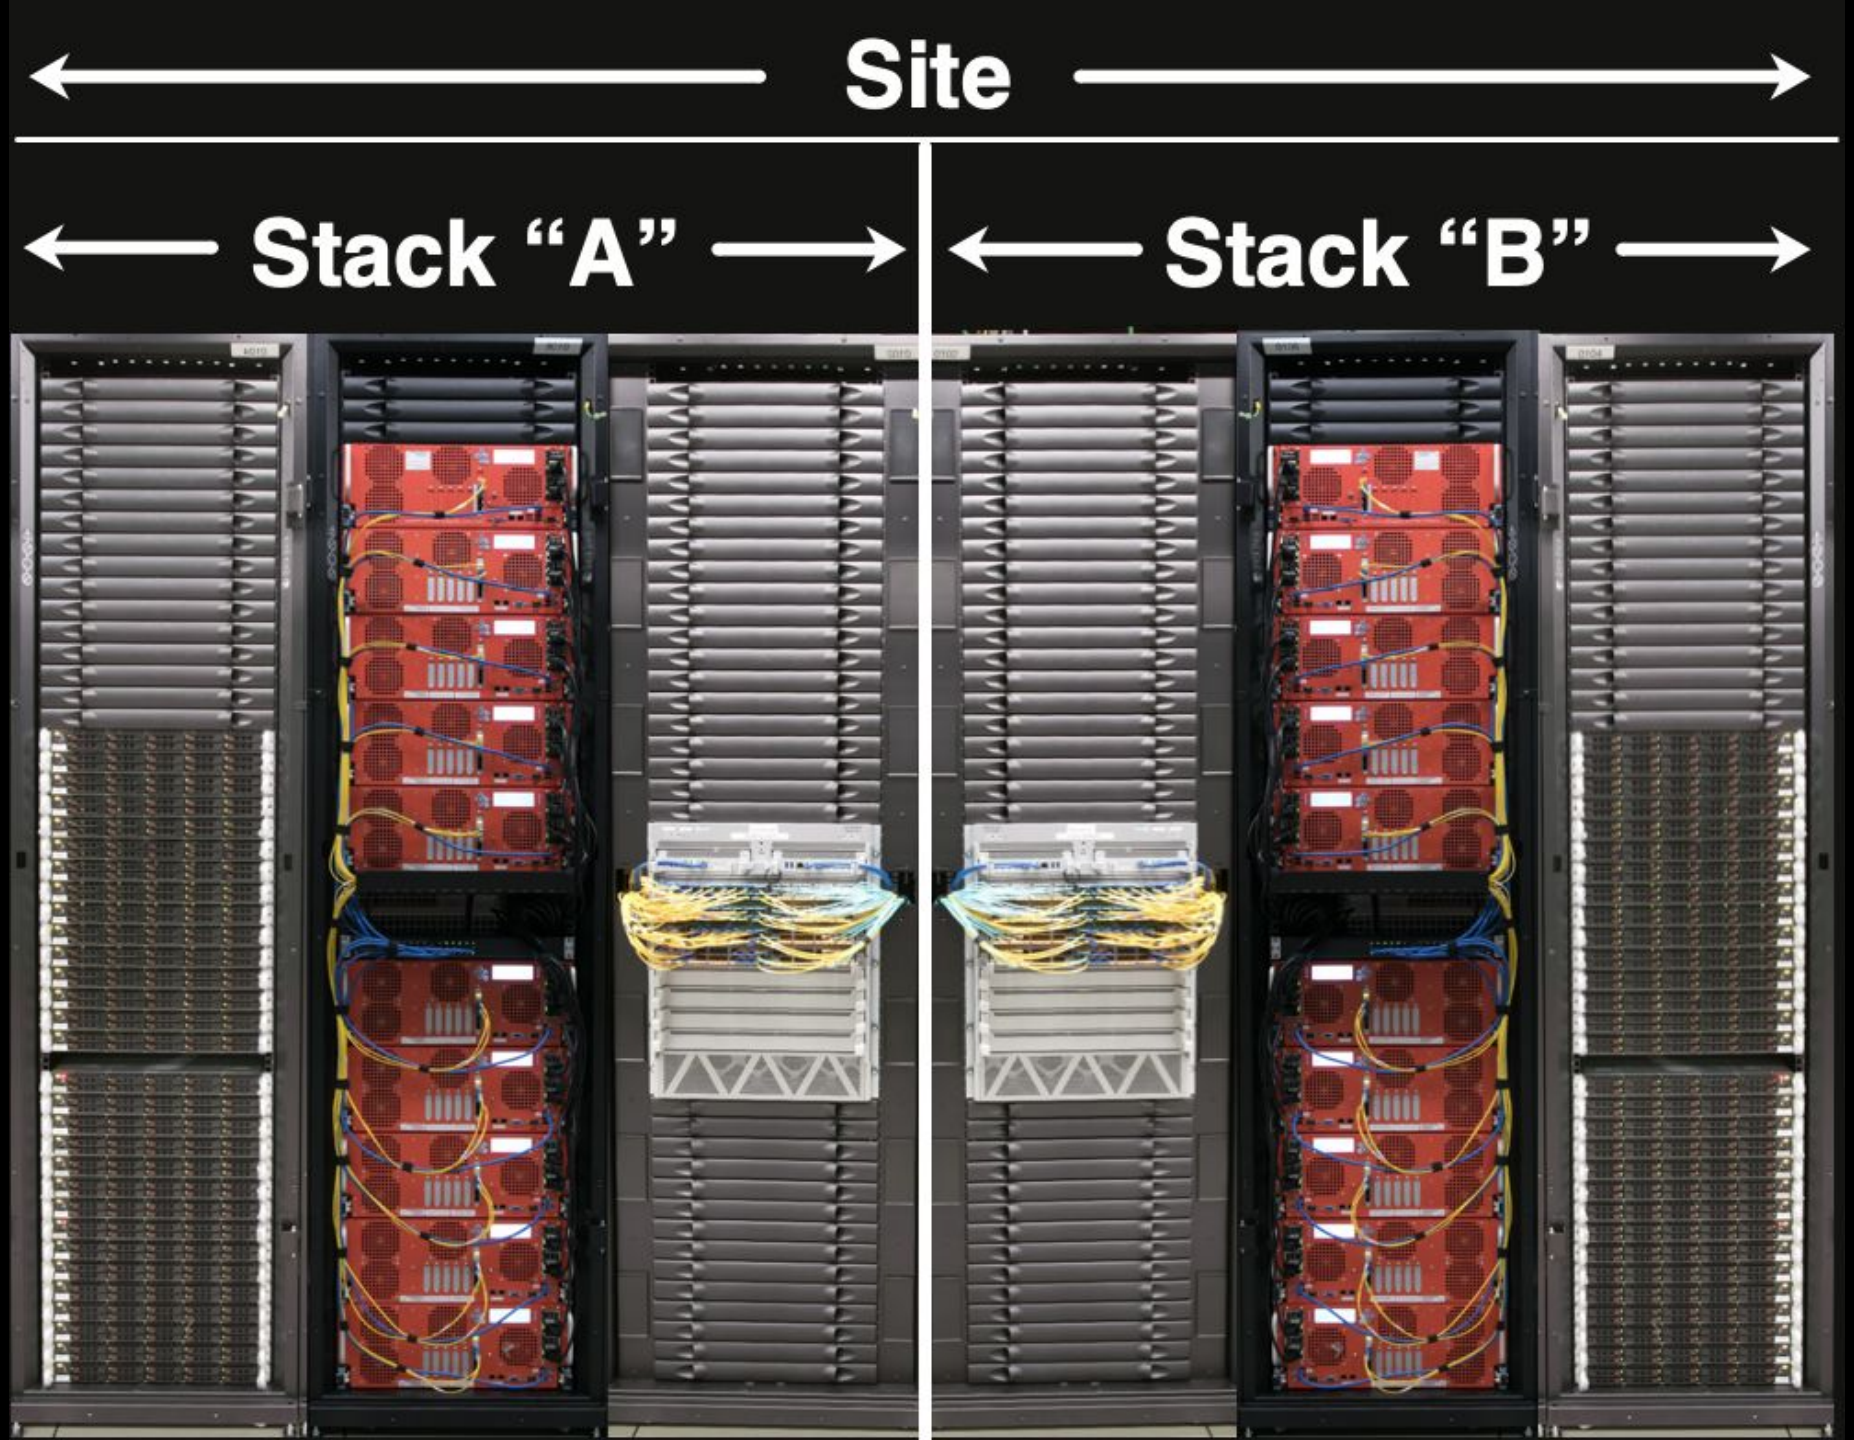
Site
Stack “A”
Stack “B”

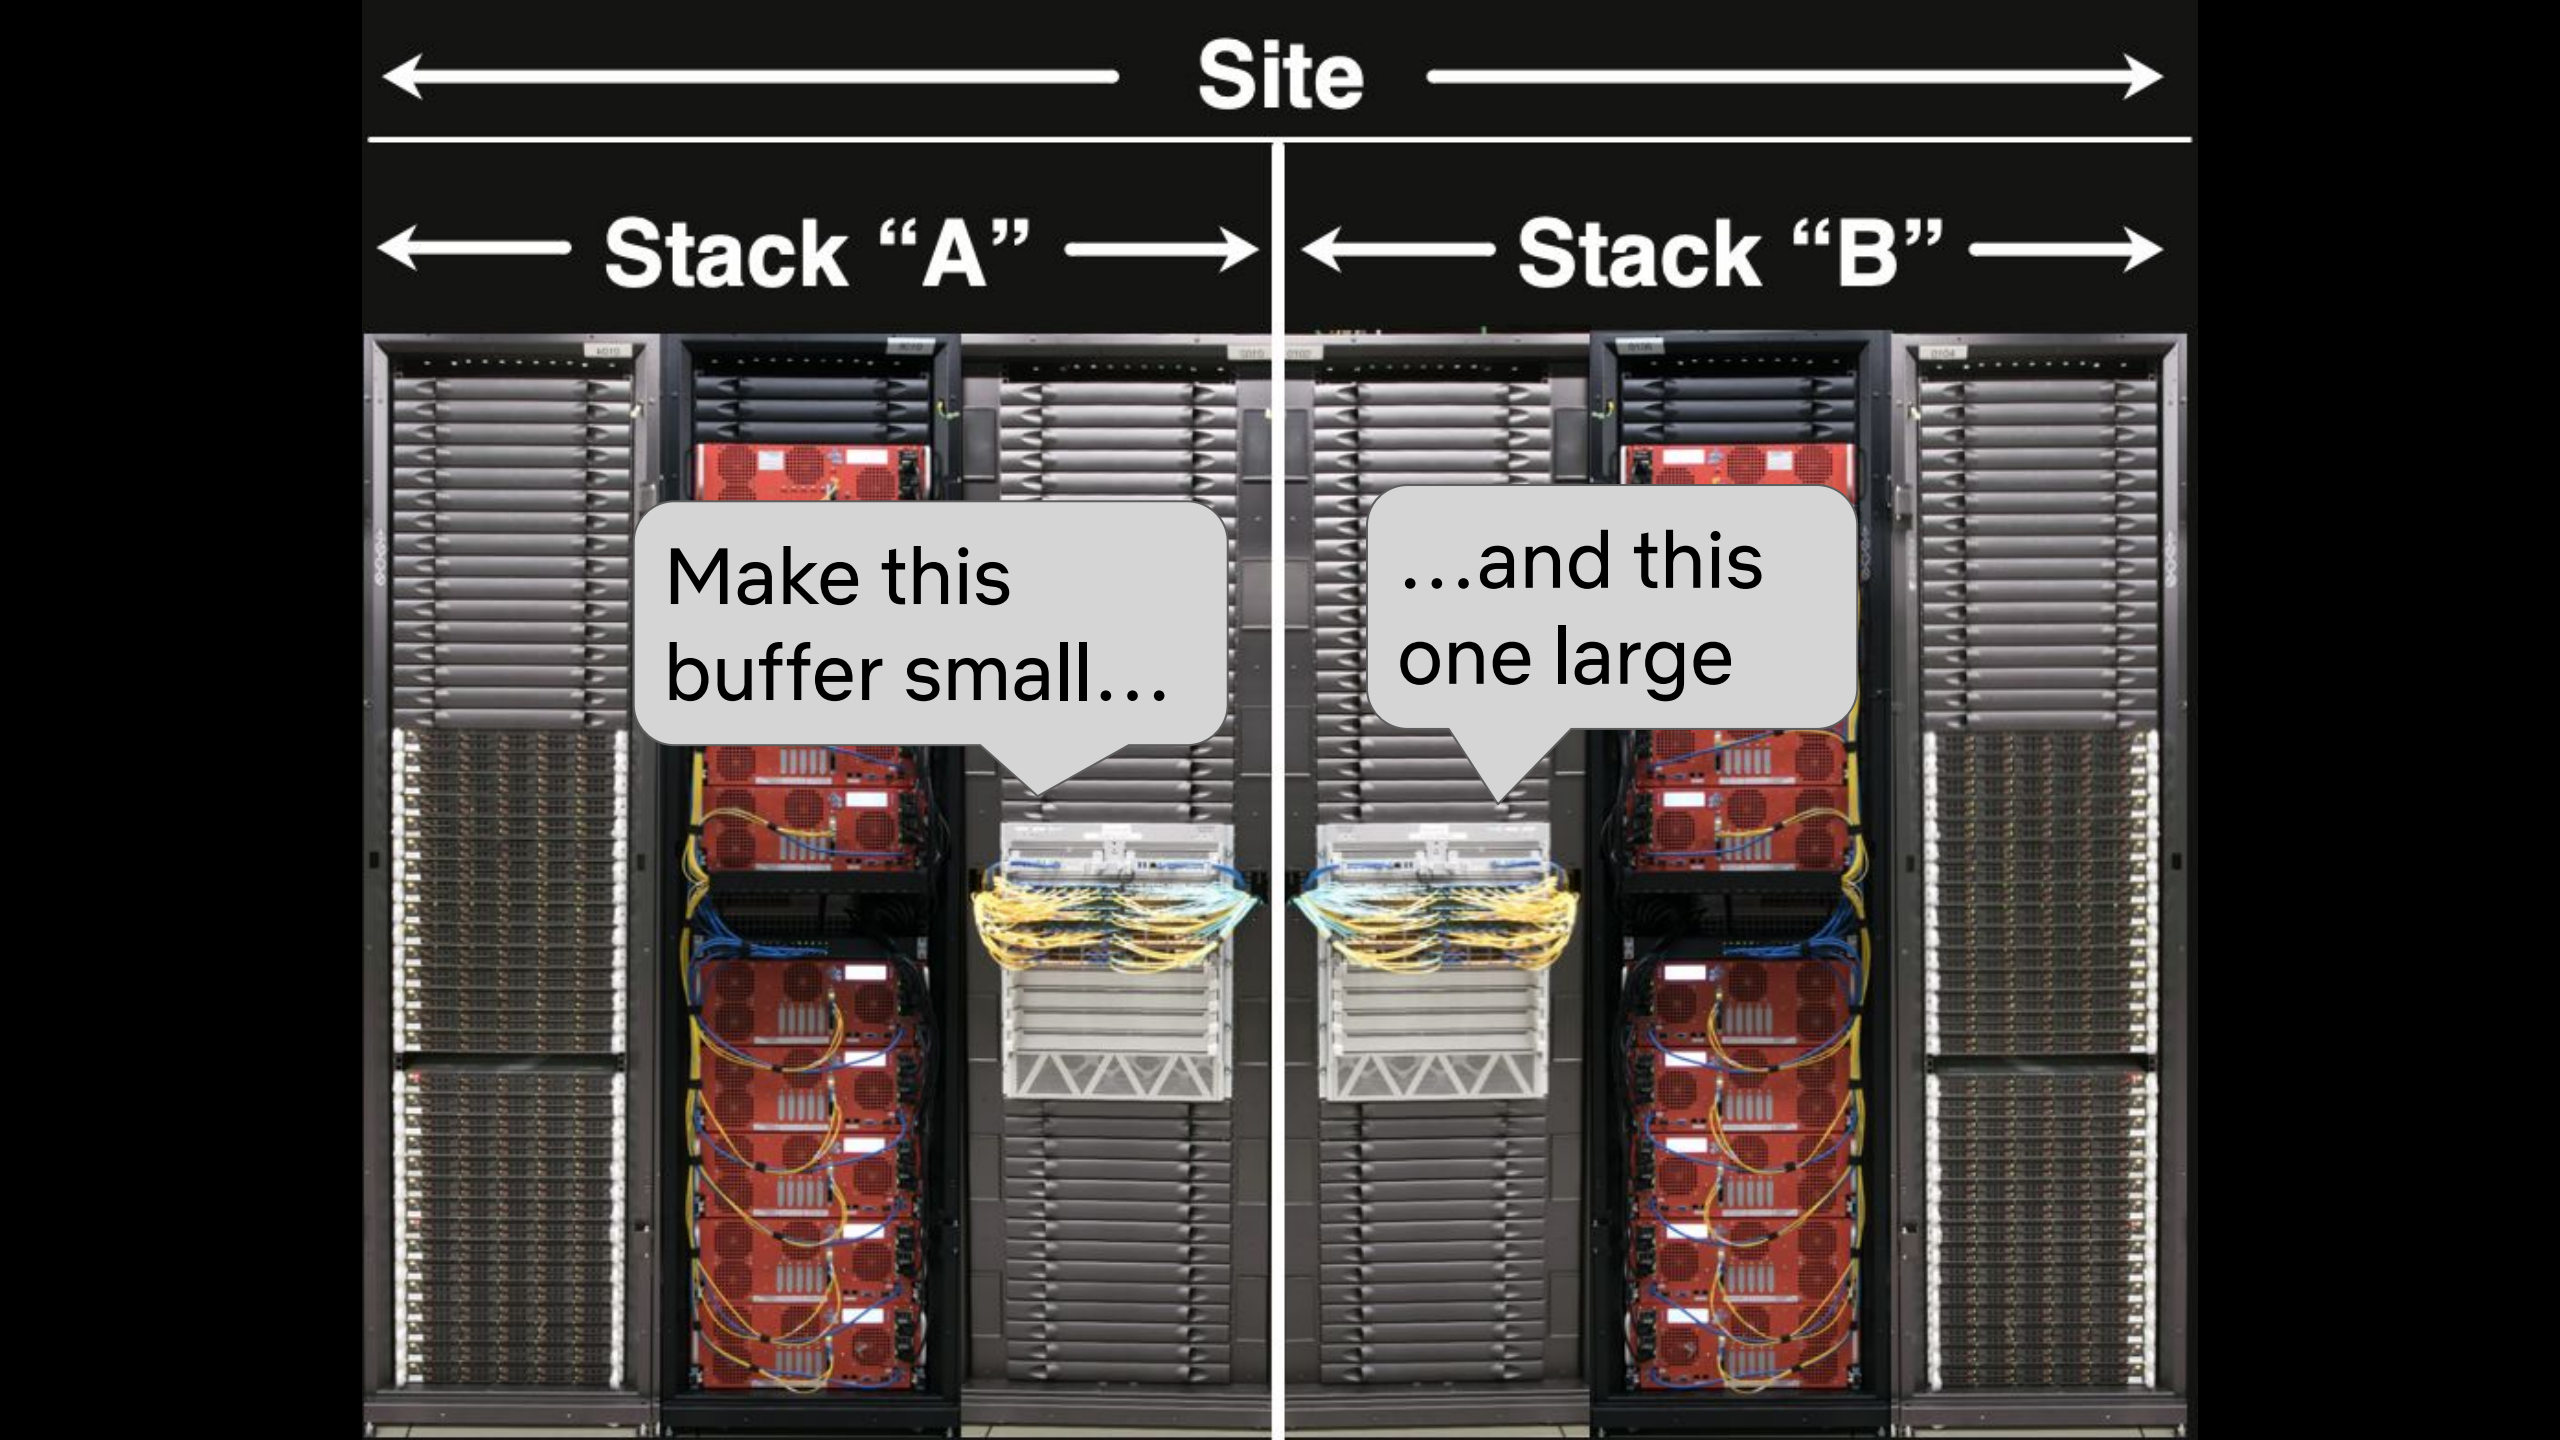
Site
Stack “A”
Stack “B”
Make this buffer small…
…and this one large

1. **TCP New Reno (mostly) behaves as expected**
2. Video performance varies
3. Real routers complicate this story

Large buffer has higher latency during congested hour

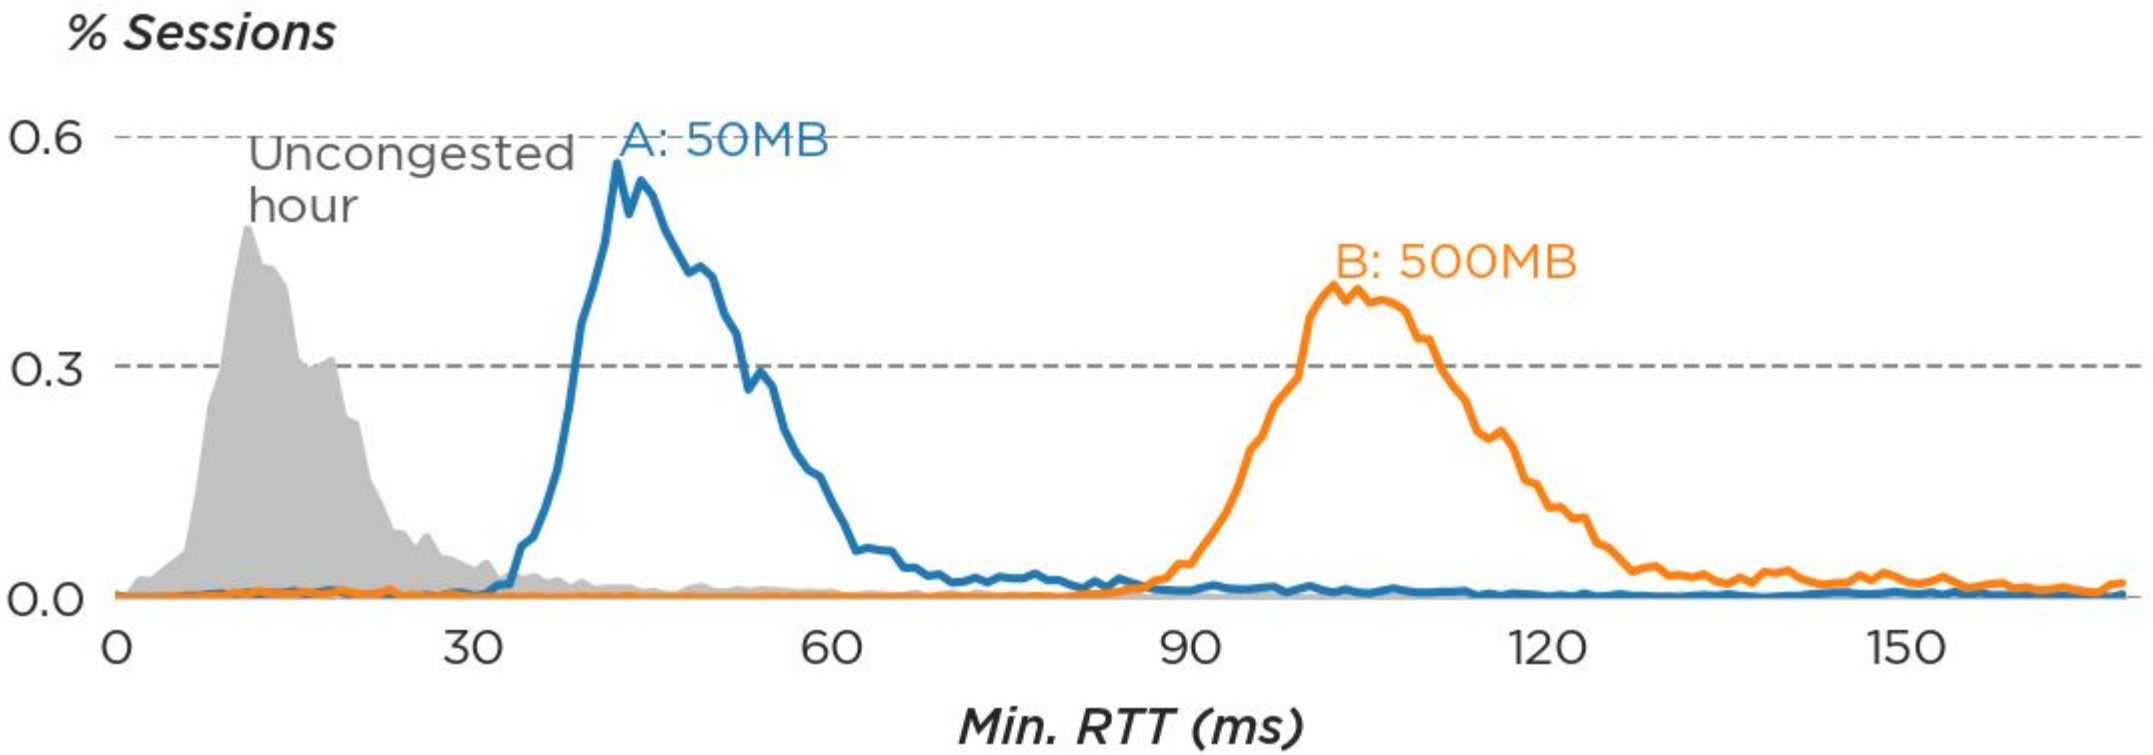

Sometimes the large buffer has much higher latency

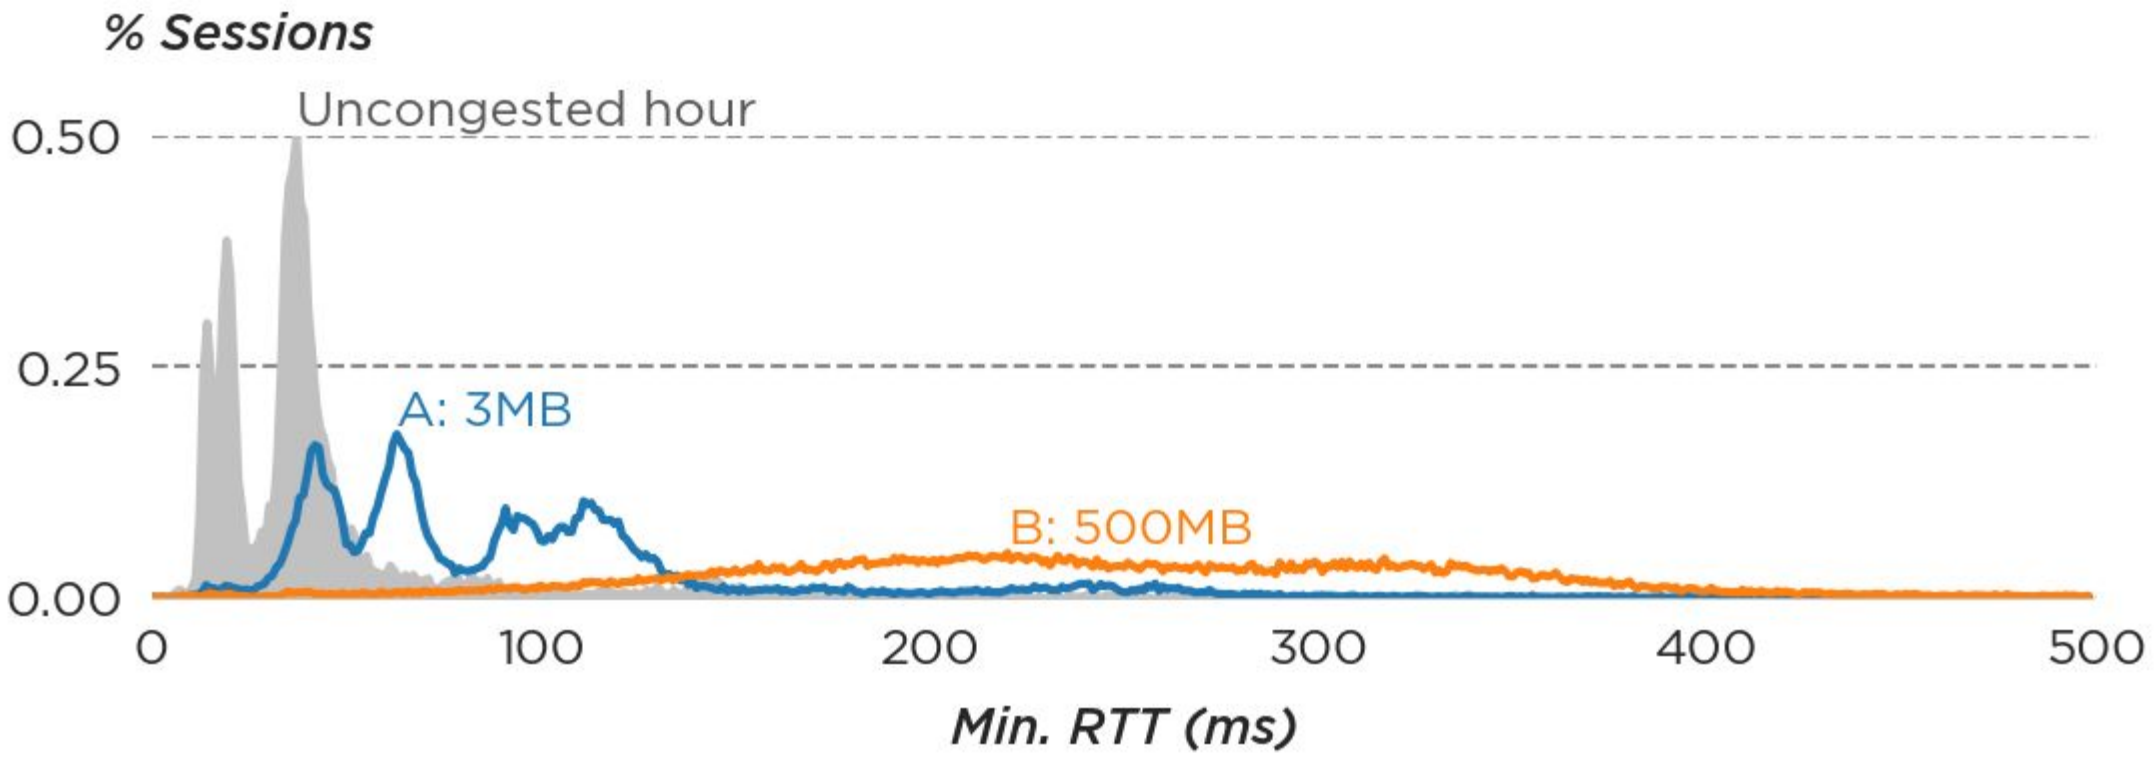

Large buffer has lower loss during congested hour

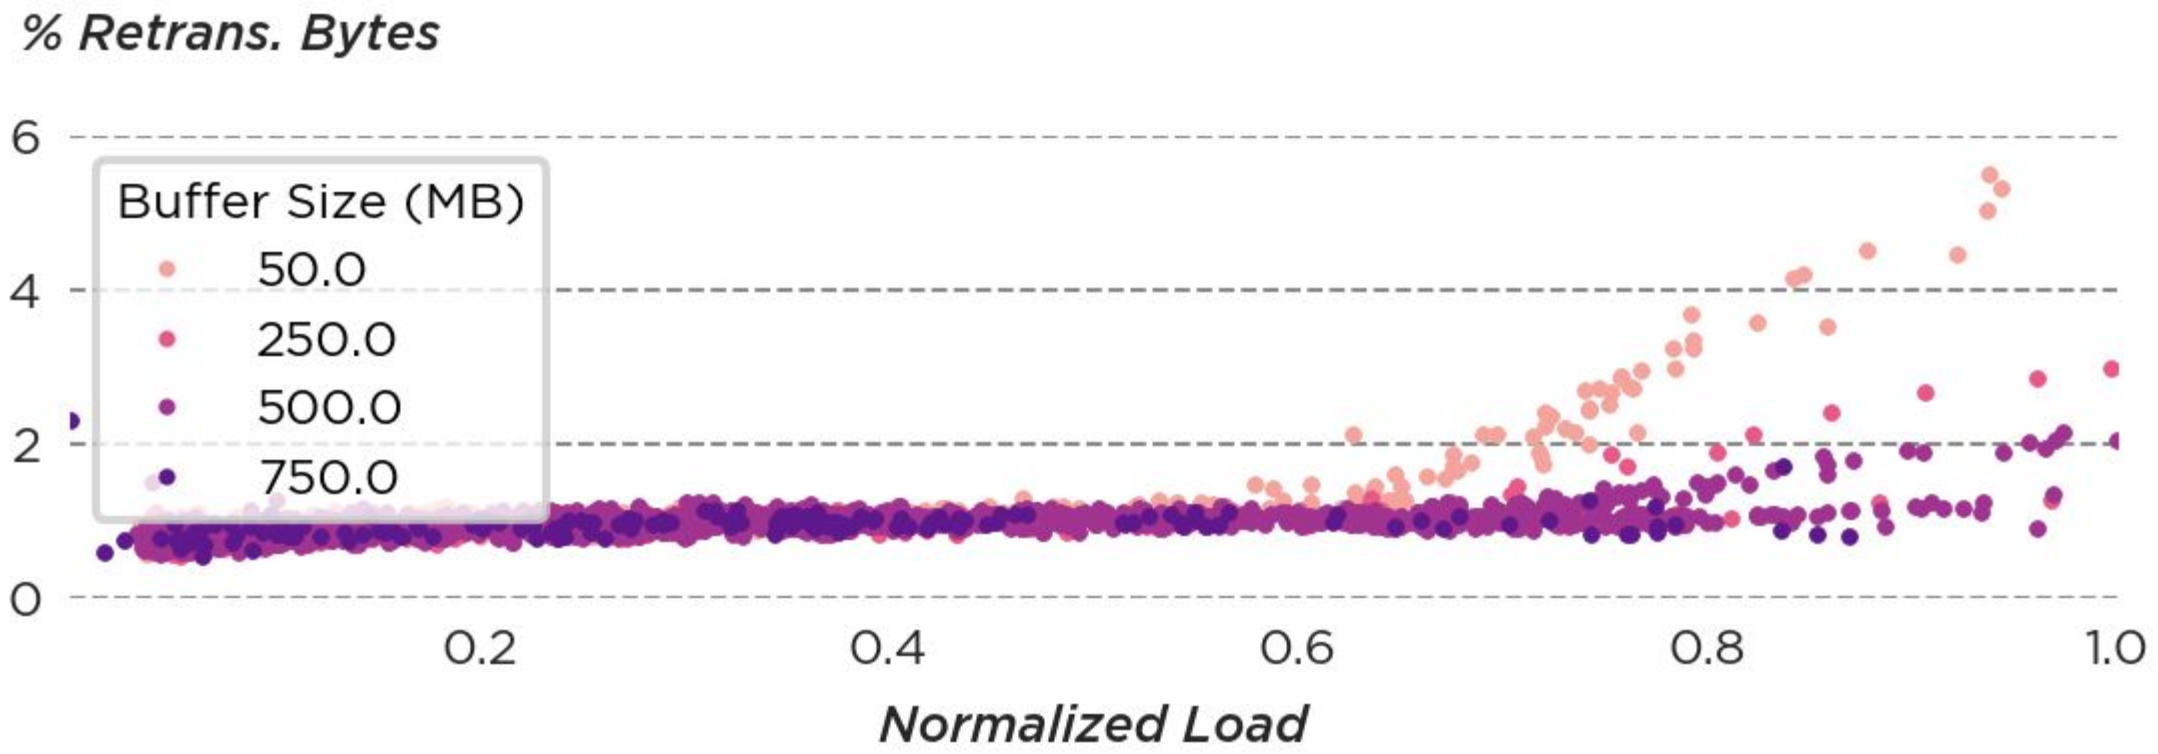

1. TCP New Reno (mostly) behaves as expected
2. **Video performance varies**
3. Real routers complicate this story

Good buffer size:

+ Fewer rebuffers
+ Better video quality
+ Videos start faster

Bad buffer size:

- More rebuffers
- Worse video quality
- Videos start slower

Good buffer size:

+ Fewer rebuffers
+ Better video quality
+ Videos start faster

Bad buffer size:

- More rebuffers
- Worse video quality
- Videos start slower

} This happens when buffer is **too large** or **too small.**

Site #2: A smaller buffer is better

Reducing the buffer from **500MB** to **25MB**

- **-15.6%** decrease in sessions with a rebuffer
- **-5.3%** decrease in low quality video
- **-13.5%** decrease in play delay

Site #3: A smaller buffer is better

Reducing the buffer from **500MB** to **50MB**

- **-22.1%** decrease in sessions with a rebuffer
- **-7.0%** decrease in low quality video
- **-14.8%** decrease in play delay

Site #1: A smaller buffer is worse

Reducing the buffer from **500MB** to **50MB**

- **+46.3%** increase in sessions with a rebuffer
- **+5.7%** increase in low quality video
- **-5.9%** decrease in play delay

1. TCP New Reno (mostly) behaves as expected
2. Video performance varies
3. **Real routers complicate this story**

Large buffer has higher latency during congested hour

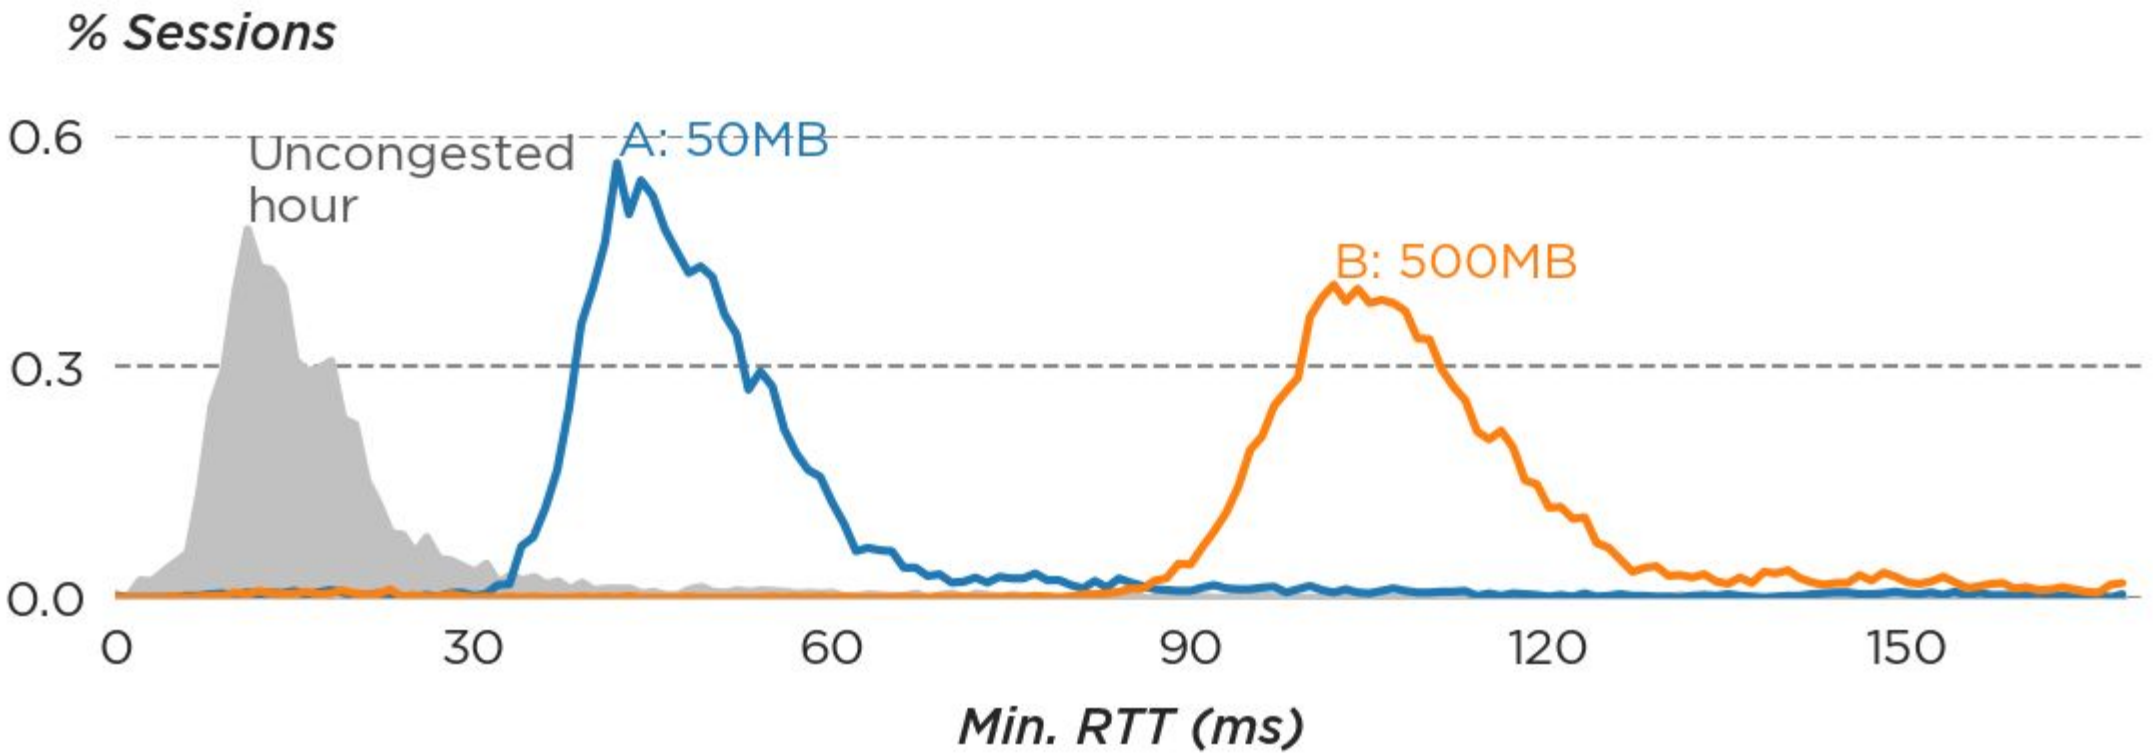

Remember how the large buffer has much higher latency...

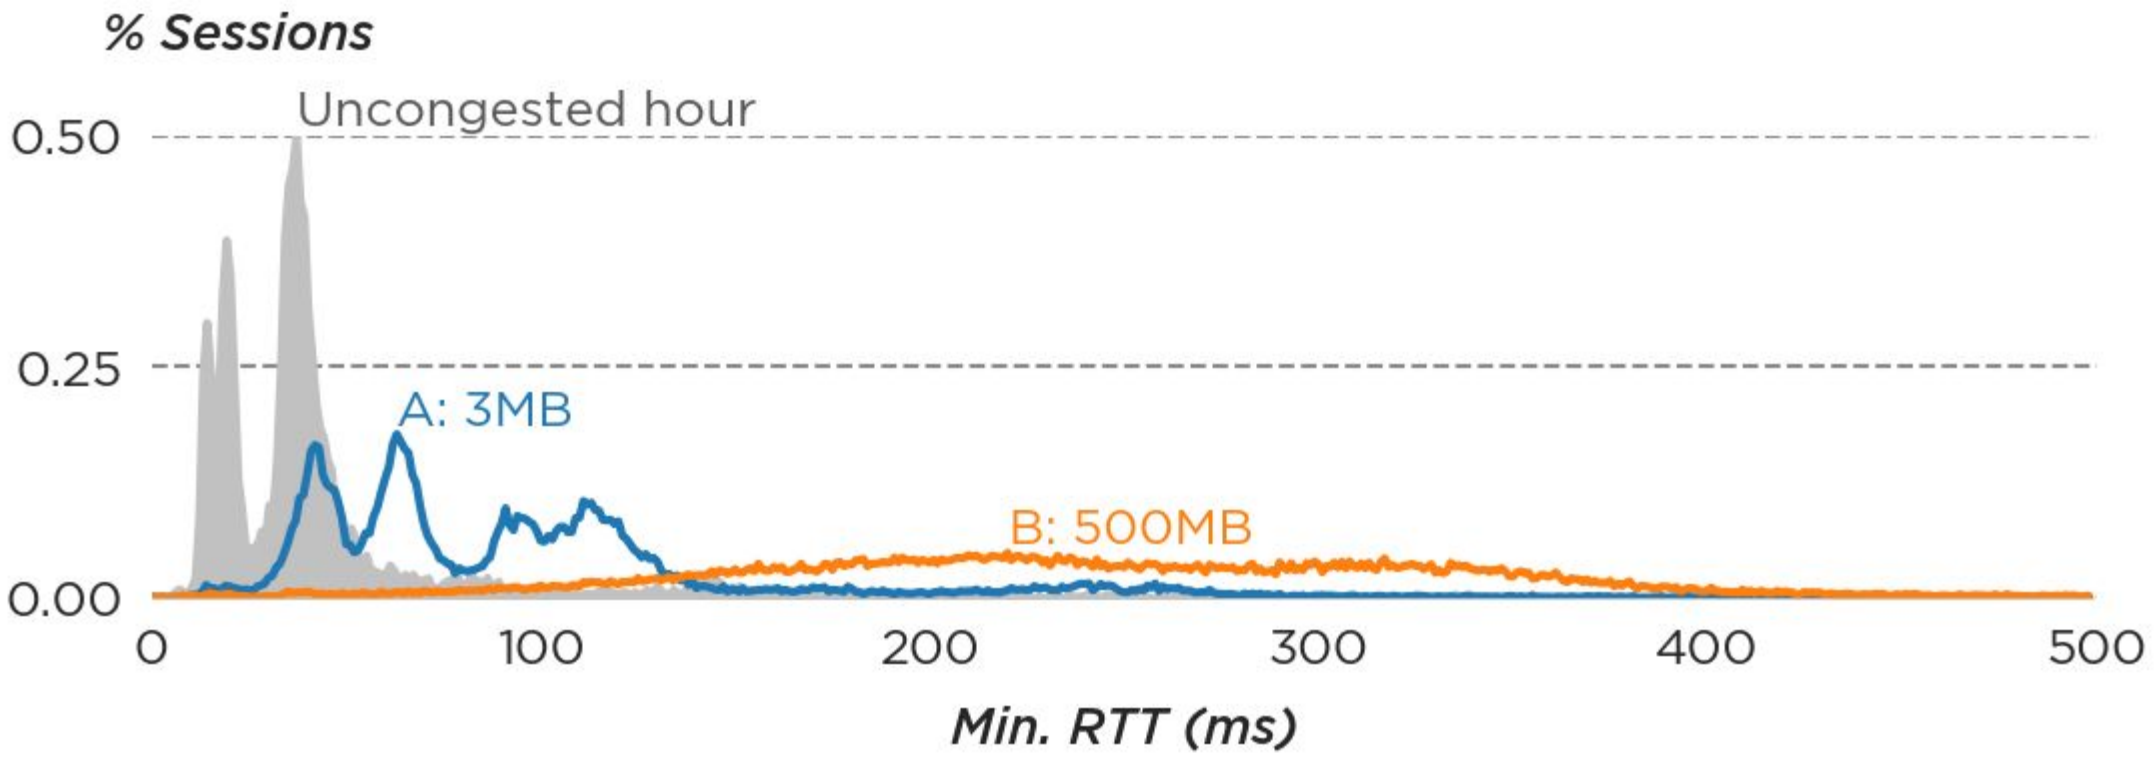

Servers have different very latency distributions

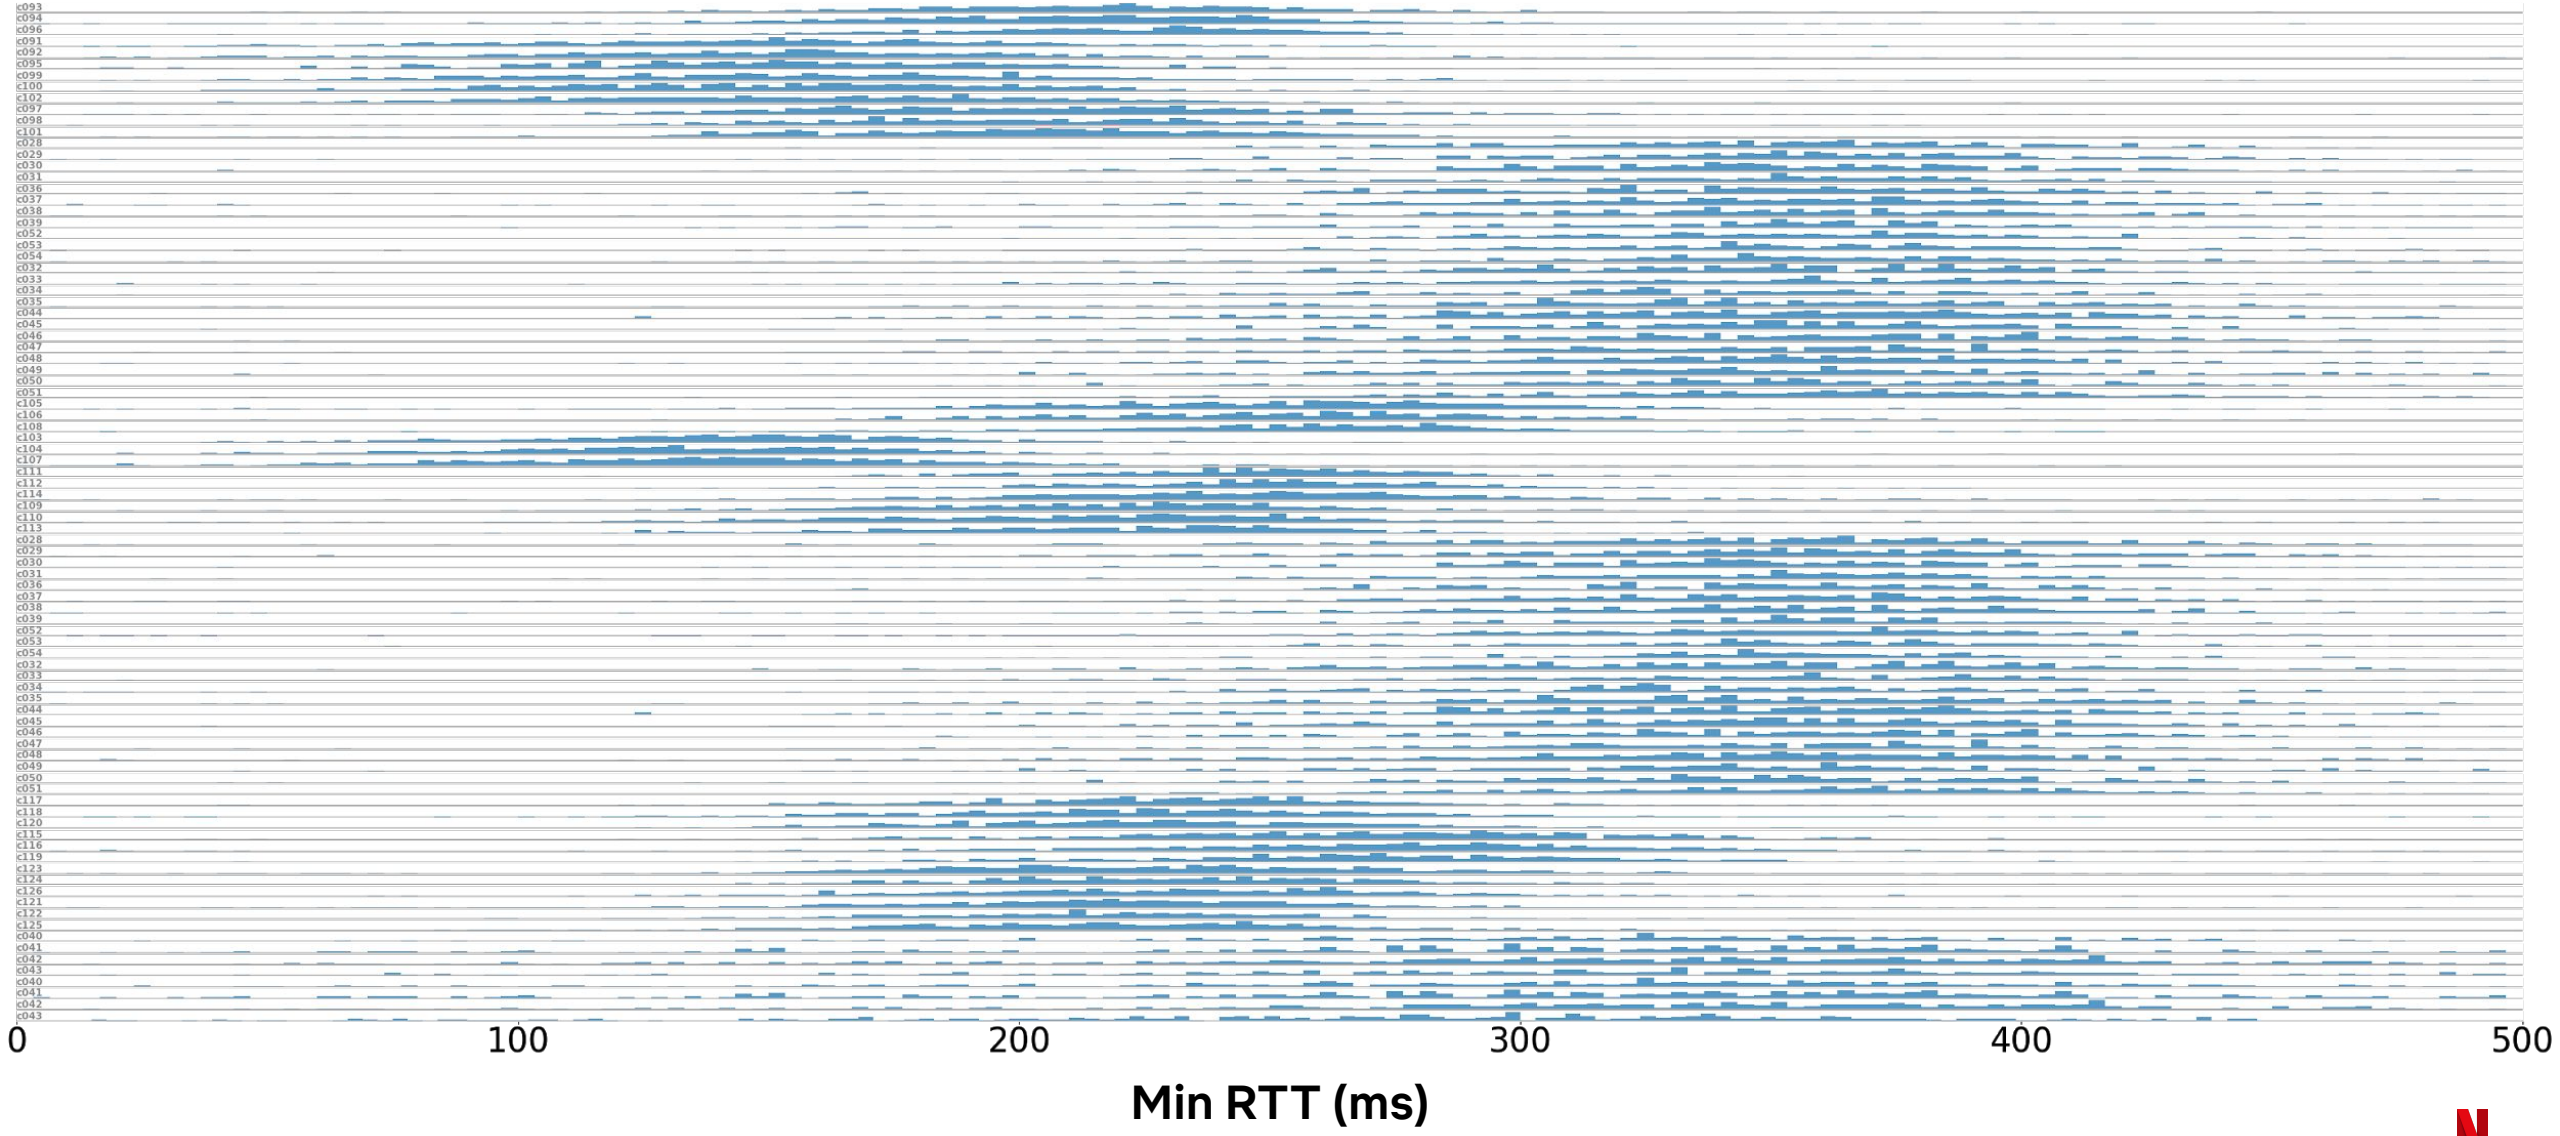

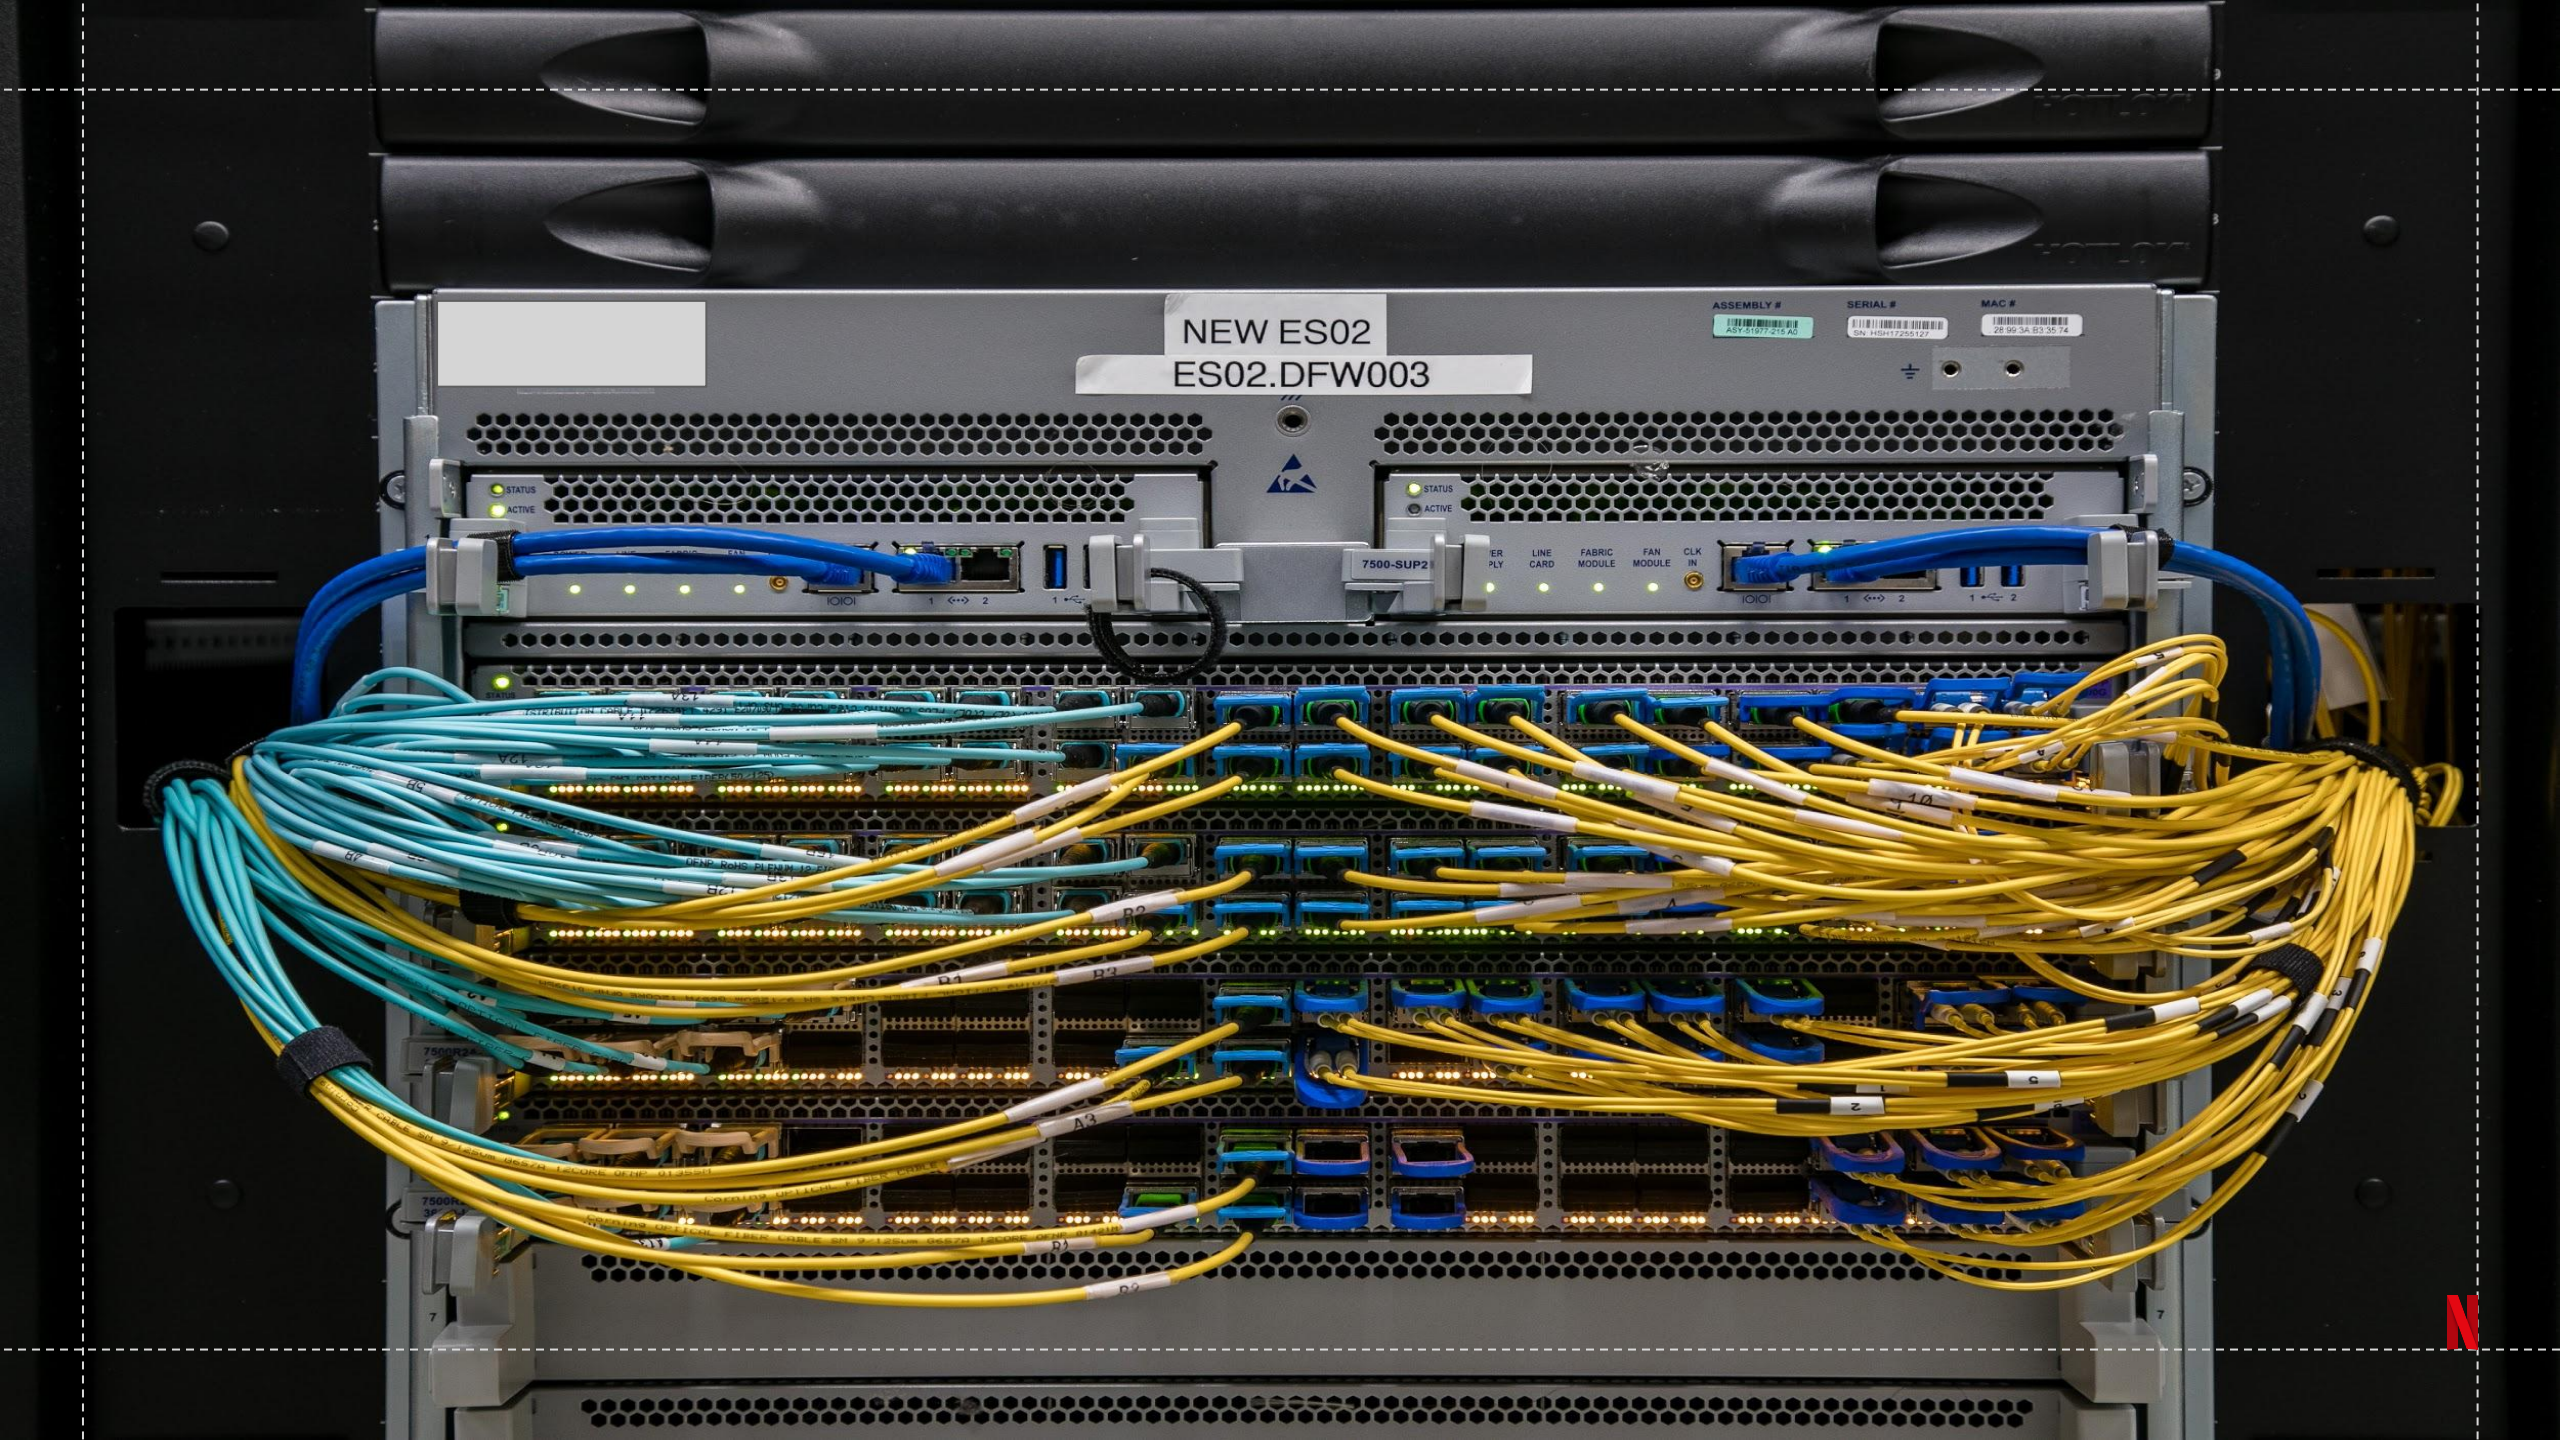
NEW ES02
ES02.DFW003
ASSEMBLY #
SERIAL #
MAC #
STATUS
ACTIVE
7500-SUP2
LINE CARD
FABRIC MODULE
FAN MODULE
CLK IN

What I imagined

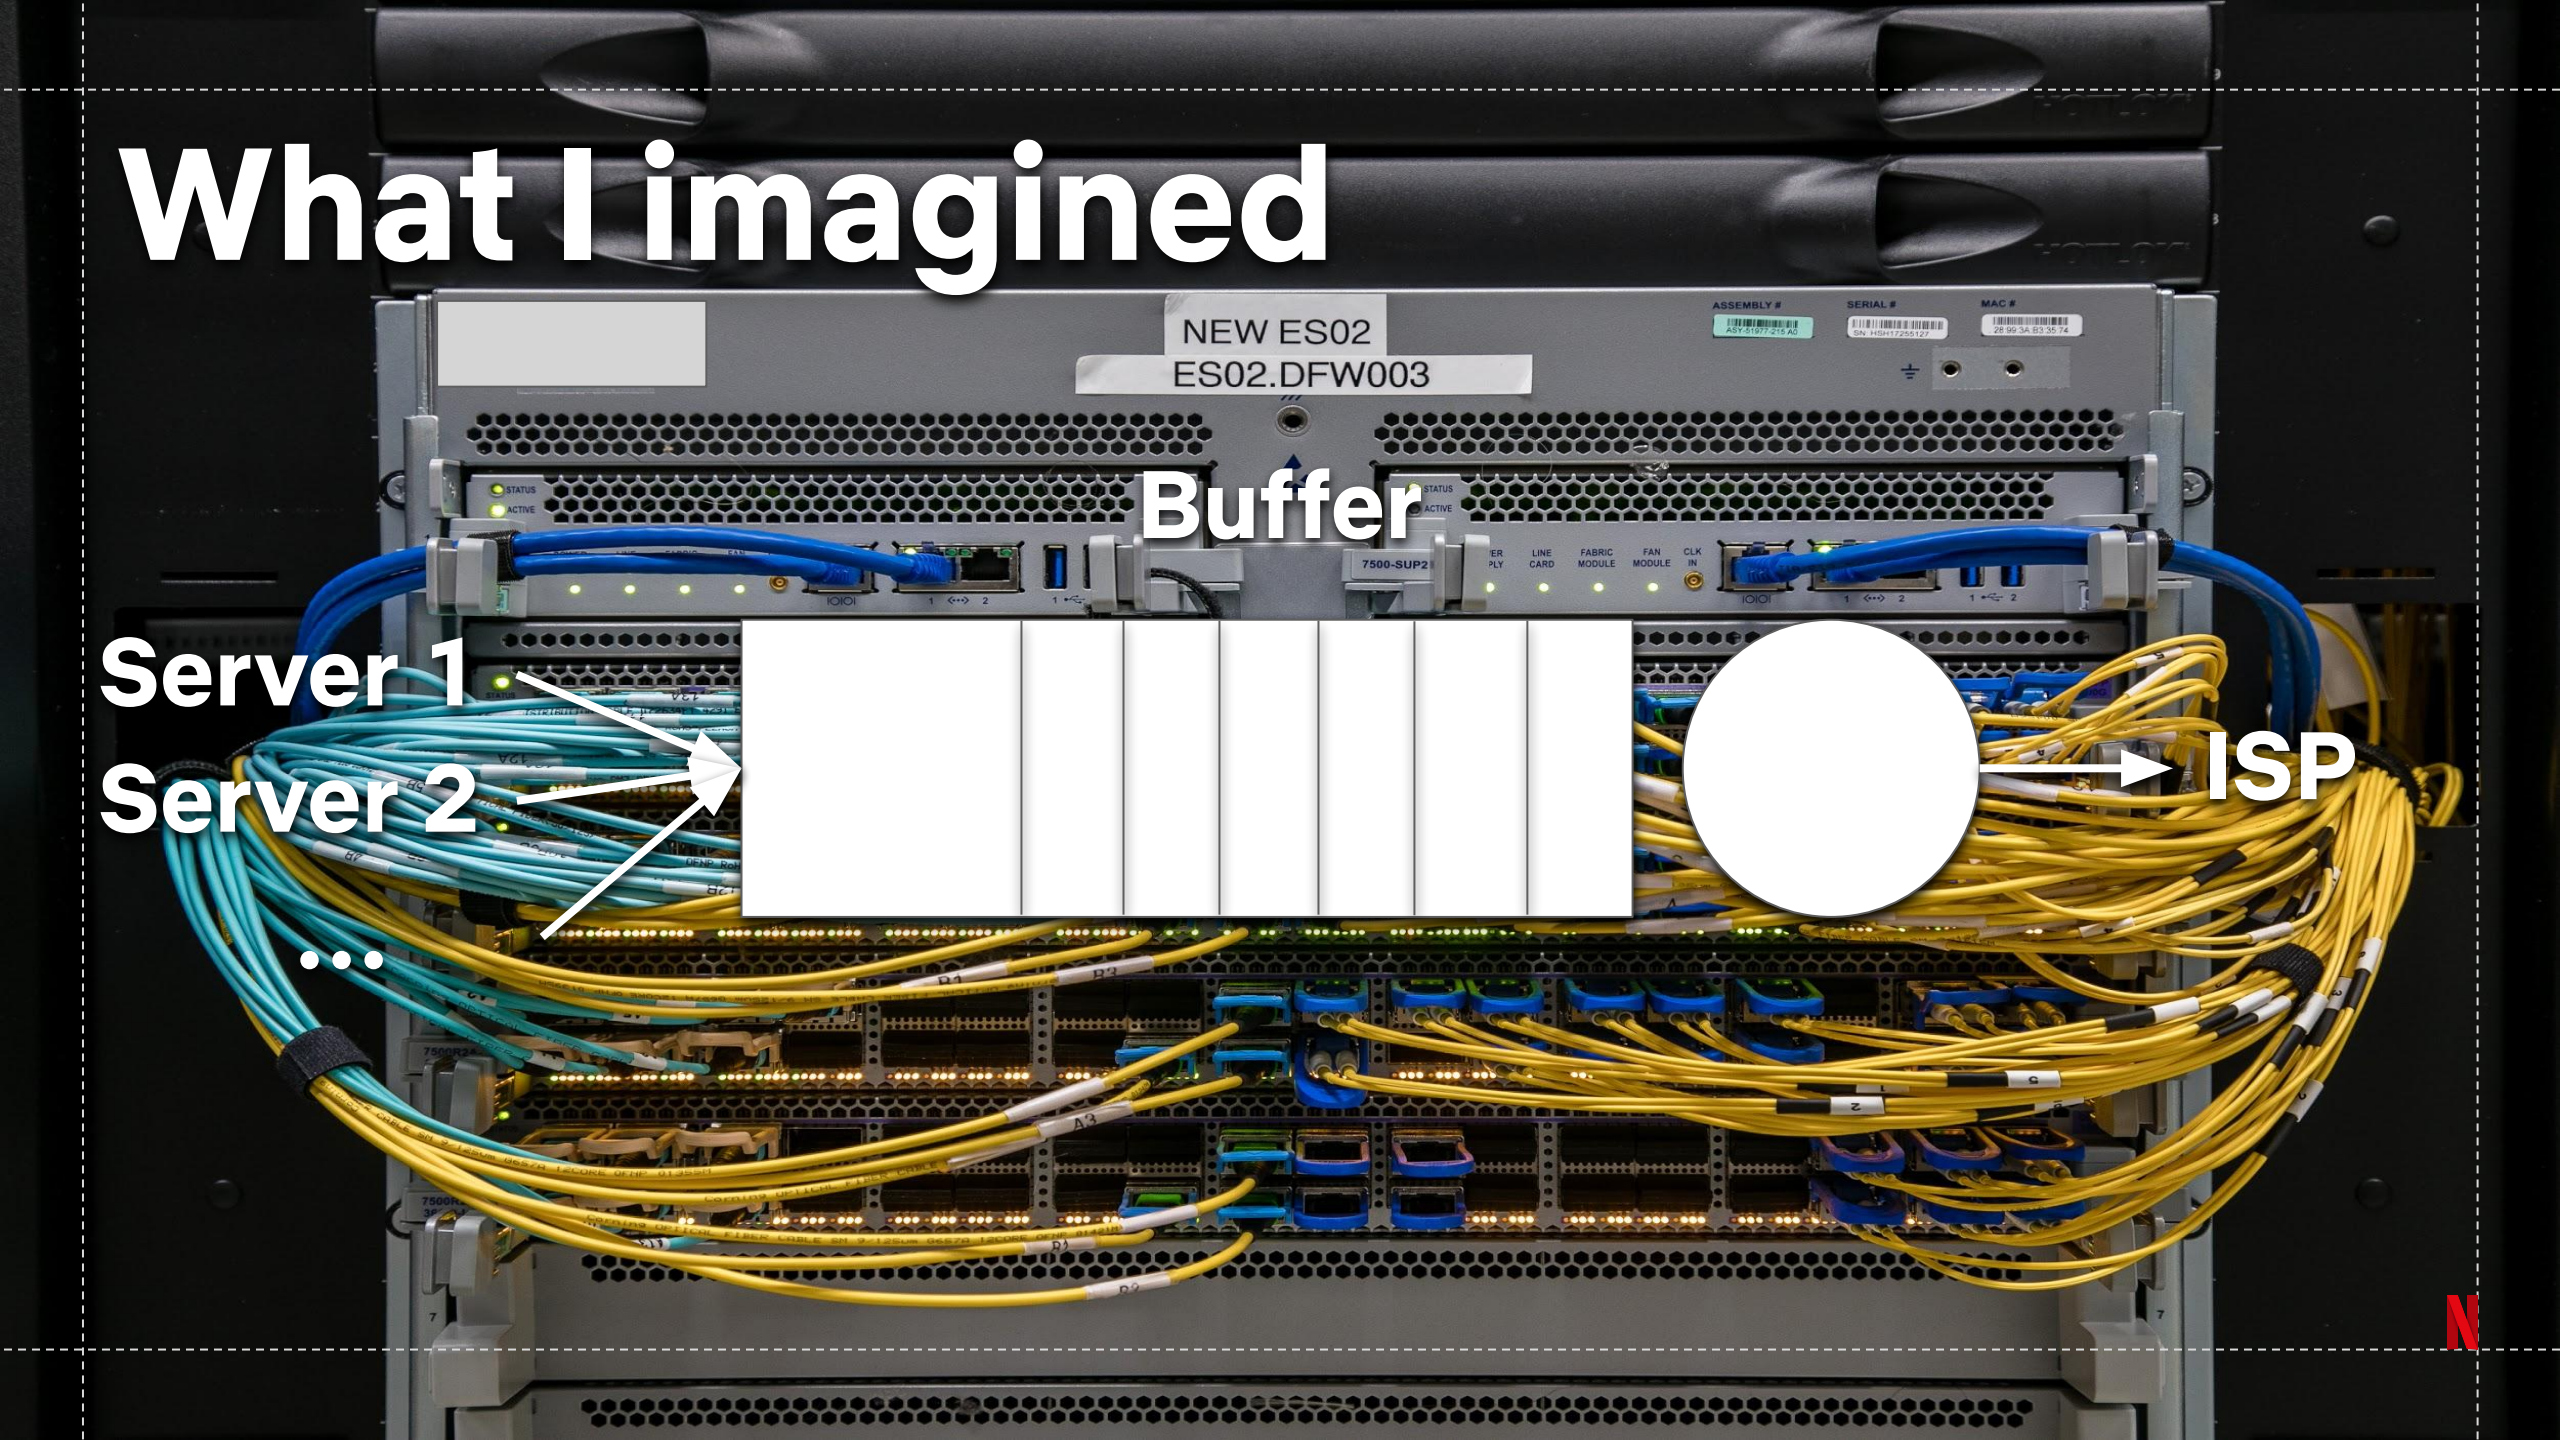

What I imagined

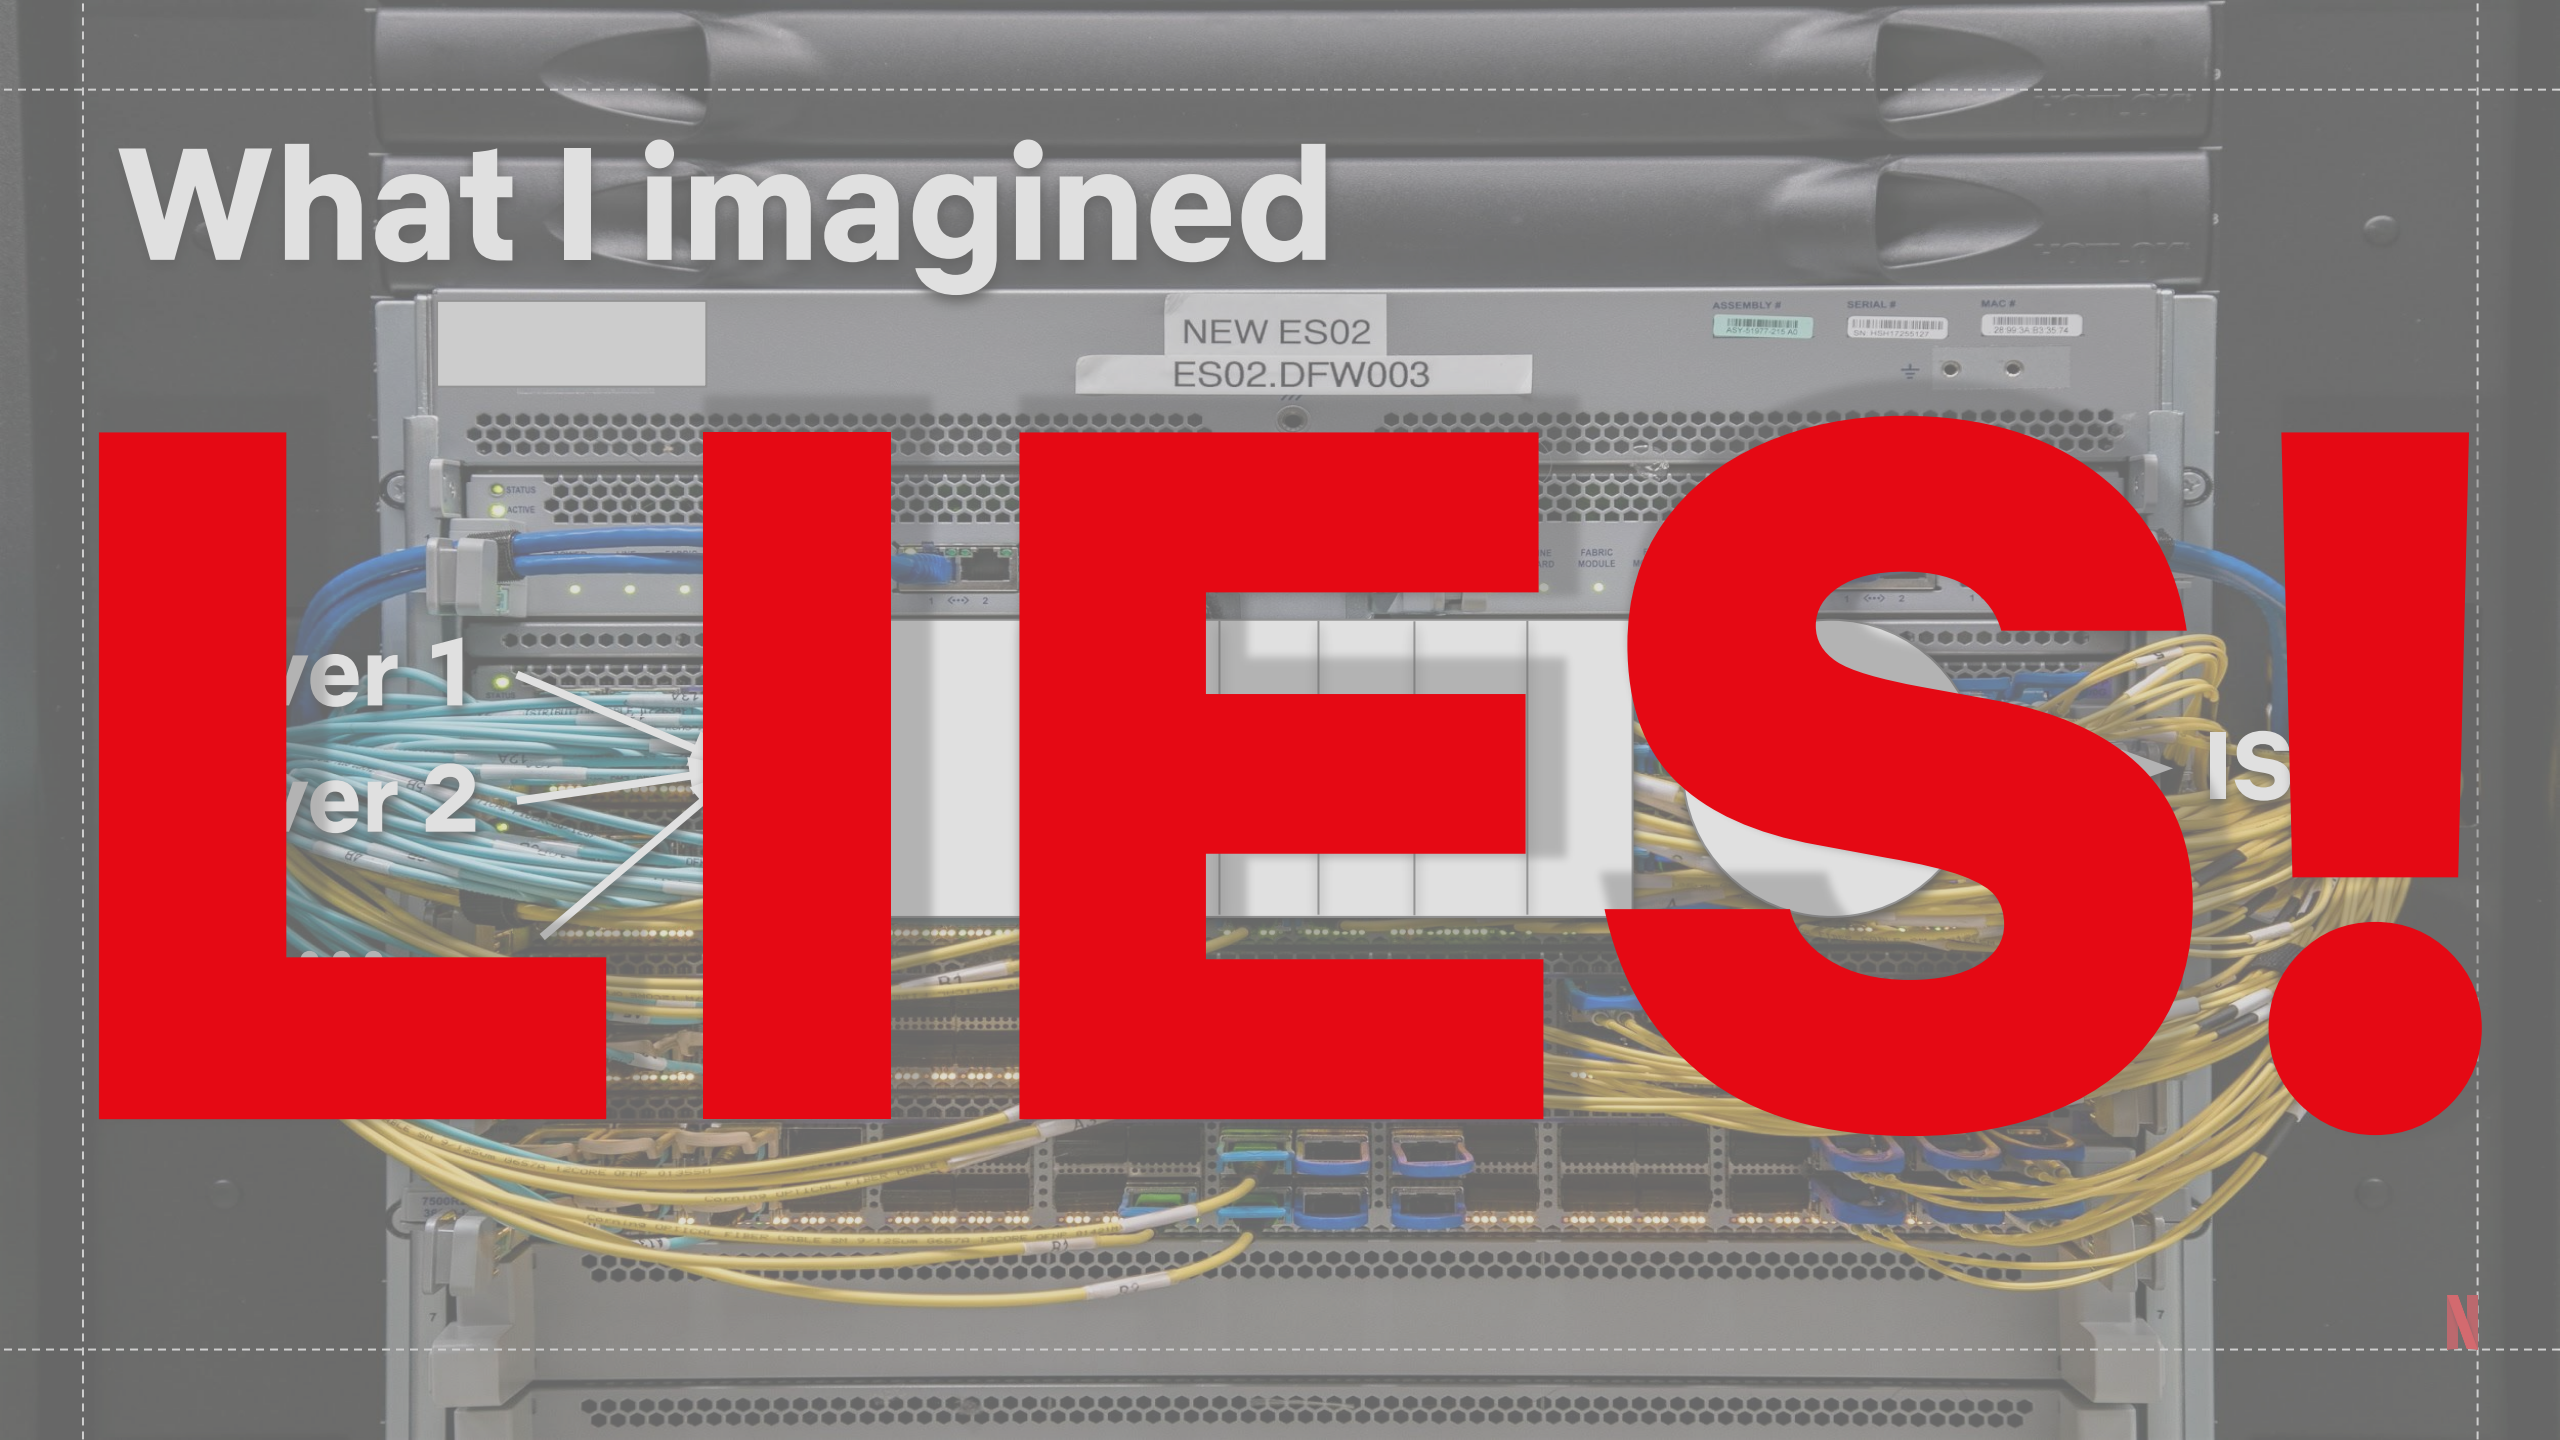

L.I.E.S.!

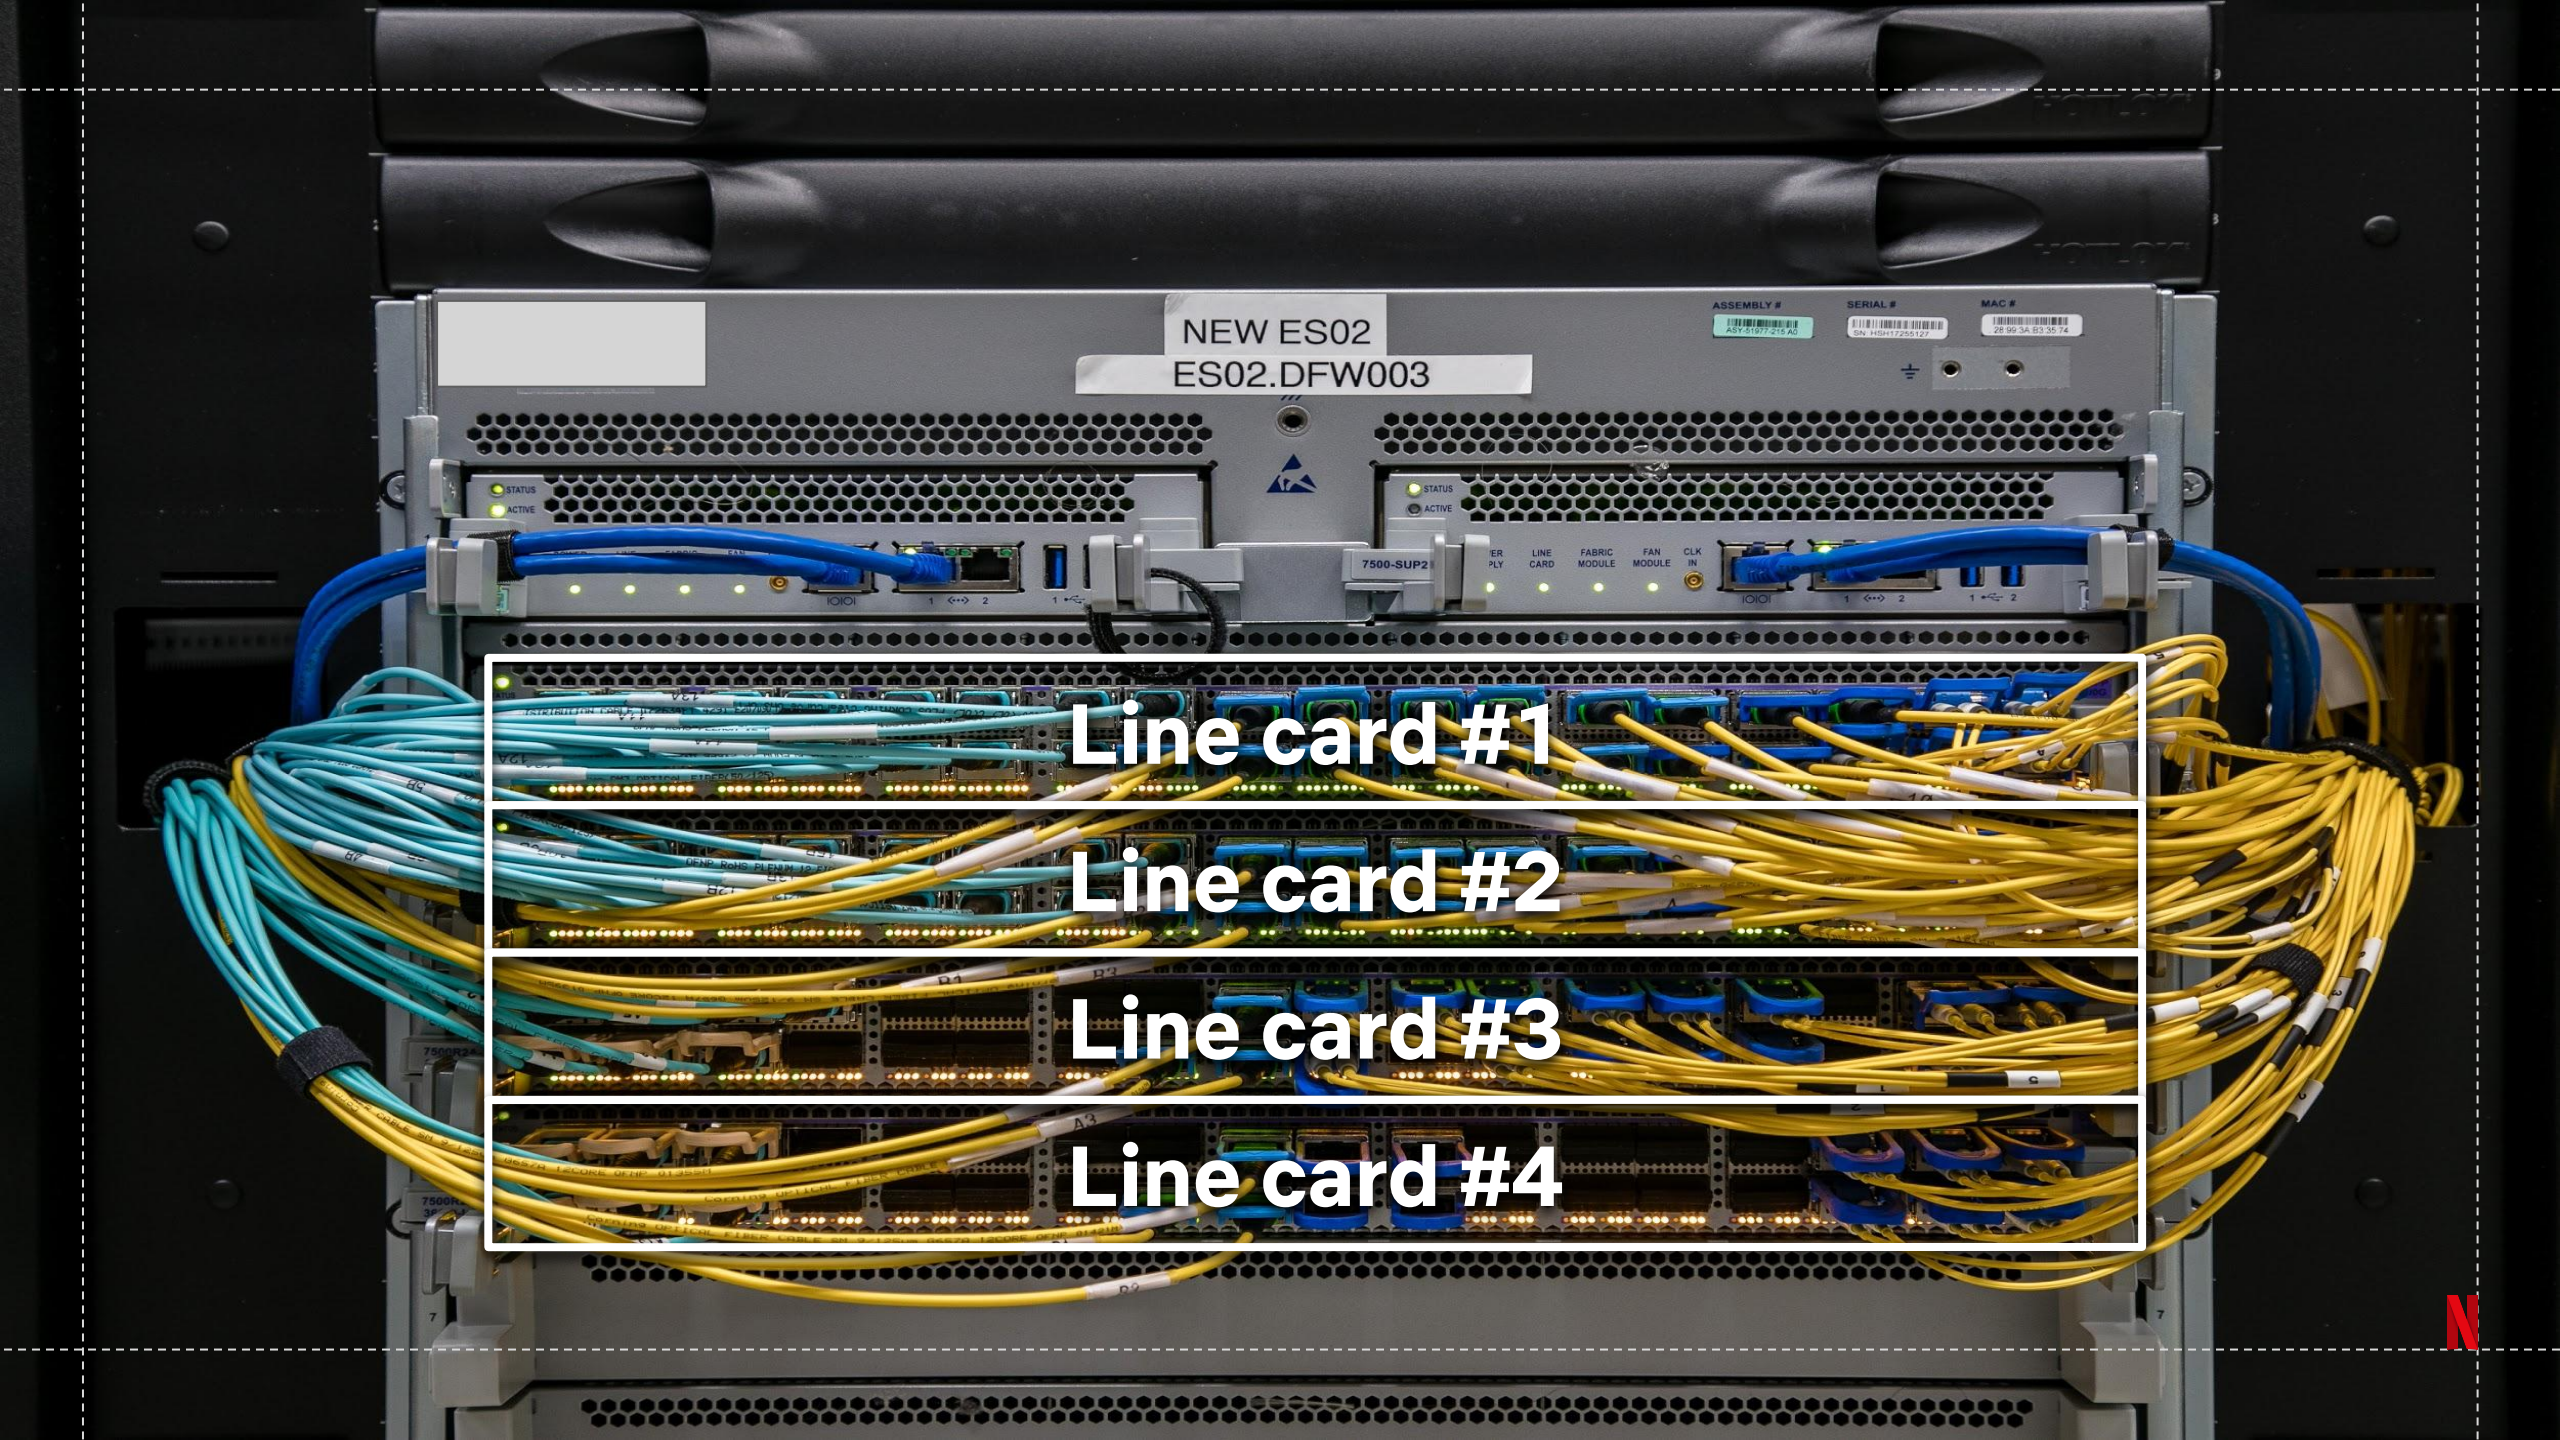
NEW ES02
ES02.DFW003
Line card #1
Line card #2
Line card #3
Line card #4

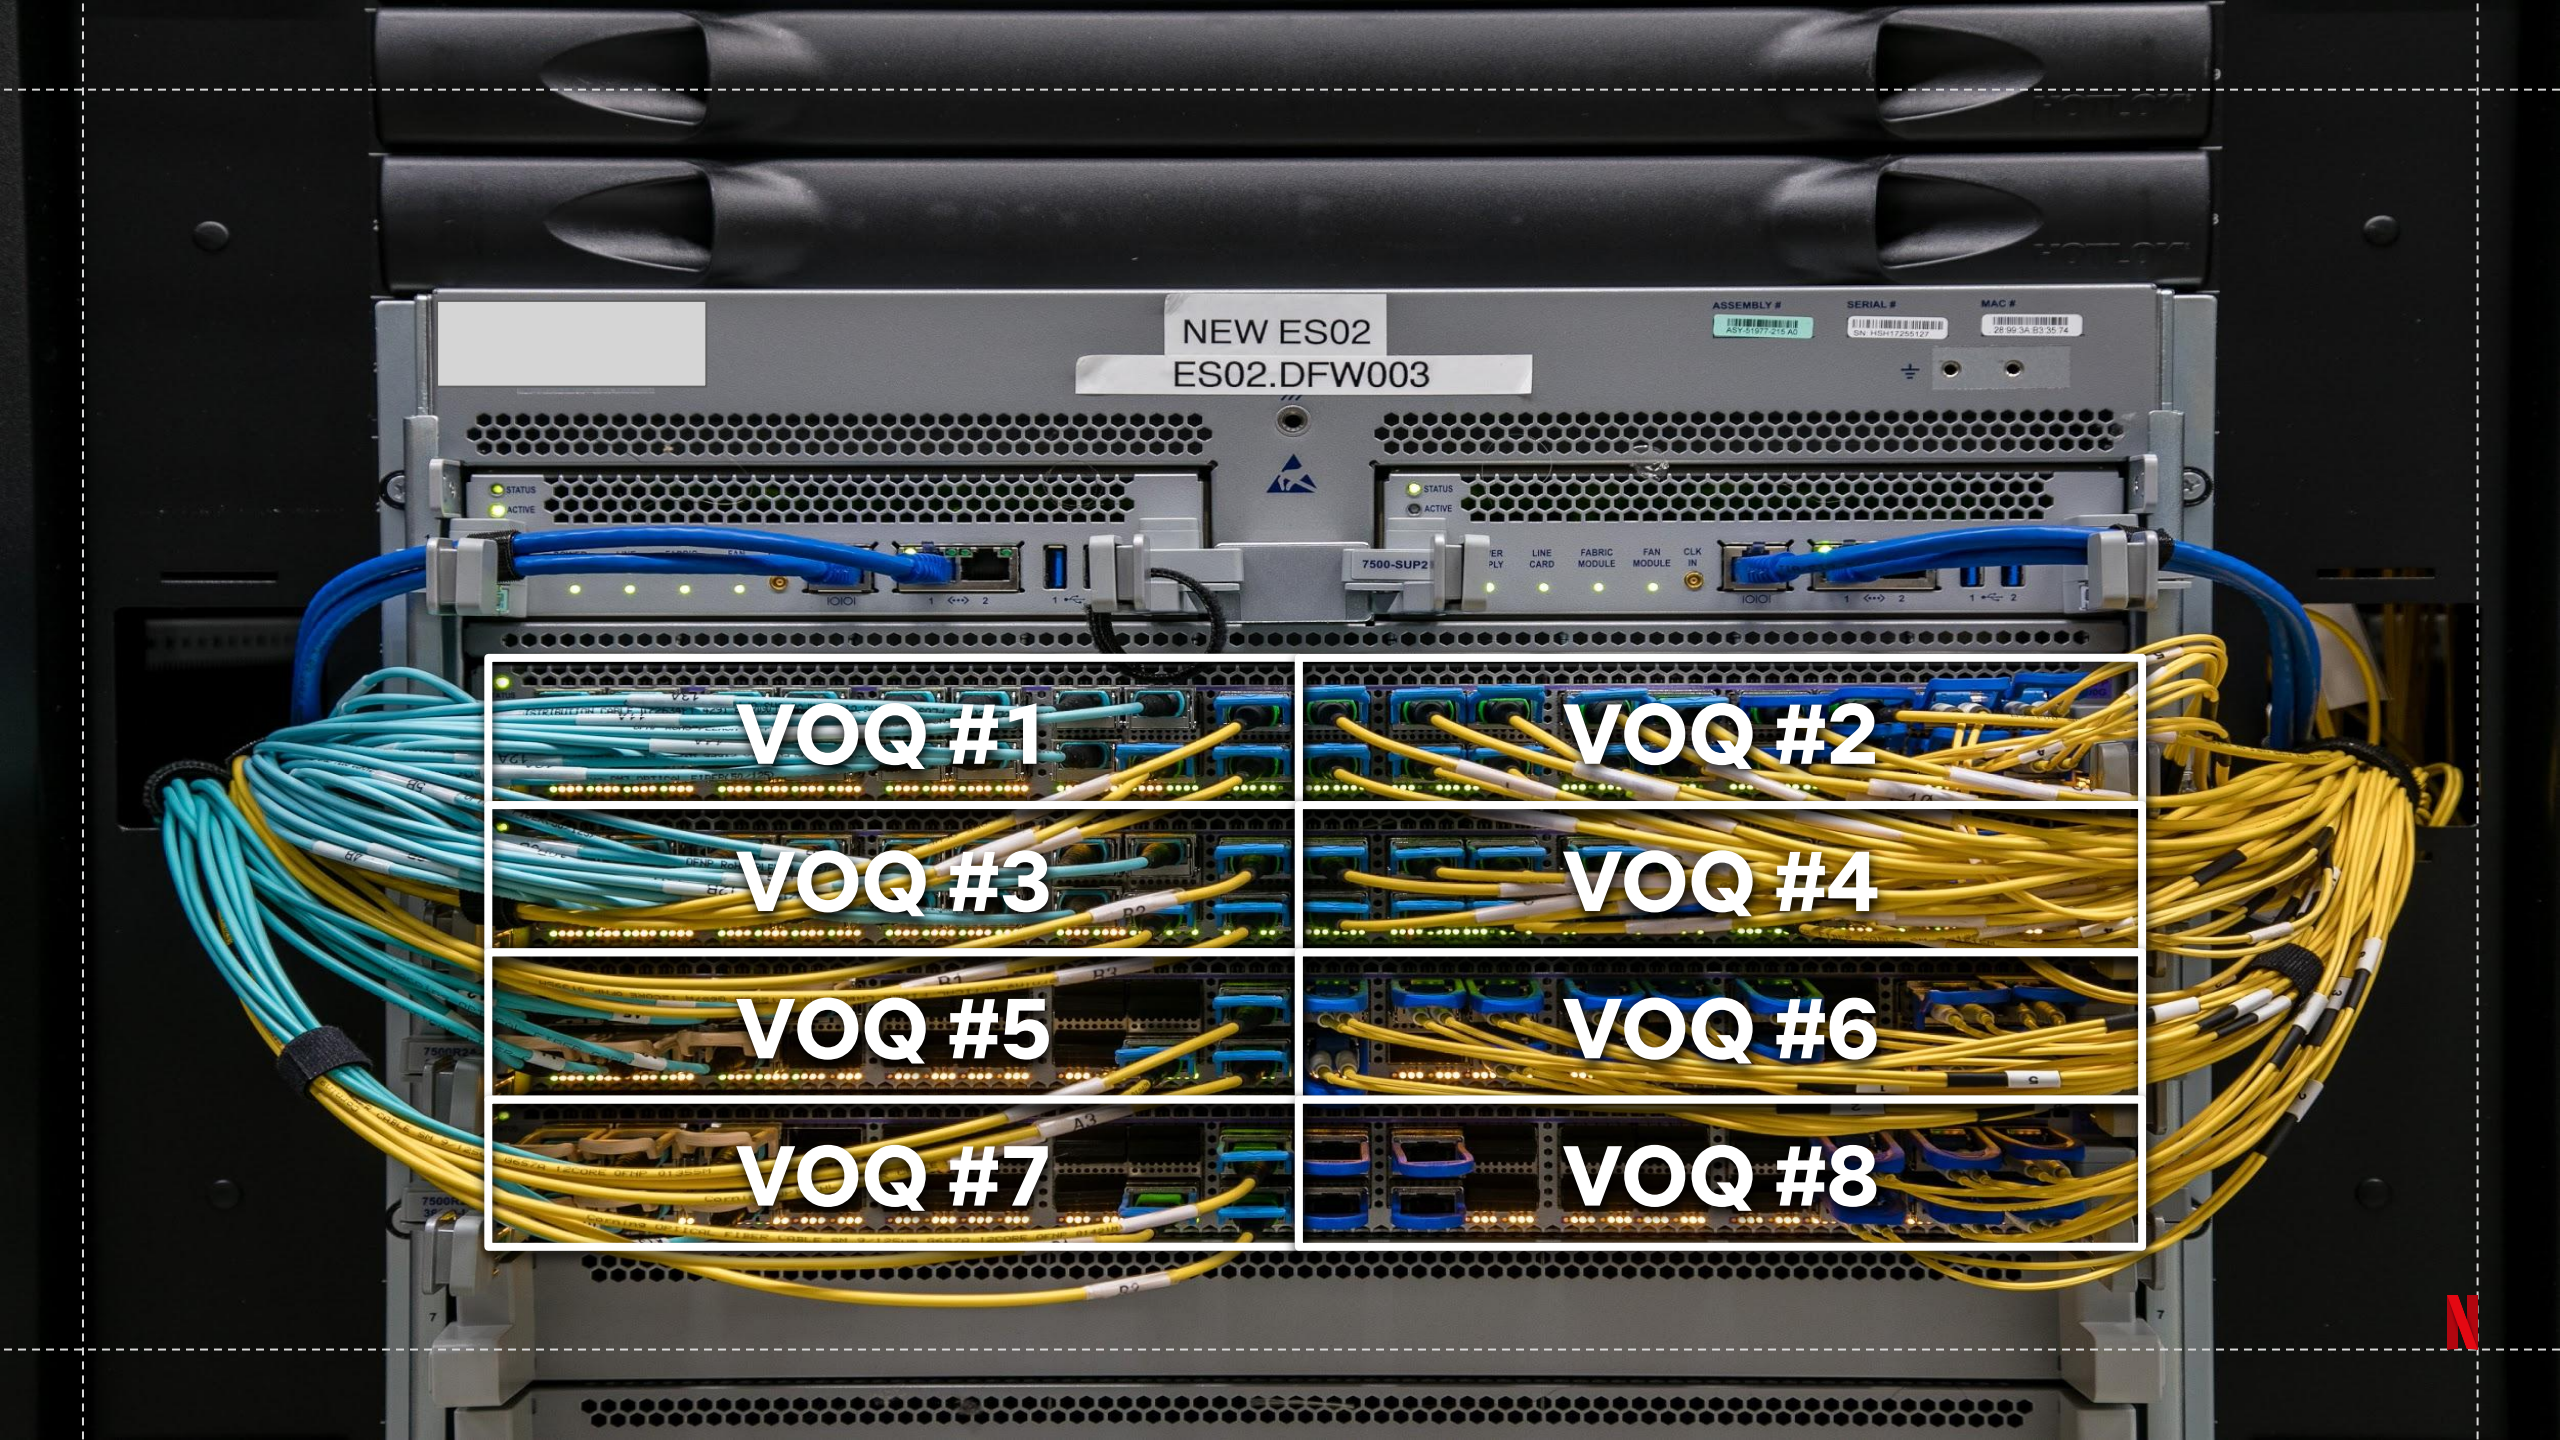
NEW ES02
ES02.DFW003
VOQ #1
VOQ #2
VOQ #3
VOQ #4
VOQ #5
VOQ #6
VOQ #7
VOQ #8

Buffer architecture

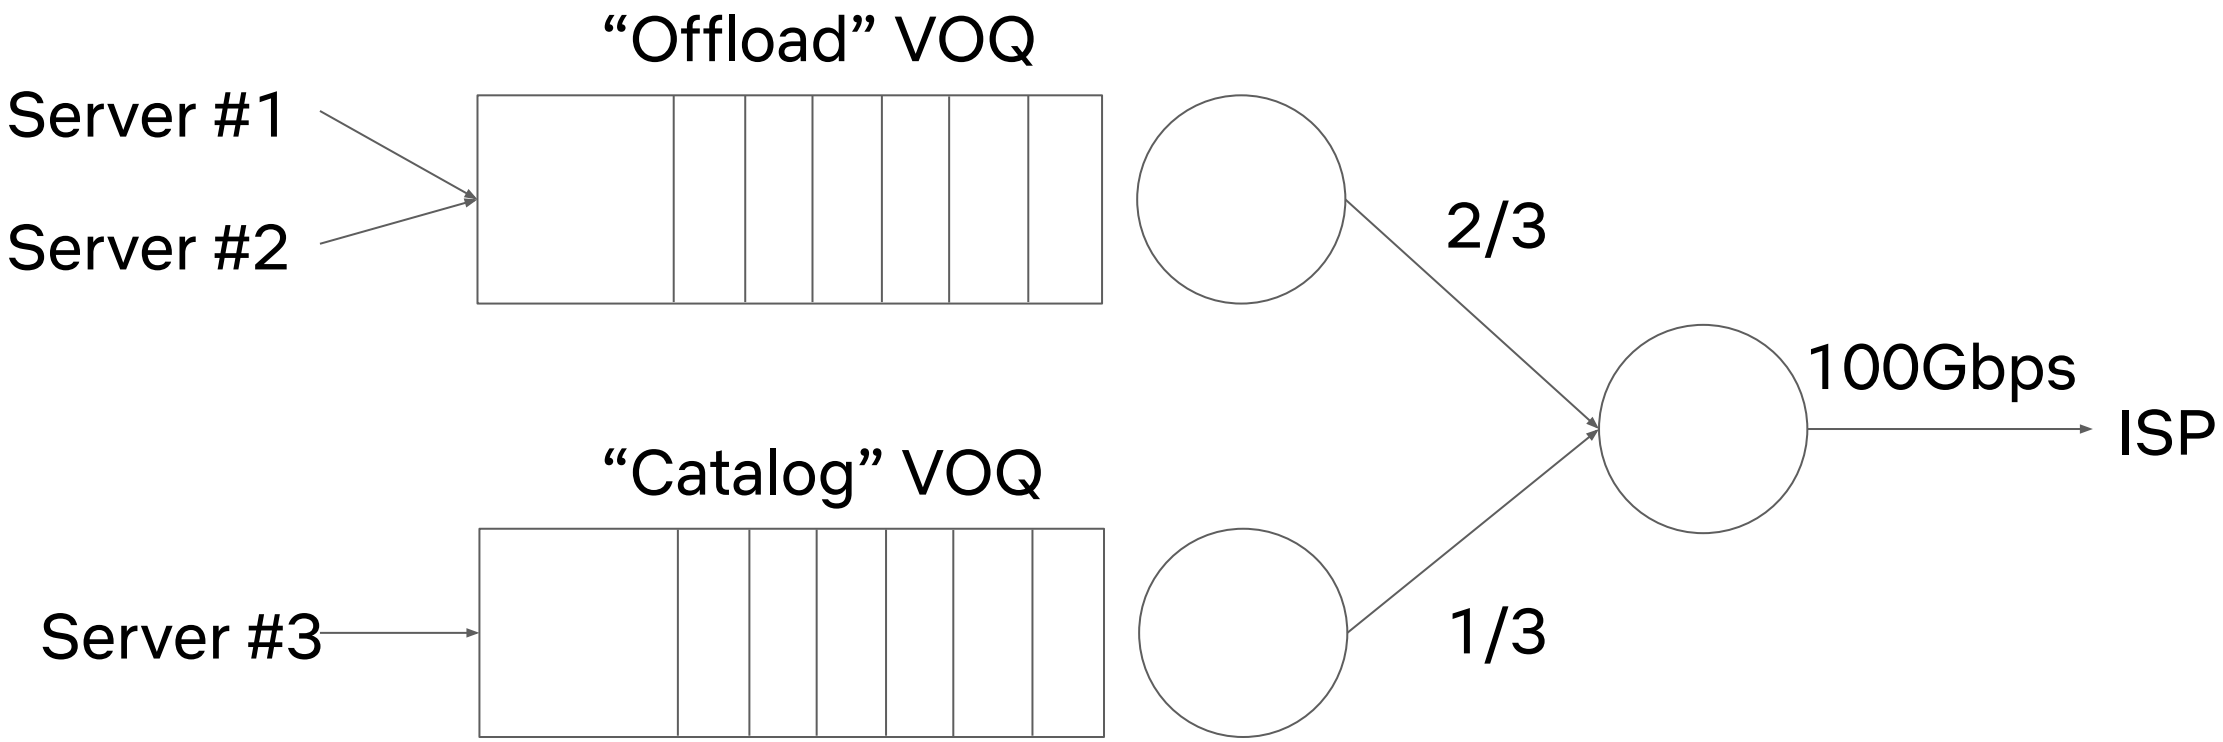

Traffic is fairly split when load is equal

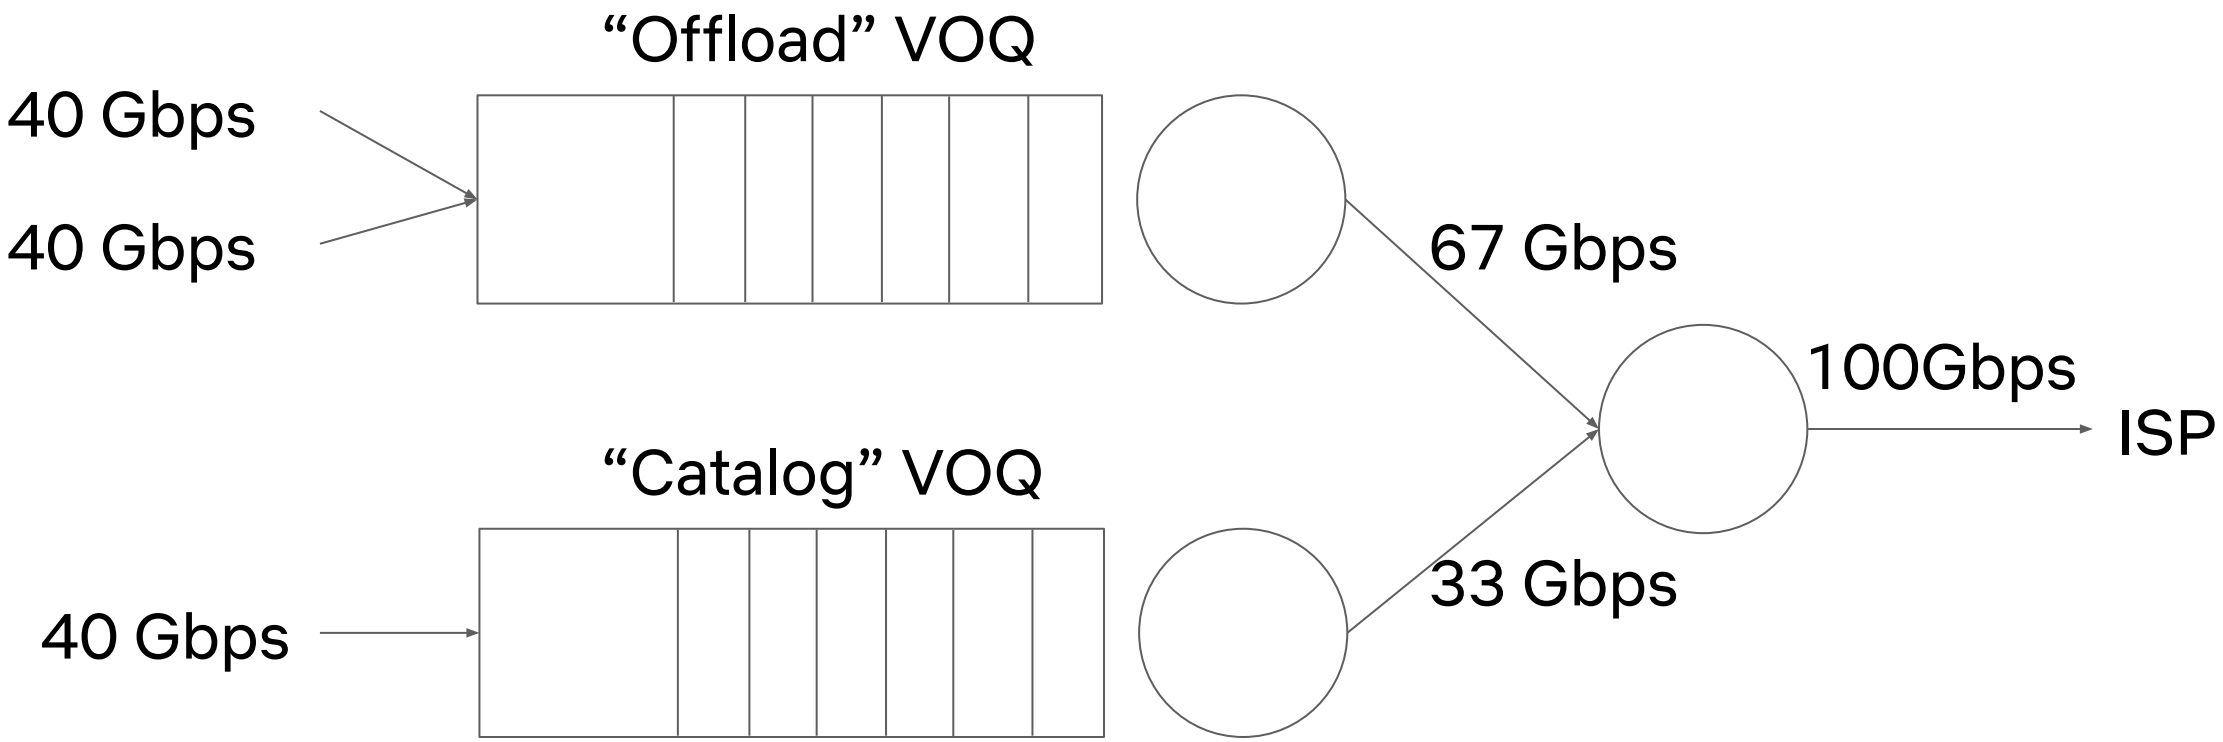

When one VOQ offers less than its fair share, it sees no congestion

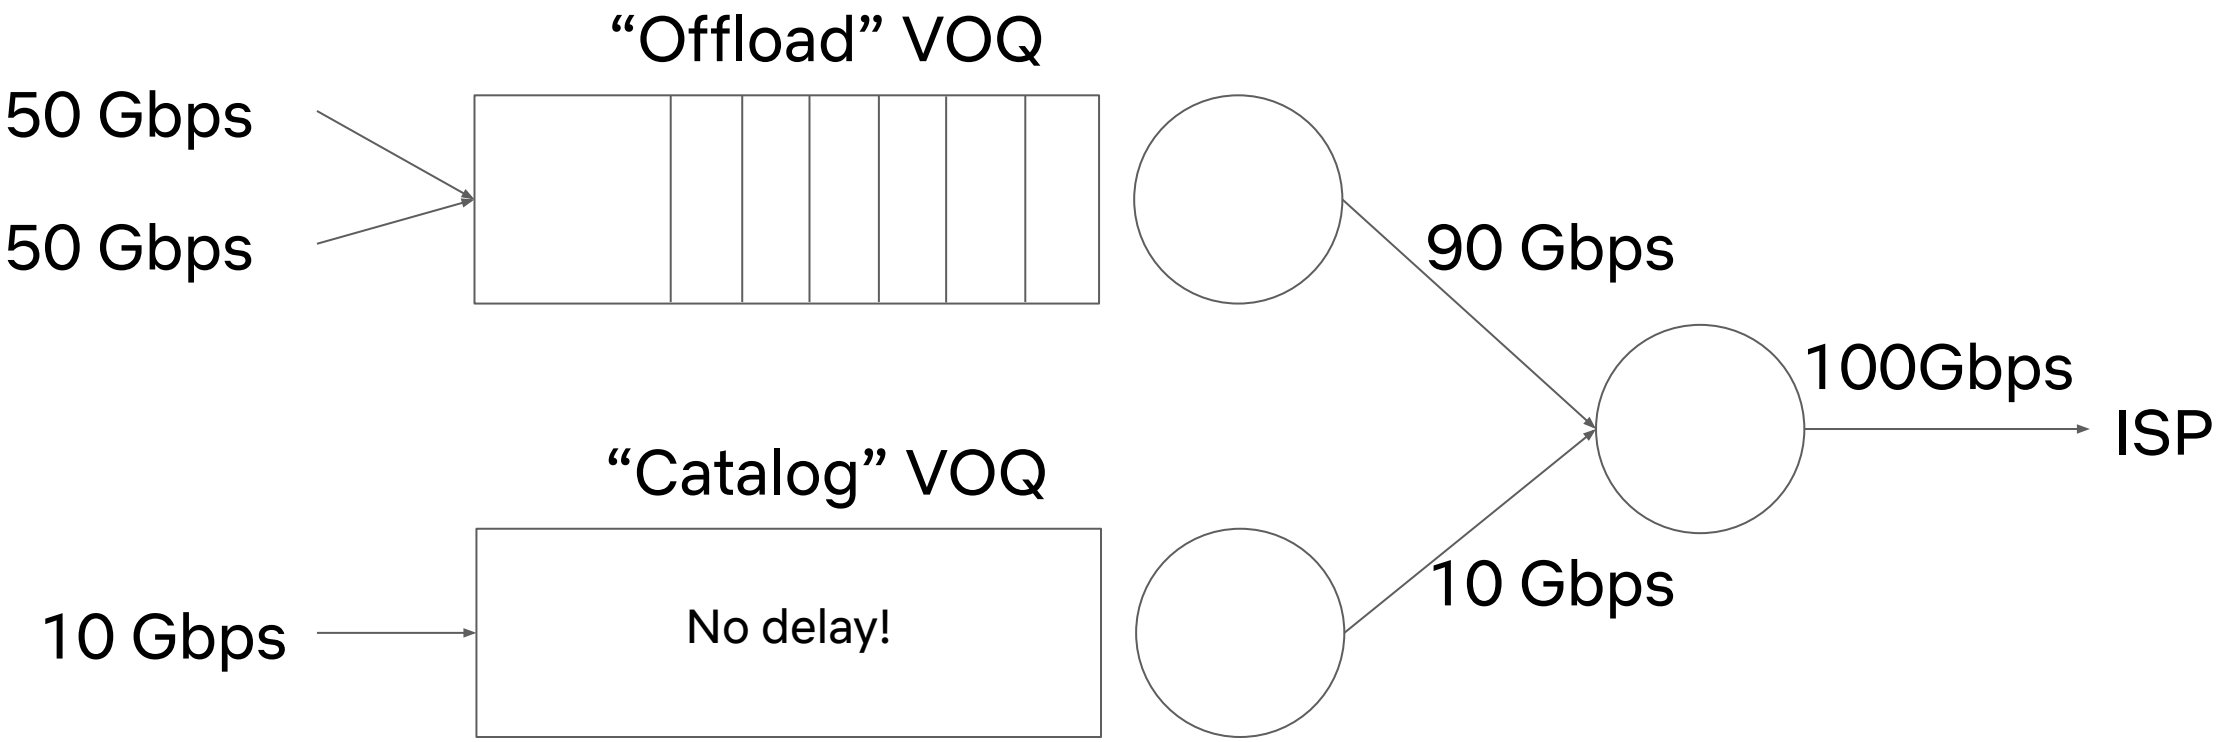

VOQs explain the RTT differences

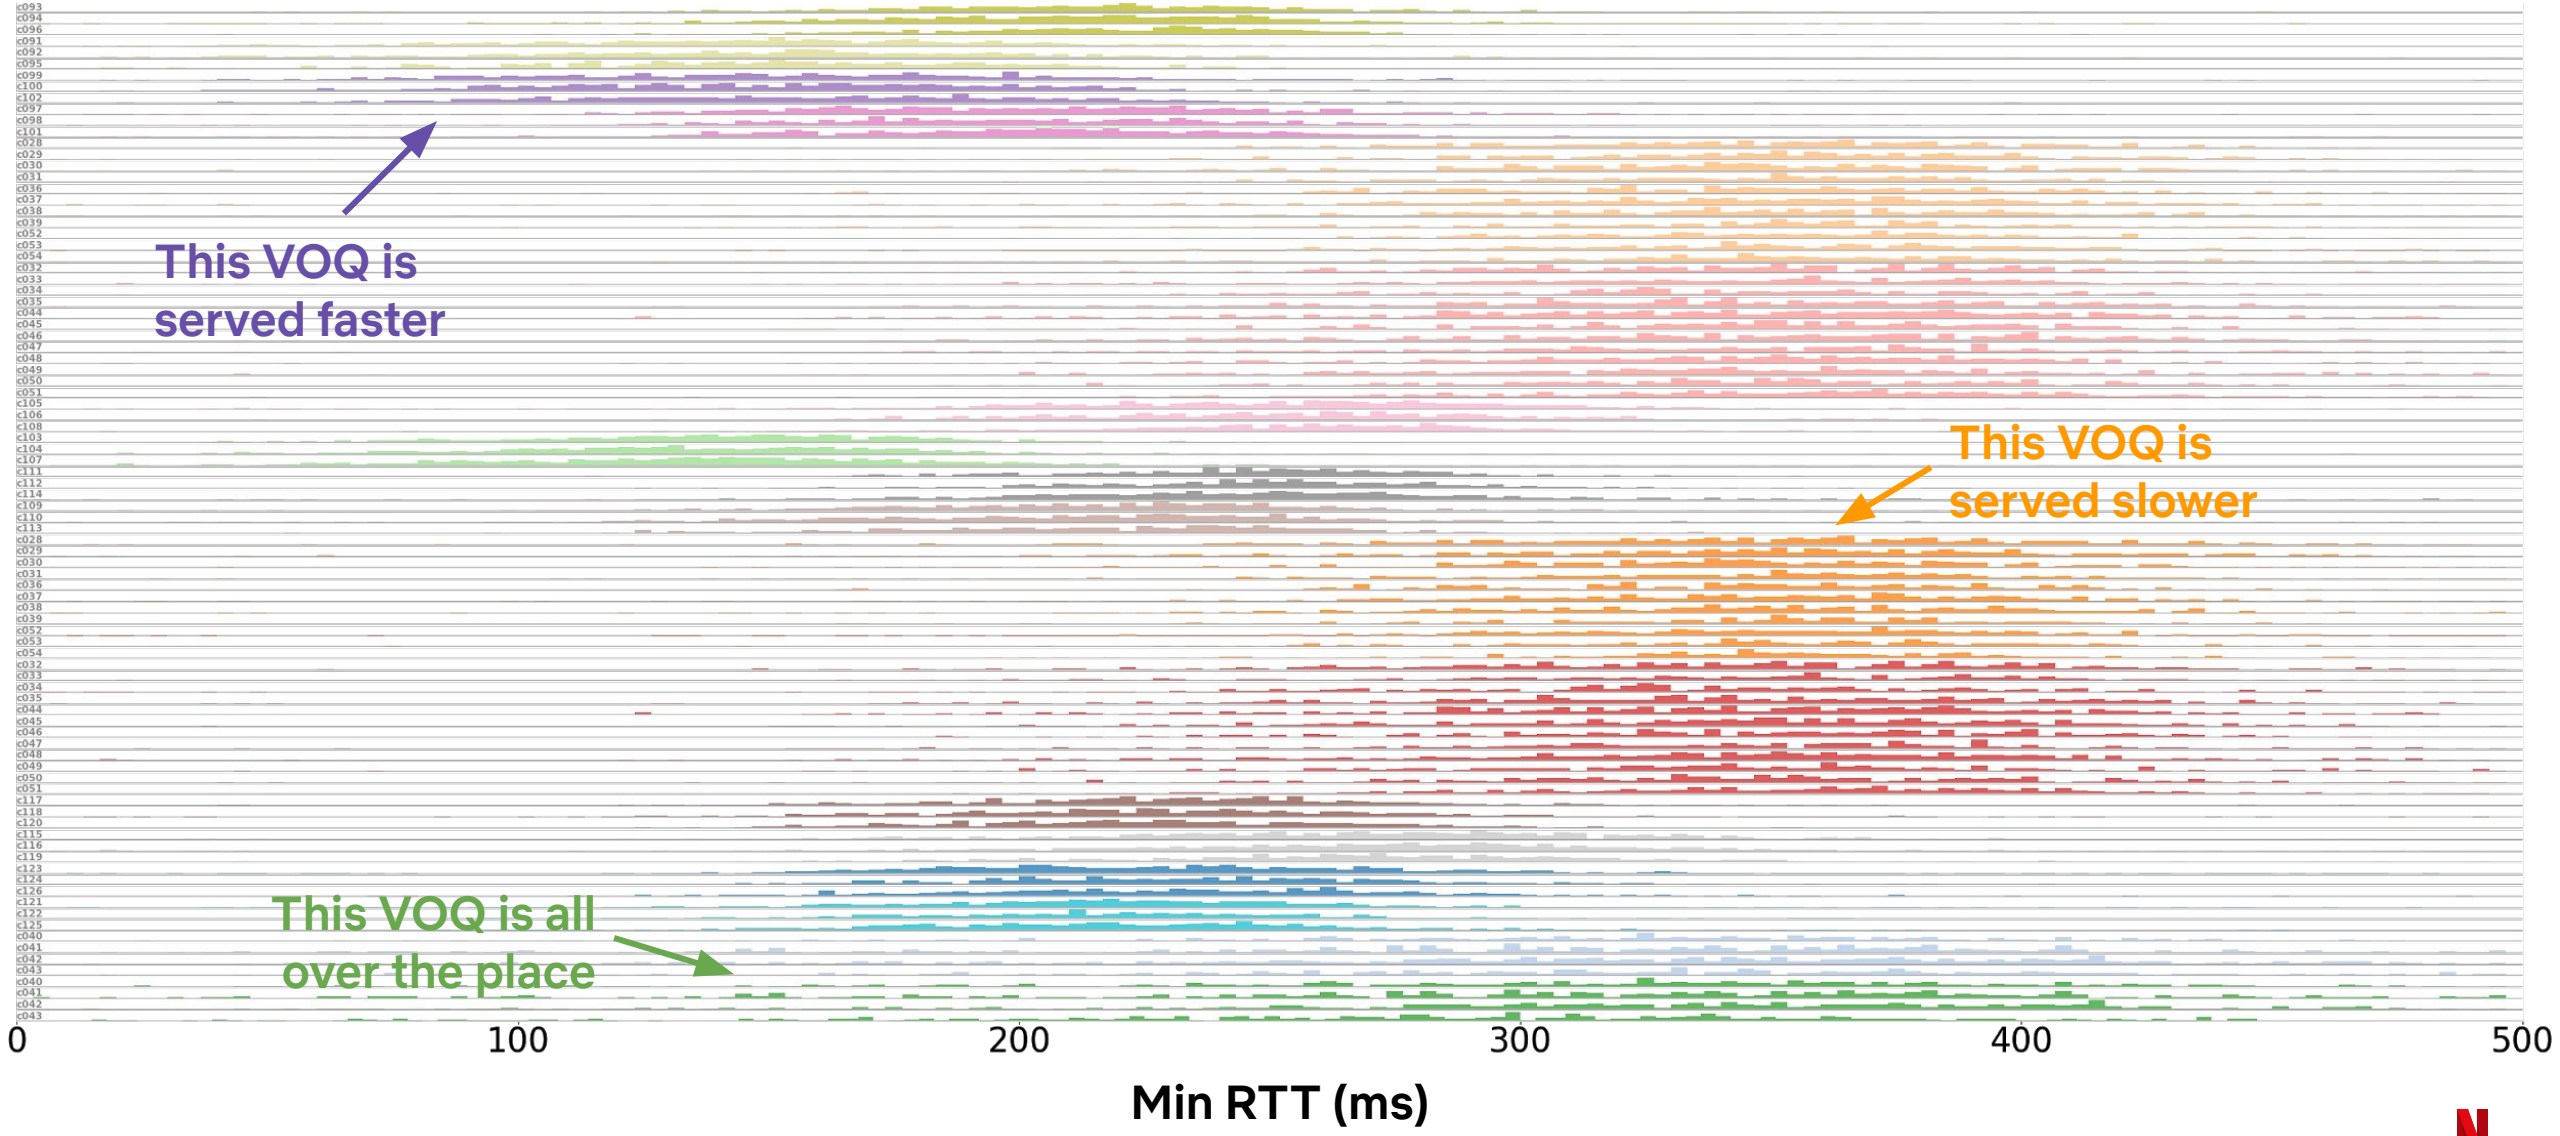

Switches prioritize long-tail content

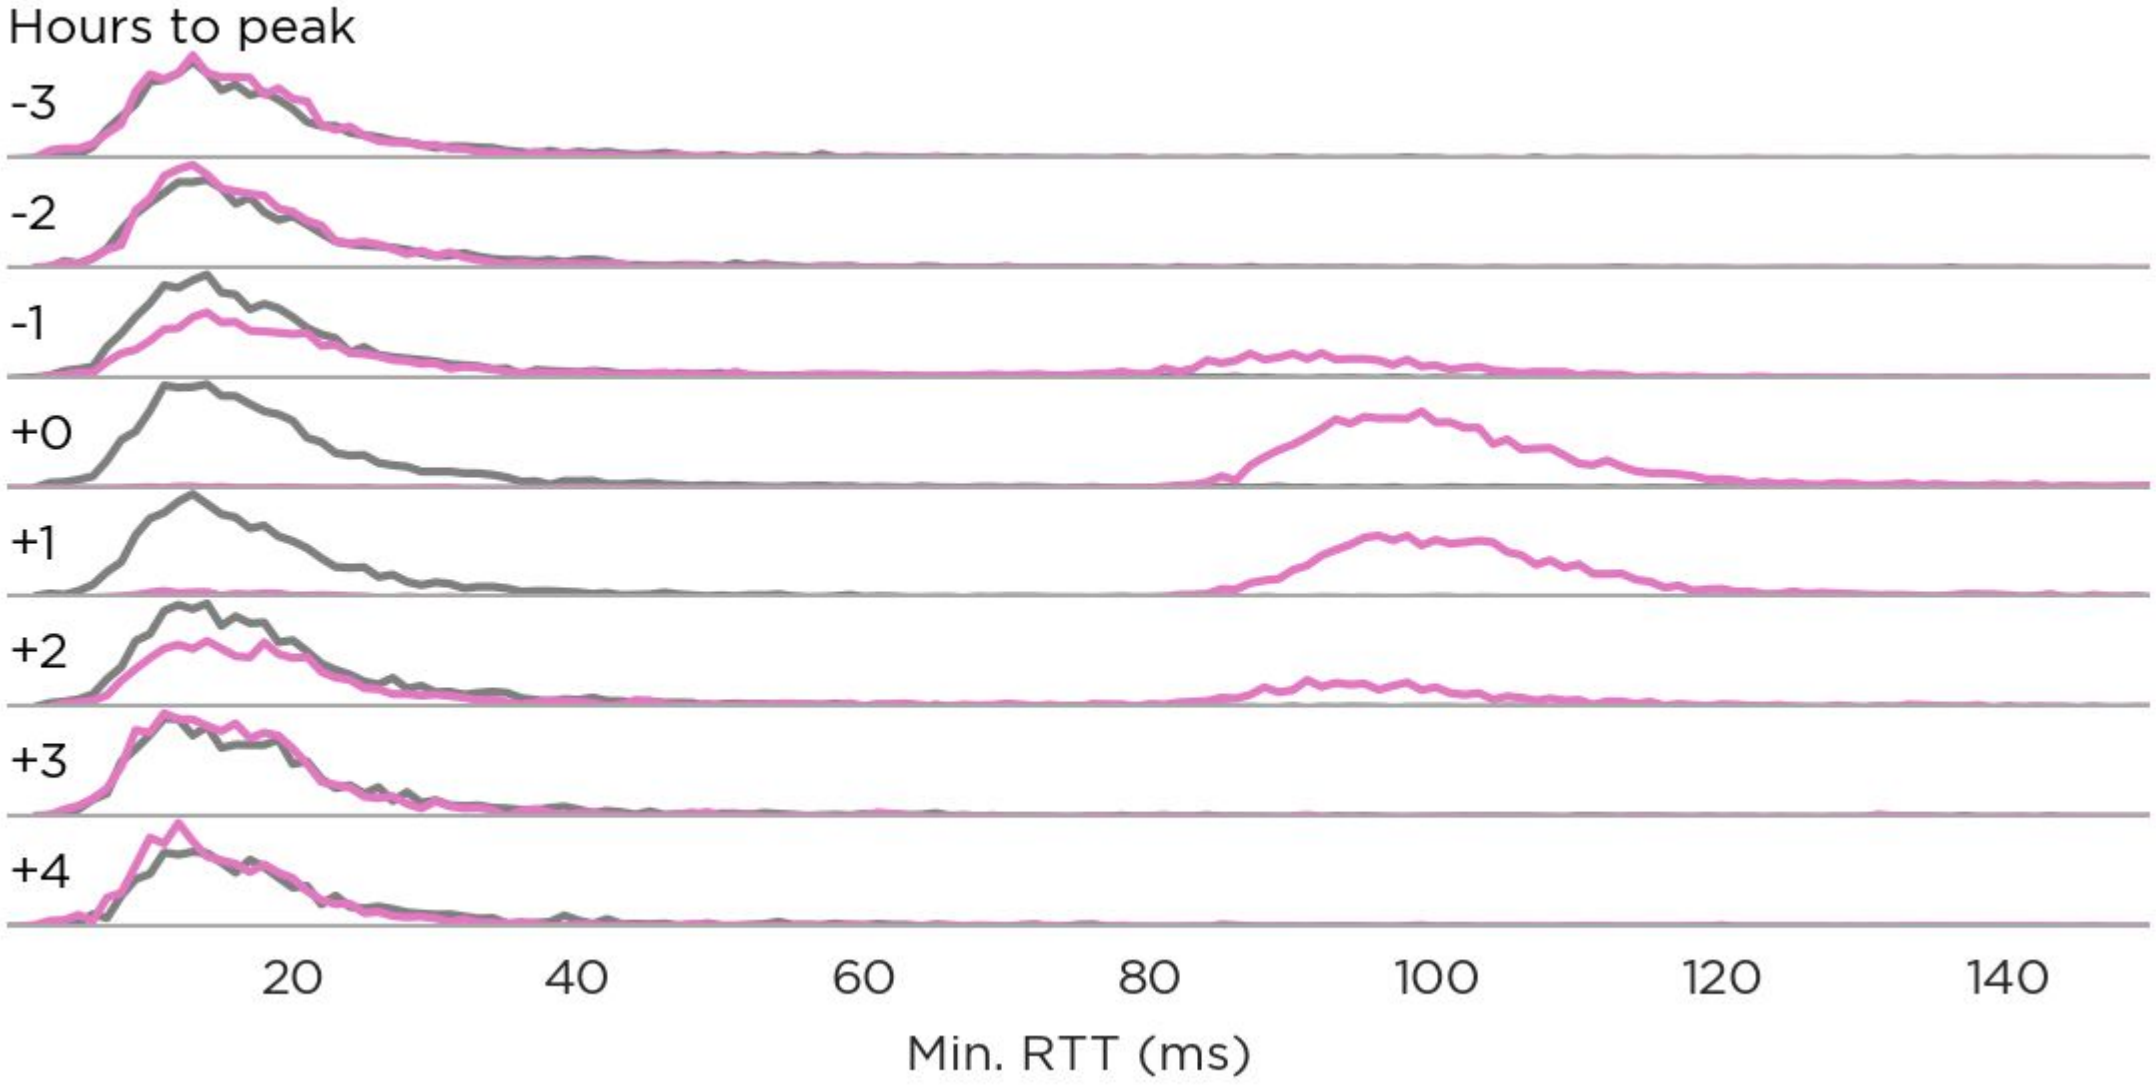

Switches prioritize long-tail content

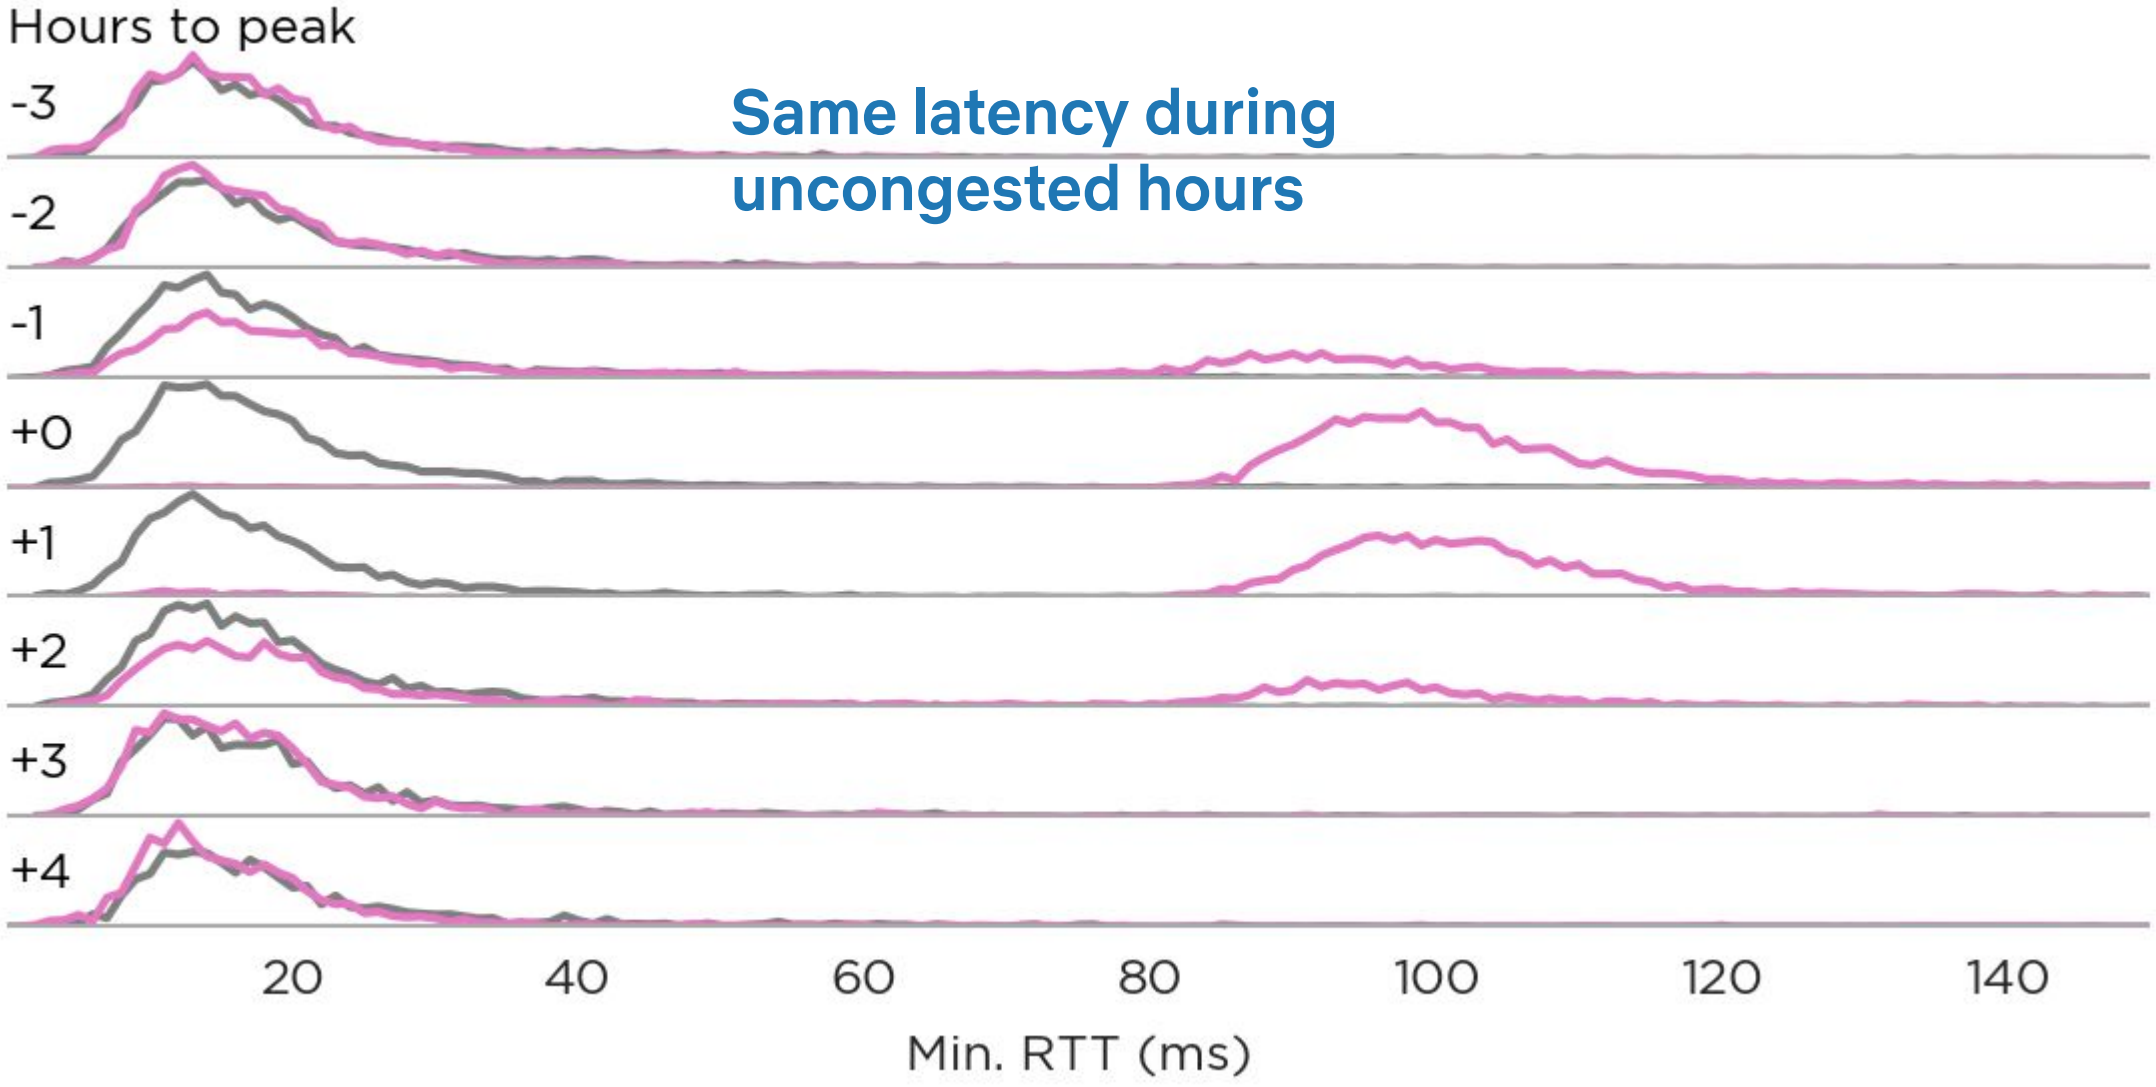

Switches prioritize long-tail content

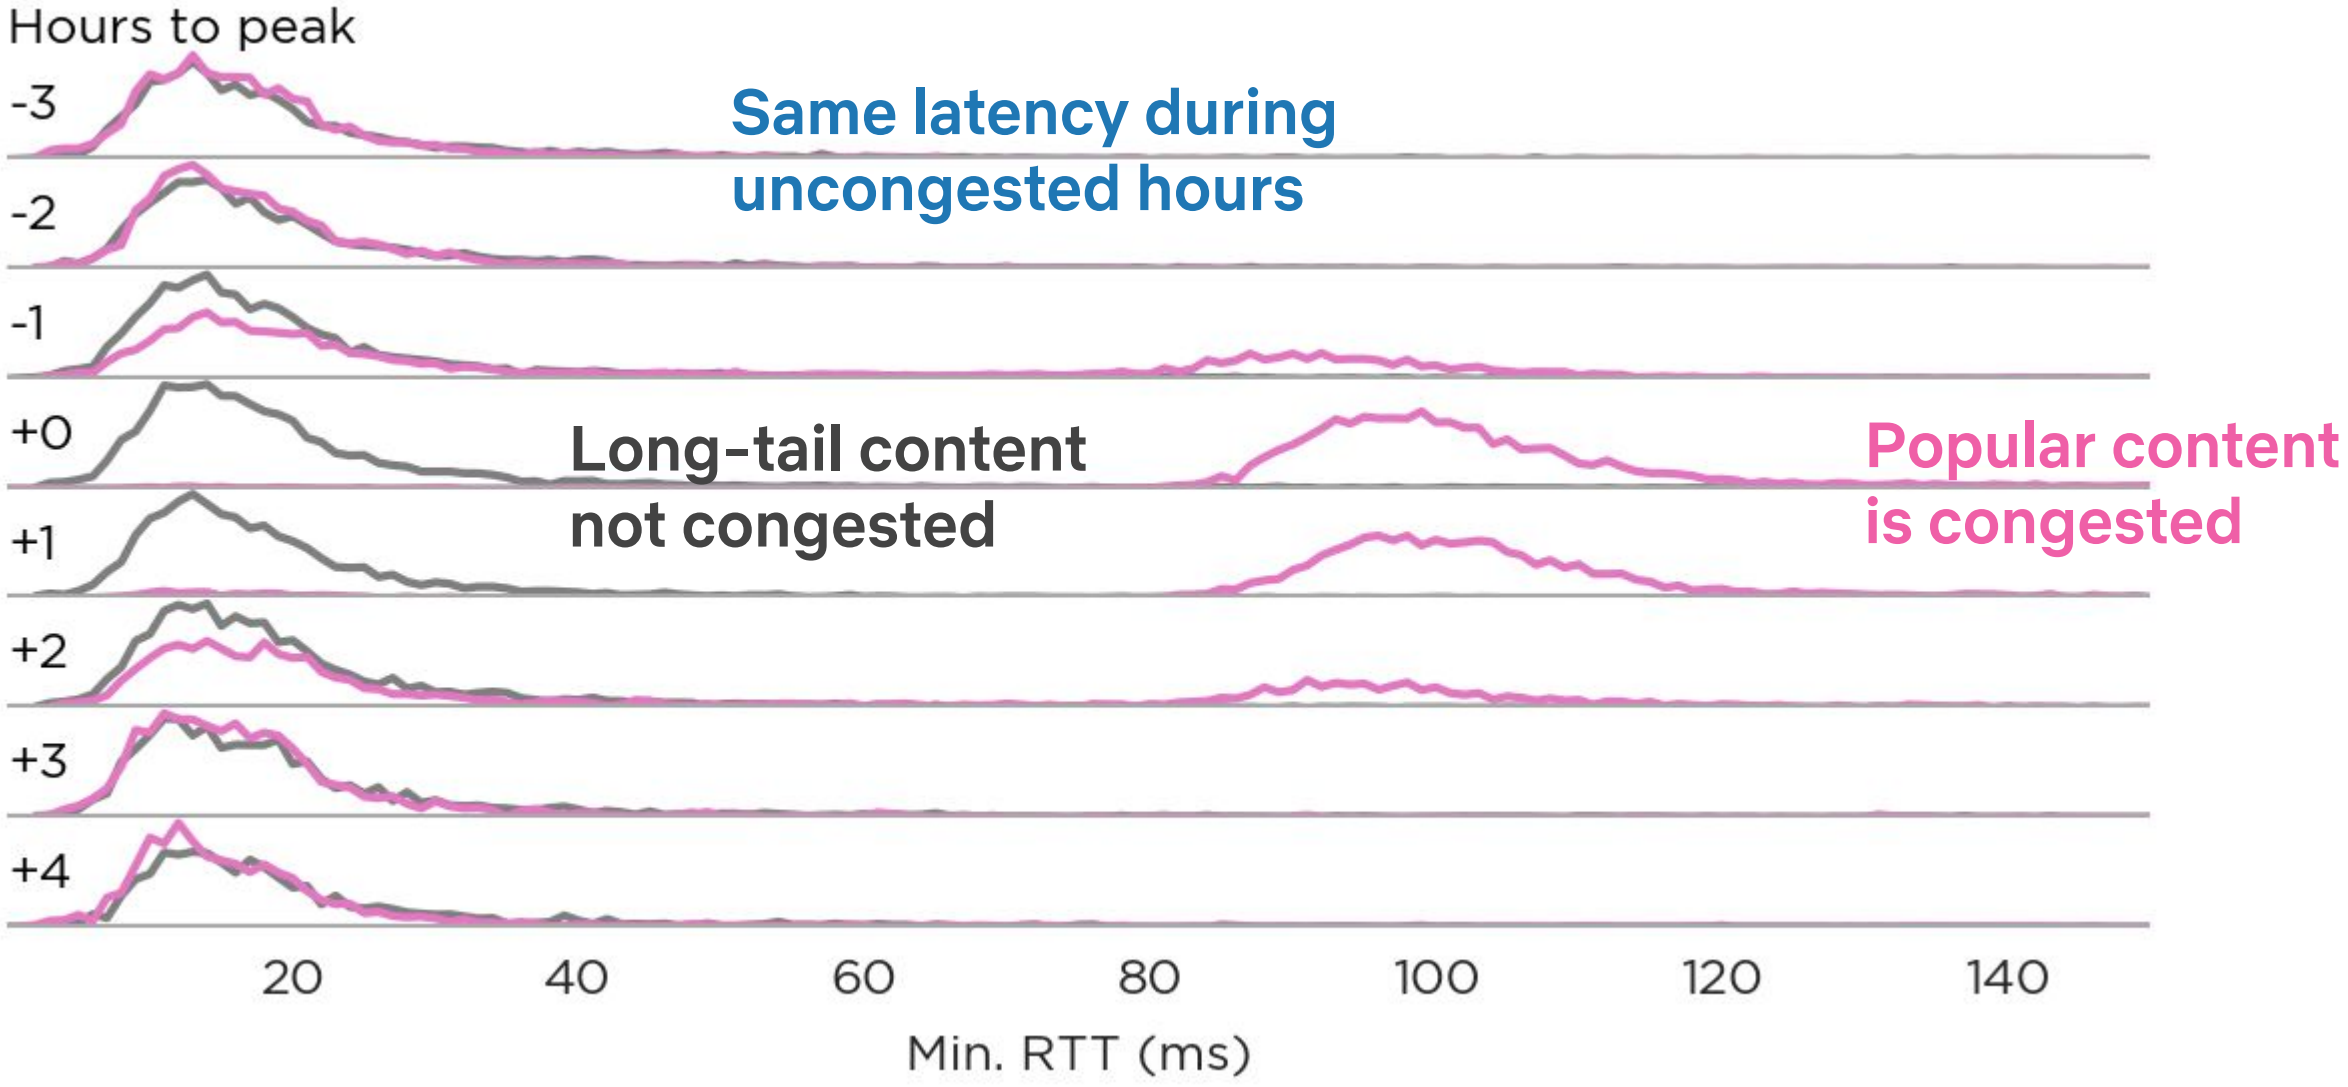

New scheduling algorithm!

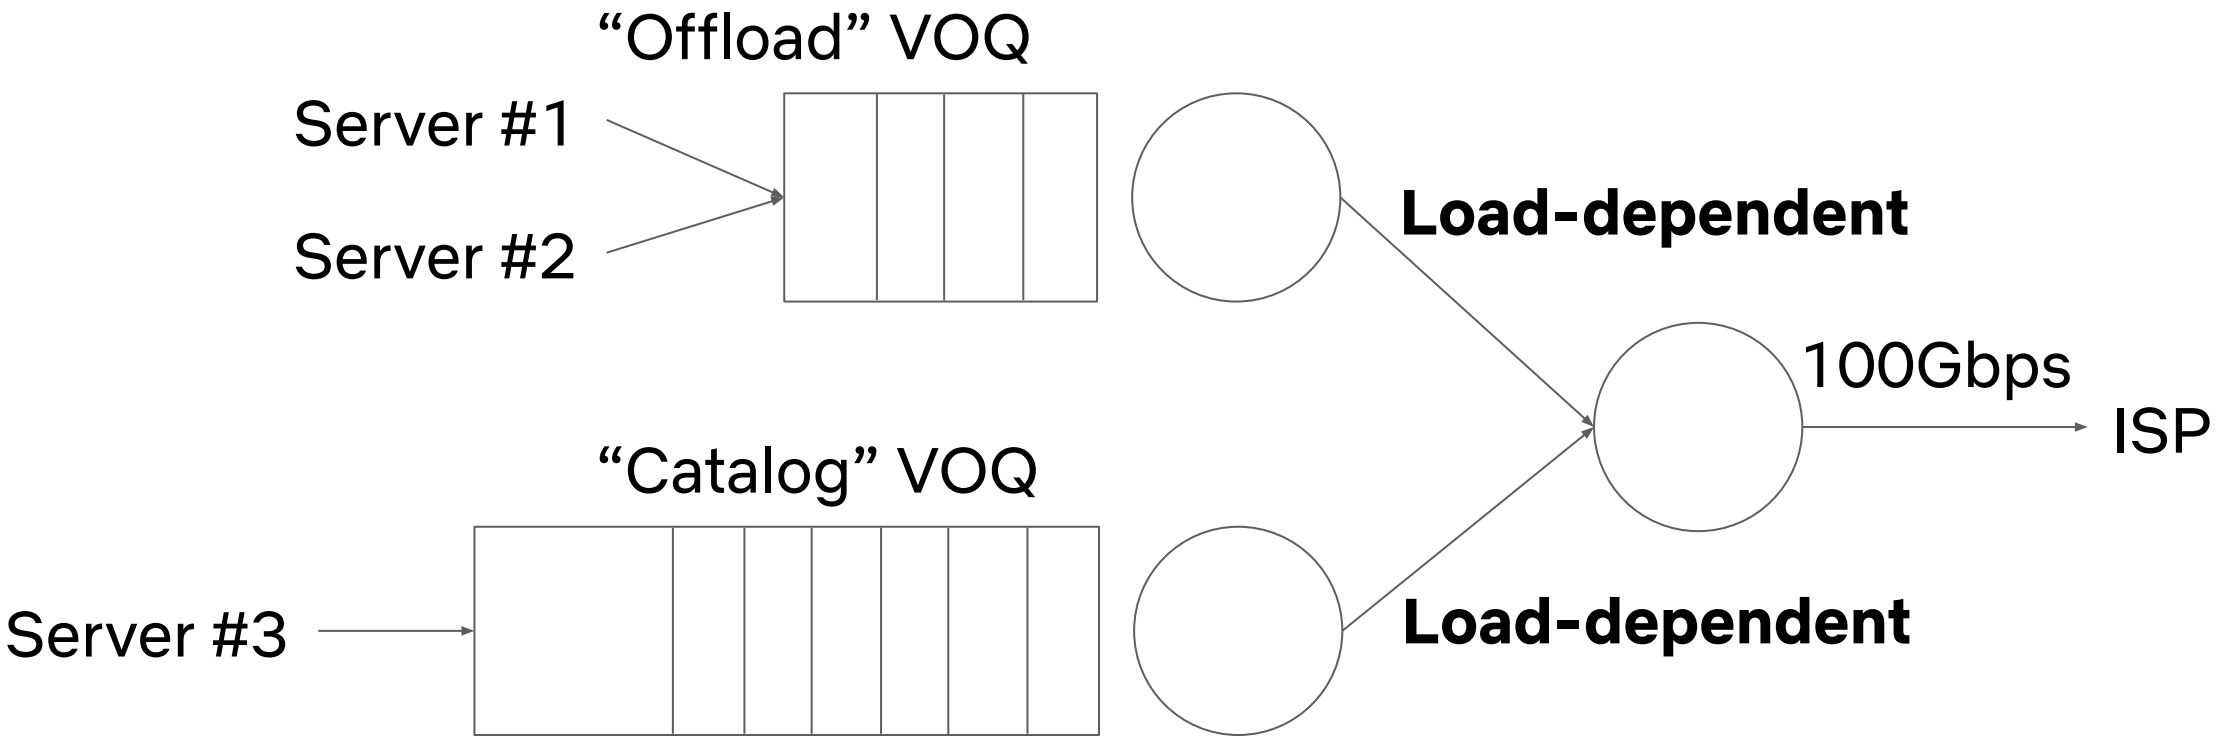

New scheduling algorithm is more consistent

Default scheduling algorithm

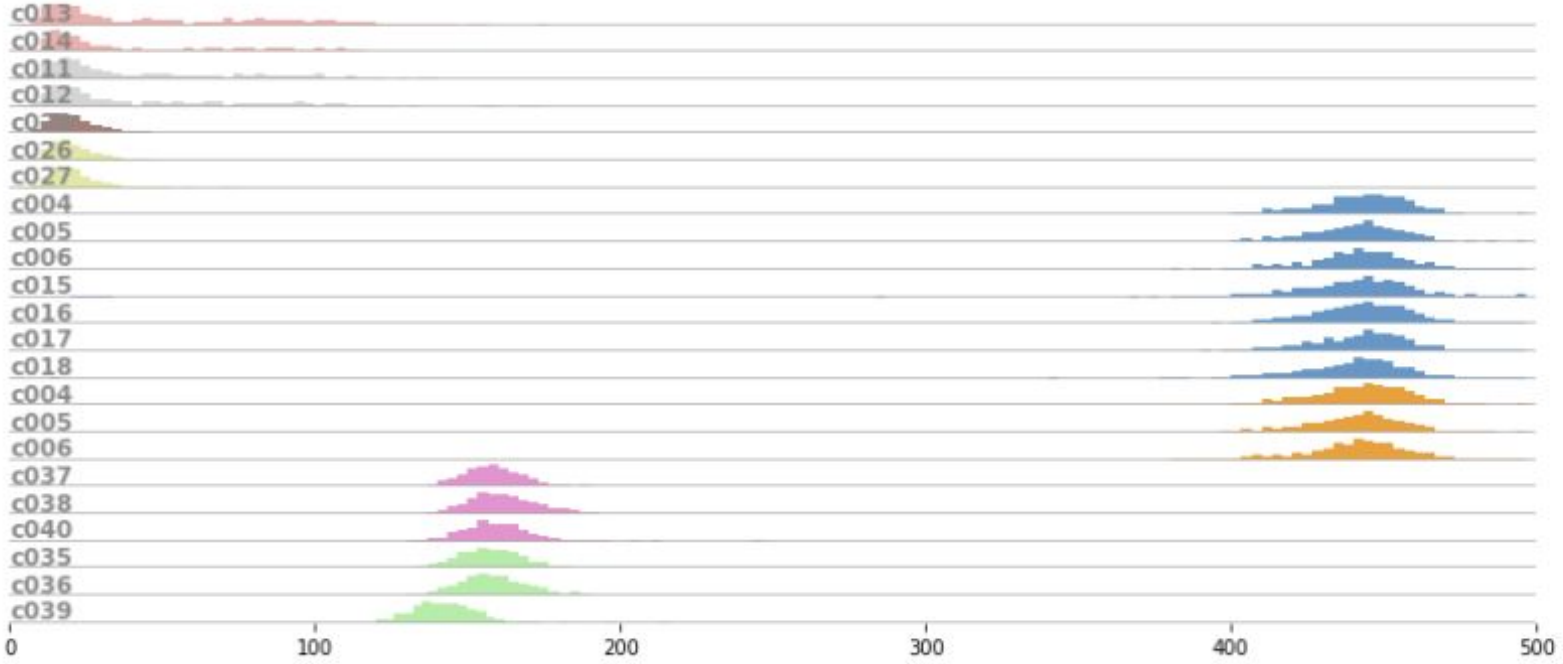

1. TCP New Reno (mostly) behaves as expected
2. Video performance varies
3. Real routers complicate this story

How big should a buffer be?

Thanks!

For more details, please see:
https://brucespang.com/papers/netflix-buffer-sizing.pdf