Towards Detecting Differential QoE

David Choffnes, Northeastern
(joint work with SBU)

Supported by a Google Grant

Traffic differentiation

Traffic differentiation

selectively changing the performance of network traffic

- Reasons for differentiation:
  - traffic engineering
  - bandwidth management
  - business reasons

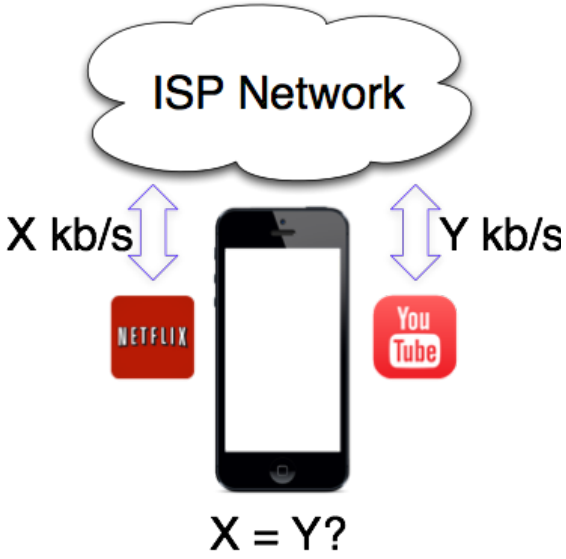

Traffic differentiation

Traffic differentiation

selectively changing the performance of network traffic

- Reasons for differentiation:
  - traffic engineering
  - bandwidth management
  - business reasons

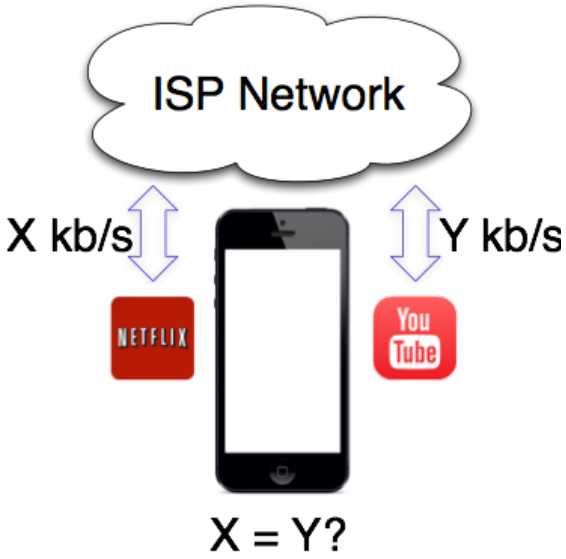

Do certain types of network traffic receive **better (or worse) QoE**?

Motivation: Differentiation in the wild

- Tested in early 2015

ISP	YouTube	Netflix	Spotify
Verizon	m	m	m
Tmobile	-	-	-
ATT	m	m	m
Sprint	m	m	m
Boost	m	m	m
BlackWireless	60%	-	-
H2O	37%	45%	65%
SimpleMobile	36%	-	-
NET10	p	p	p

- m: content modified on the fly
- p: translucent proxies change connection behavior

Motivation: Differentiation in the wild

- Tested in early 2015

ISP	YouTube	Netflix	Spotify
Verizon	m	m	m
Tmobile	-	-	-
ATT	m	m	m
Sprint	m	m	m
Boost	m	m	m
BlackWireless	60%	-	-
H2O	37%	45%	65%
SimpleMobile	36%	-	-
NET10	p	p	p

- m: content modified on the fly
- p: translucent proxies change connection behavior

Motivation: Differentiation in the wild

- Tested in early 2015

ISP	YouTube	Netflix	Spotify
Verizon	m	m	m
Tmobile	-	-	-
ATT	m	m	m
Sprint	m	m	m
Boost	m	m	m
BlackWireless	60%	-	-
H2O	37%	45%	65%
SimpleMobile	36%	-	-
NET10	p	p	p

- m: content modified on the fly
- p: translucent proxies change connection behavior

Motivation: Differentiation in the wild

- Tested in early 2015
- Again in late 2015: No observed differences

ISP	YouTube	Netflix	Spotify
Verizon	m	m	m
Tmobile	-	-	-
ATT	m	m	m
Sprint	m	m	m
Boost	m	m	m
BlackWireless	60%	-	-
H2O	37%	45%	65%
SimpleMobile	36%	-	-
NET10	p	p	p

- m: content modified on the fly
- p: translucent proxies change connection behavior

Key Questions

What do you test for differentiation?

- How do you generate traffic?
- What triggers differentiation?

How can you *tell* if there is differentiation?

- How do you do a controlled experiment?
- How do you rule out other reasons for differential service?

What is the *impact* on QoE?

- How do you map observed degradation to QoE?
- How do you scale this to arbitrary applications?

How do you test for differentiation?

- What triggers differentiation?
  - *We don't know*

- They might trigger on
  - IP addresses
  - ports
  - payload signatures
  - total number of connections
  - total bandwidth
  - time of day

  (This is consistent with online manuals for DPI boxes)

How do you test for differentiation?

- What triggers differentiation?
  - ***We tested using carrier-grade DPI boxes***

- They might trigger on
  - IP addresses
  - **ports**
  - **payload signatures**
  - total number of connections
  - total bandwidth
  - time of day

What triggers differentiation?

- HTTP
  - **Host:** and **GET** fields, typically **regex**
  - Examples: youtube, facebook, netflix
- HTTPS
  - Server cert is sent in plaintext
  - Searches on **SNI**, **CN** fields
- Other protocols
  - Can identify Skype using some knowledge of handshake format
  - **Open question:** how to reverse engineer classifiers?
- What is **not** important
  - IP addresses don't seem to matter!

Triggering differentiation

Triggering differentiation

NETFLIX

Triggering differentiation

end user

Triggering differentiation

end user

Triggering differentiation

Triggering differentiation

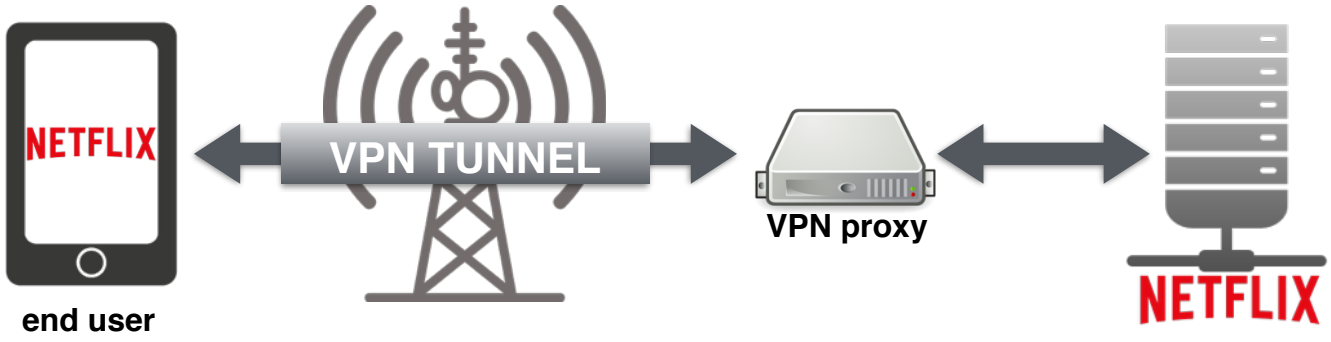

Triggering differentiation

Record:

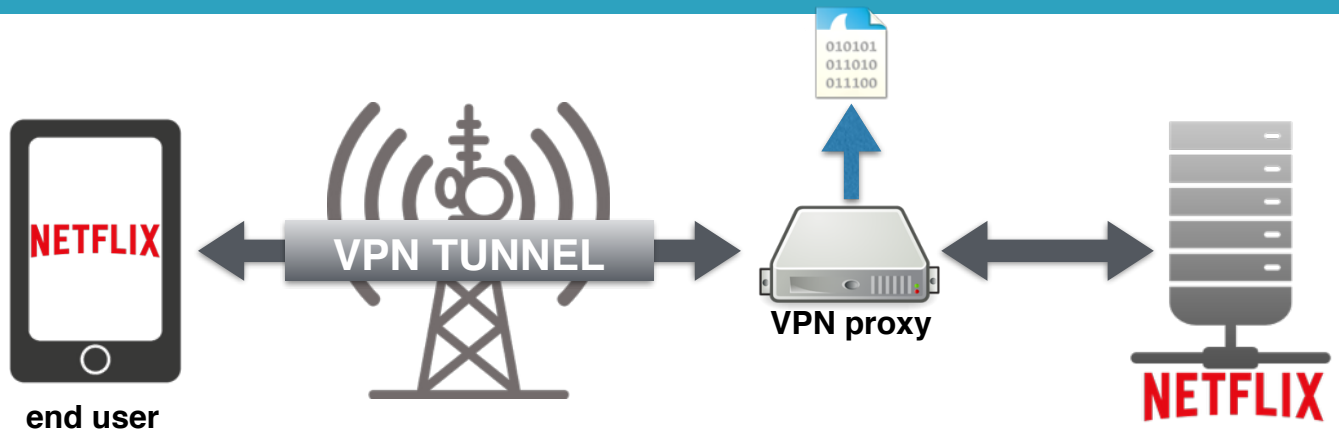

Triggering differentiation

Record:

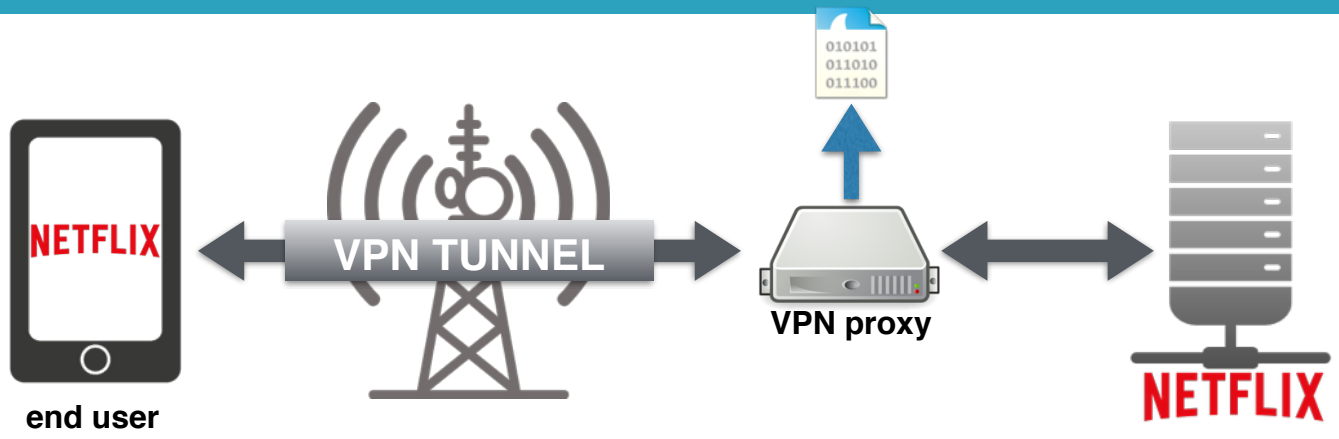

Replay:

Triggering differentiation

Record:

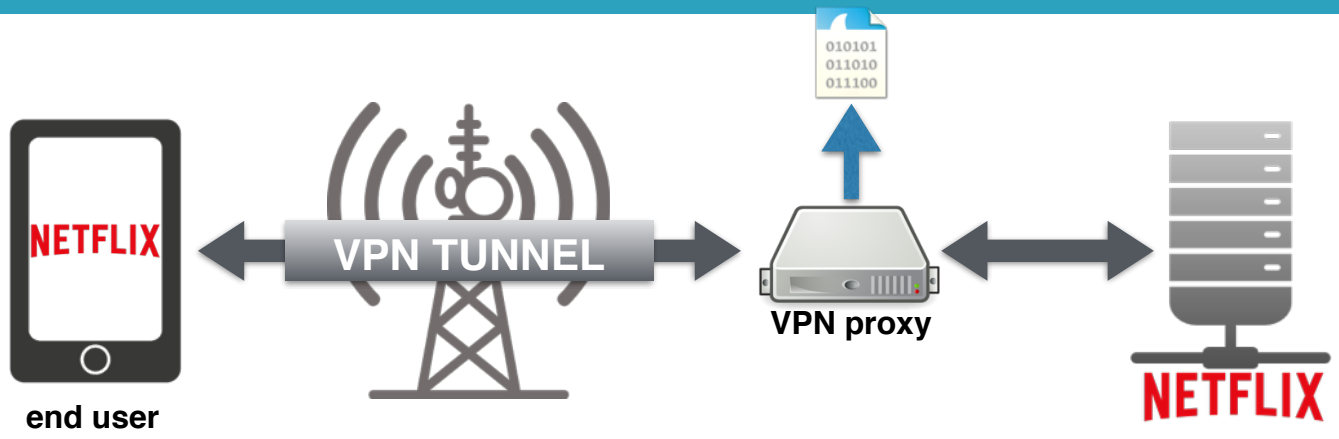

Replay:

Triggering differentiation

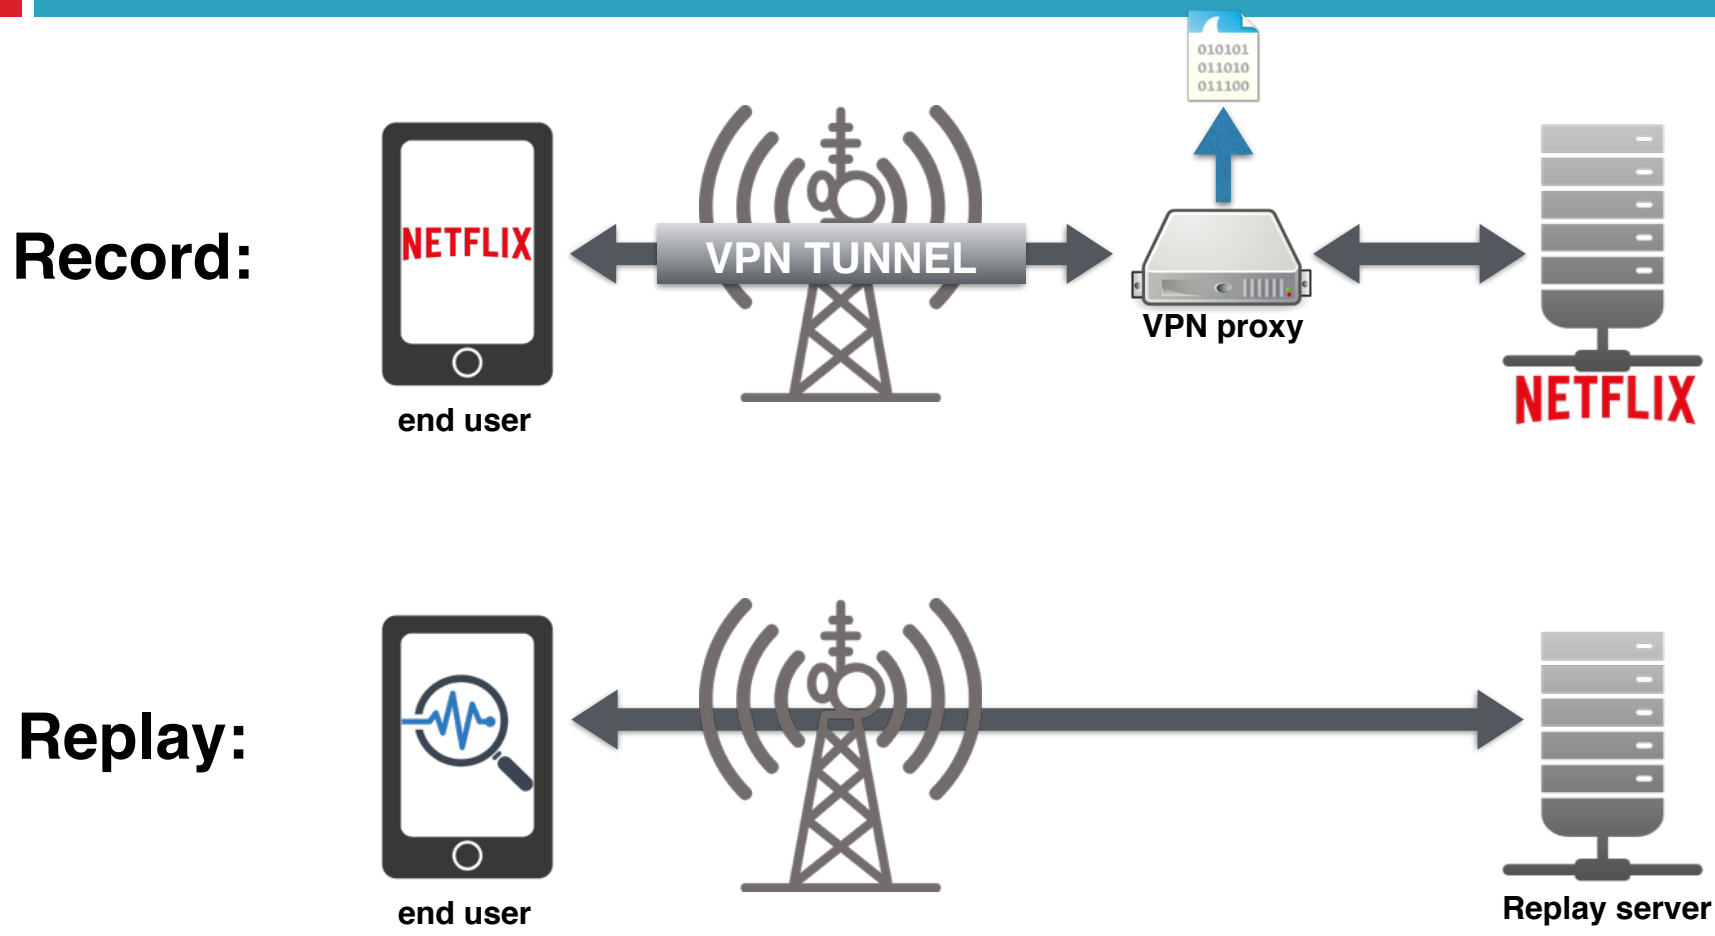

Triggering differentiation

Record:

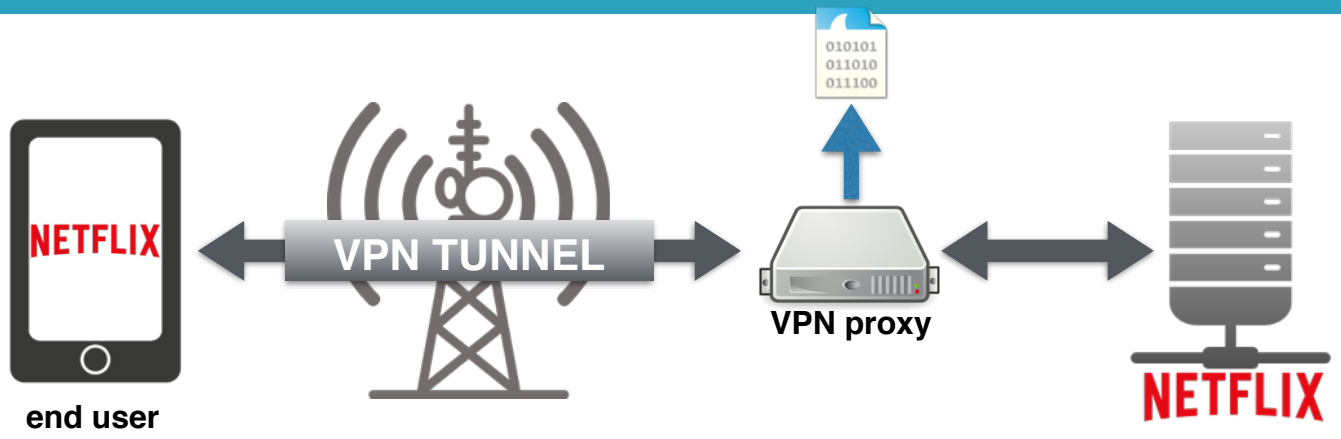

Replay:

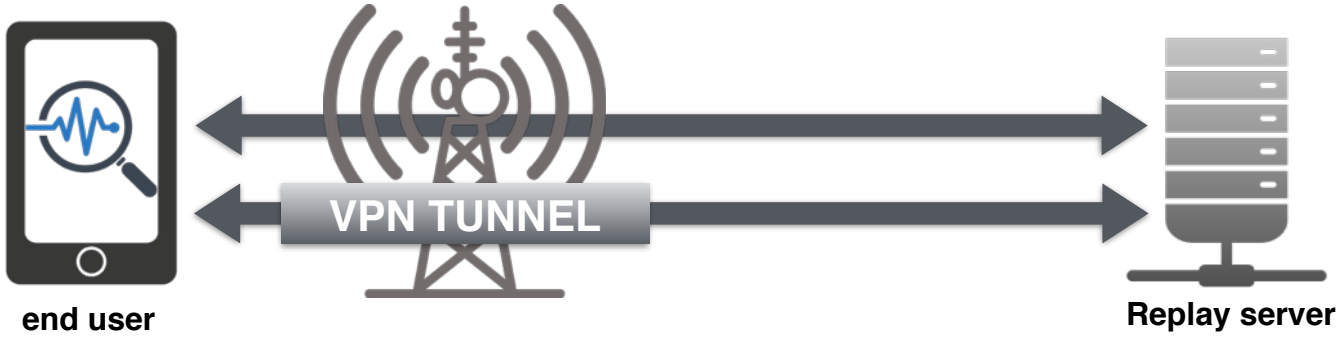

Triggering differentiation

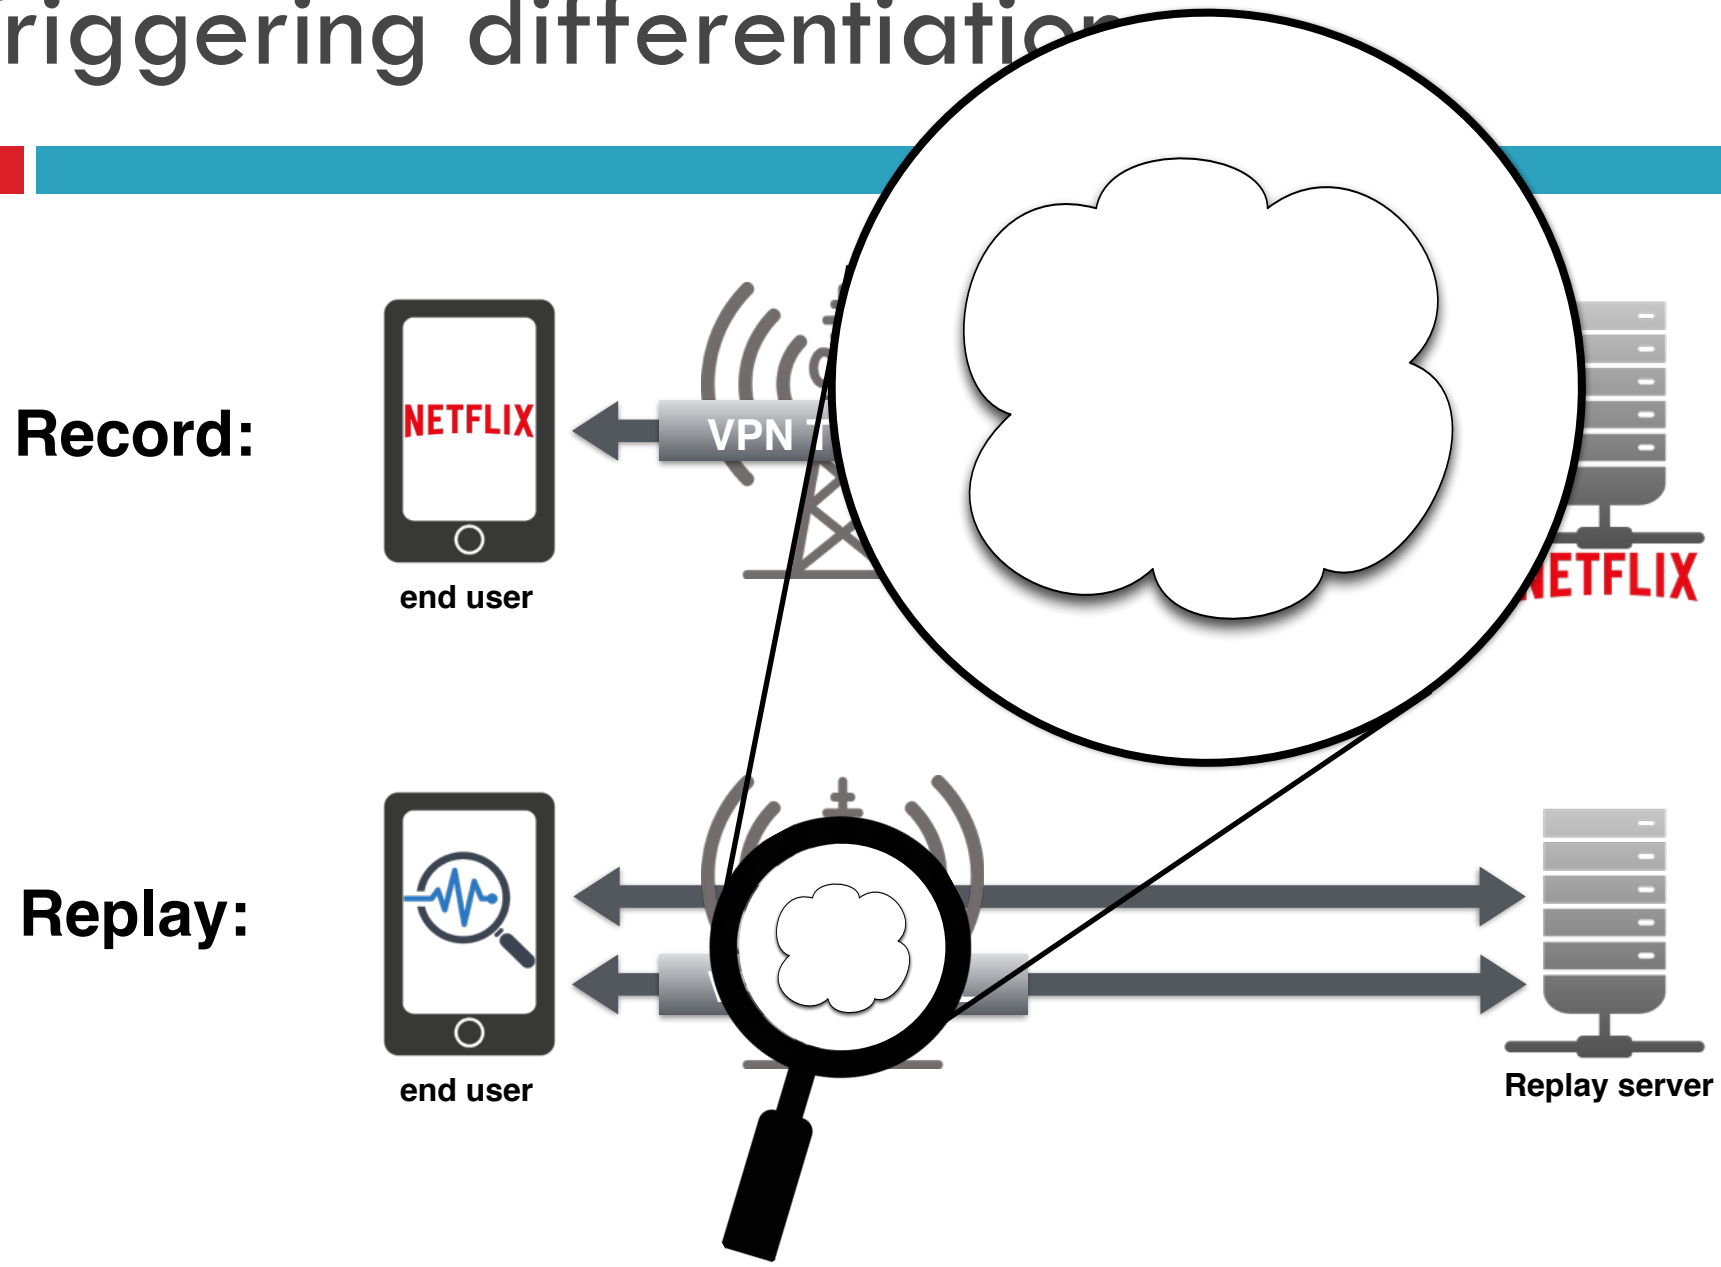

Triggering differentiation

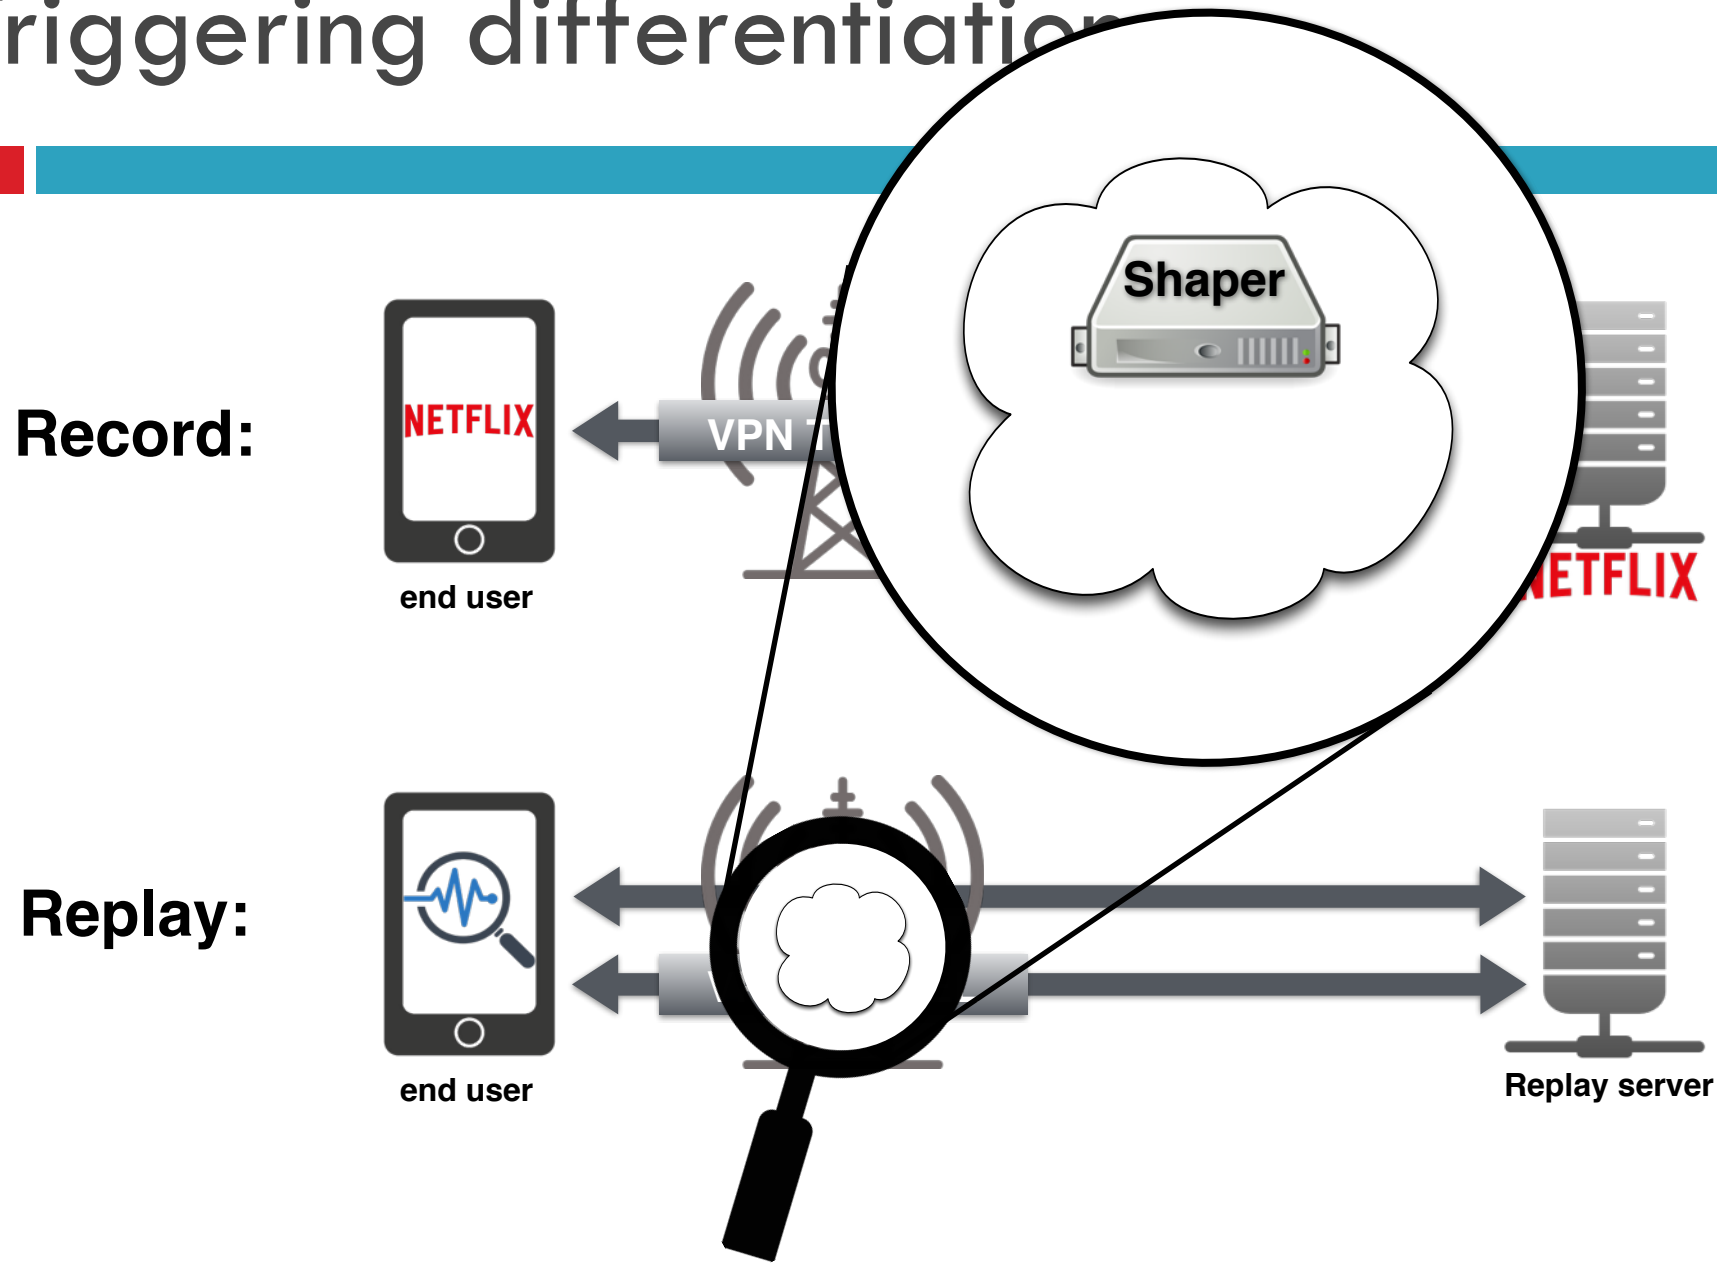

Triggering differentiation

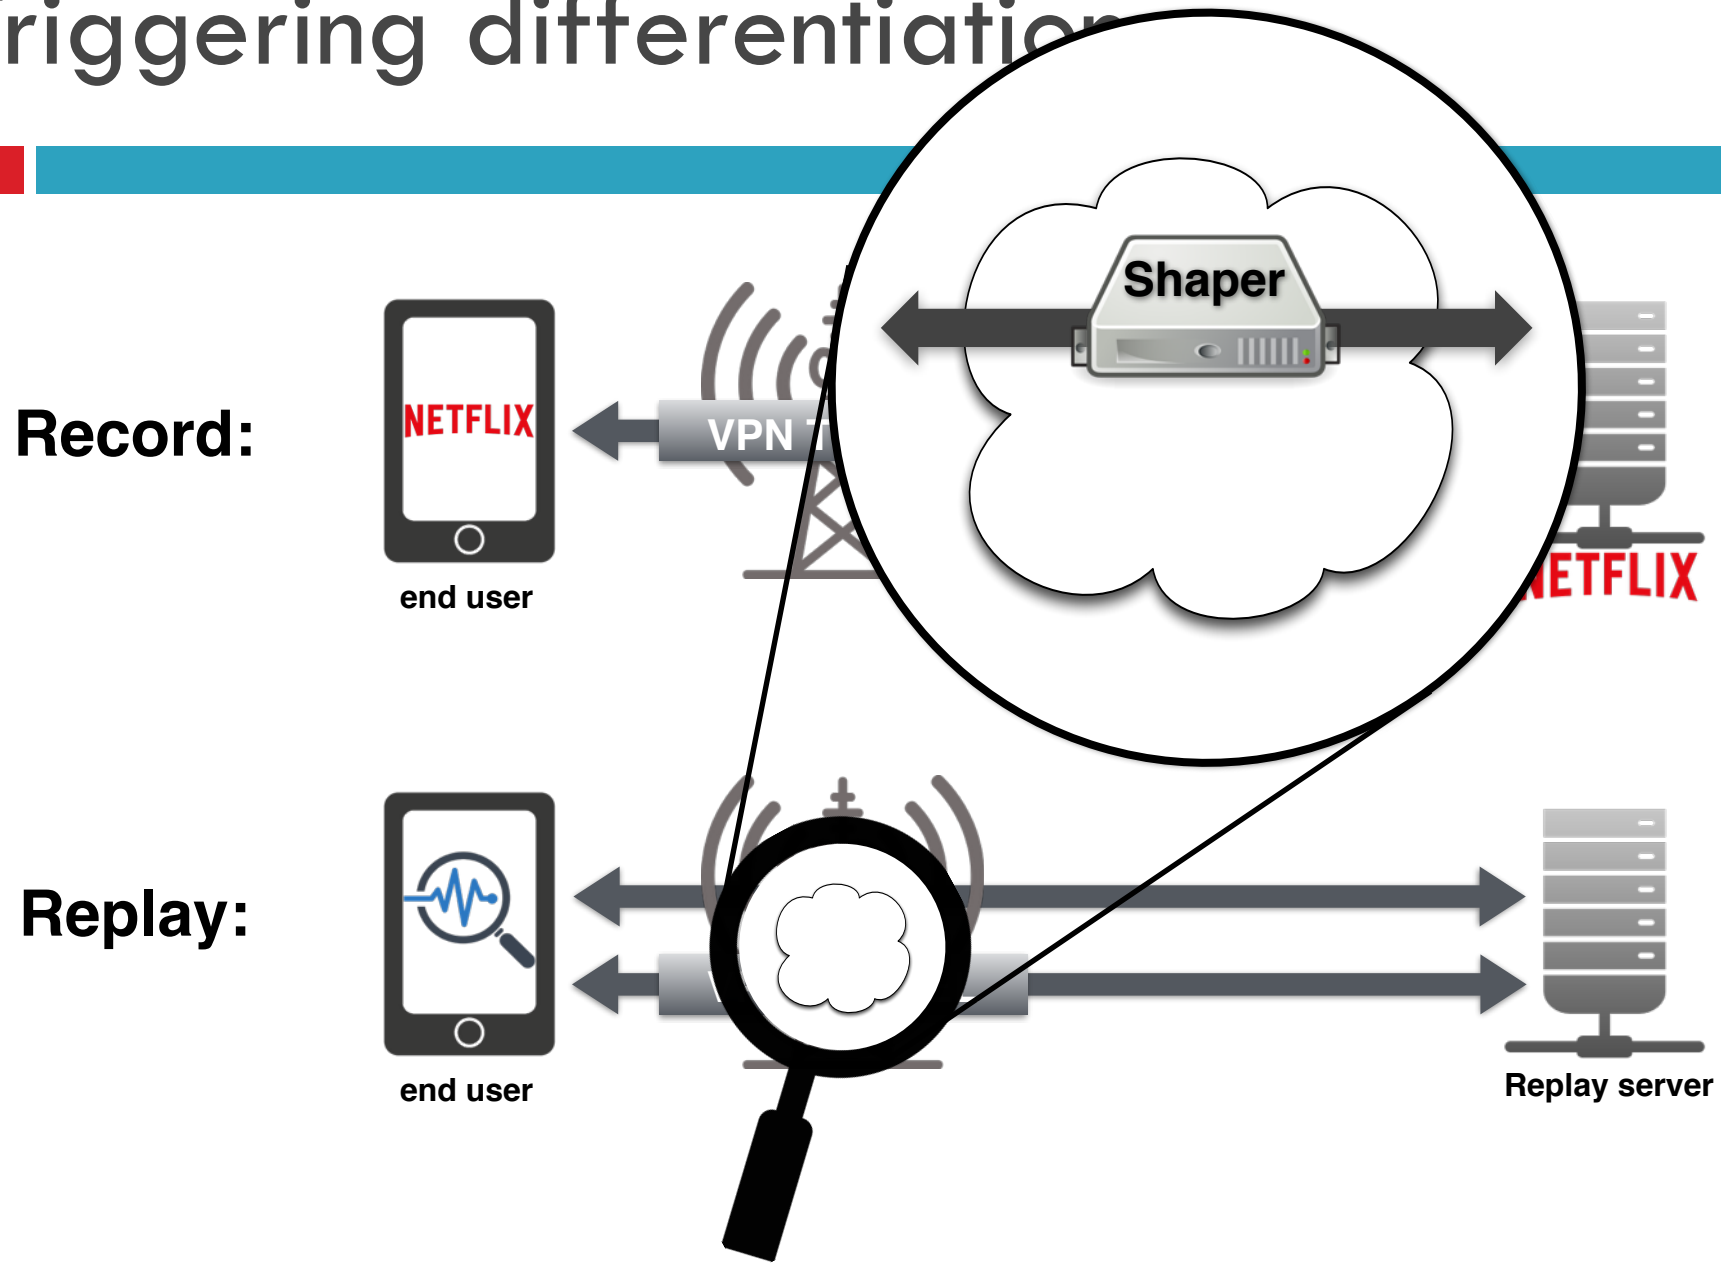

Triggering differentiation

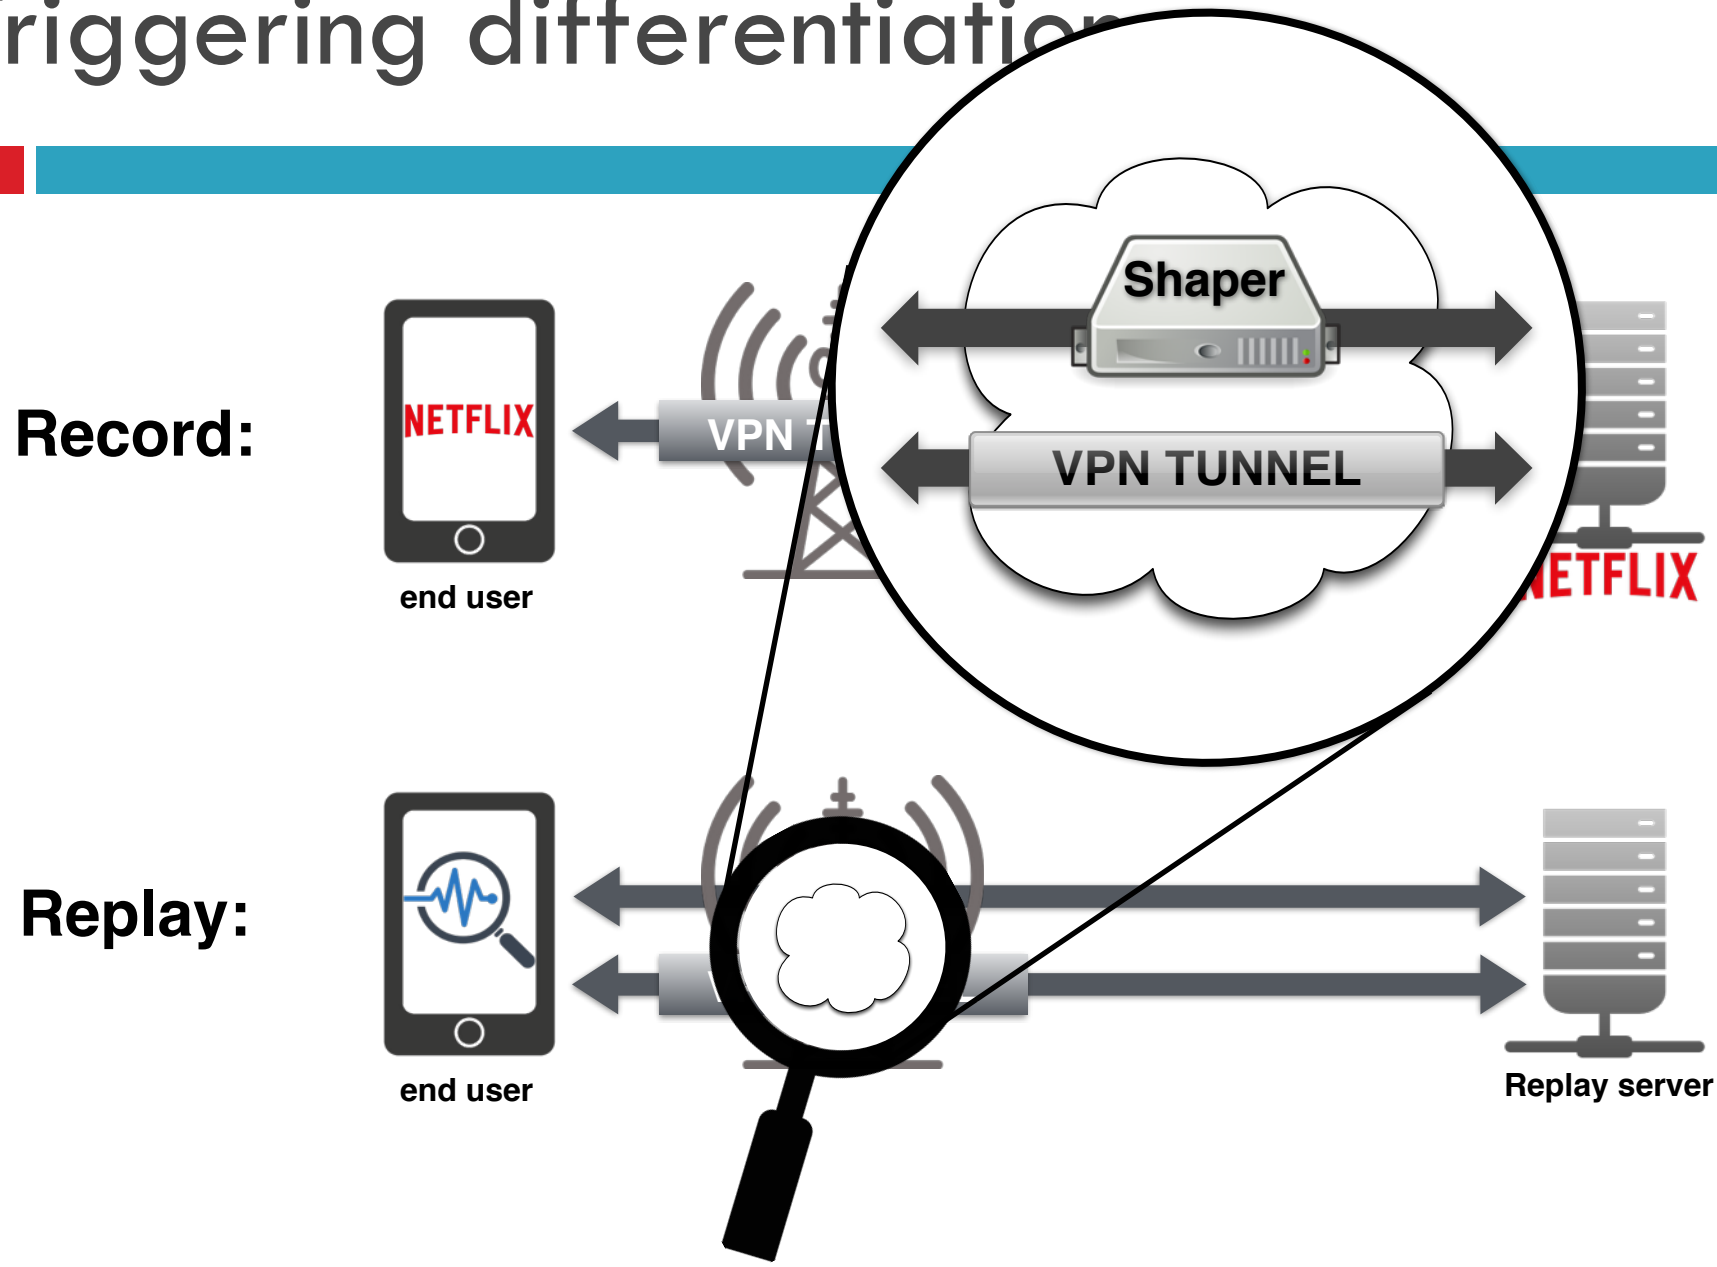

How to tell if there is differentiation?

- Developed and validated new detection technique
  - Back-to-back tests of same trace
    - Includes VPN and random payload (but not ports)
  - Send only at recorded rate (this is not a speed test)
  - Statistical tests to reliably find differentiation

How to tell if there is differentiation?

- Developed and validated new detection technique
  - Back-to-back tests of same trace
    - Includes VPN and random payload (but not ports)
  - Send only at recorded rate (this is not a speed test)
  - Statistical tests to reliably find differentiation

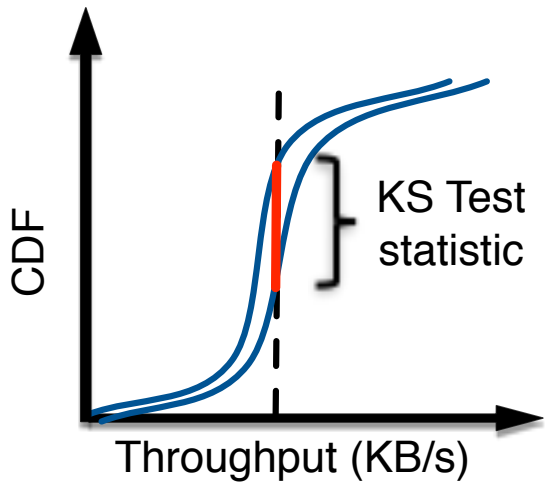

How to tell if there is differentiation?

- Developed and validated new detection technique
  - Back-to-back tests of same trace
    - Includes VPN and random payload (but not ports)
  - Send only at recorded rate (this is not a speed test)
  - Statistical tests to reliably find differentiation

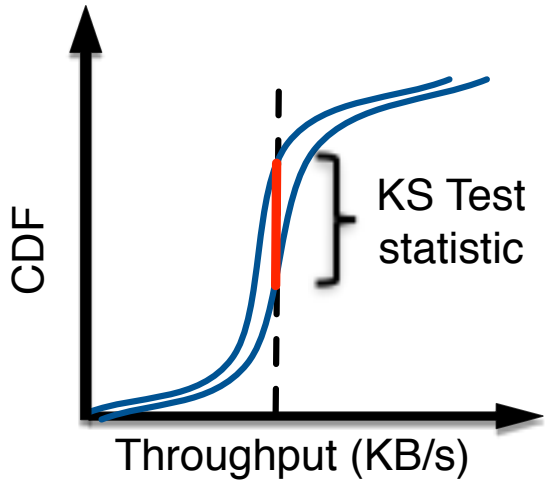

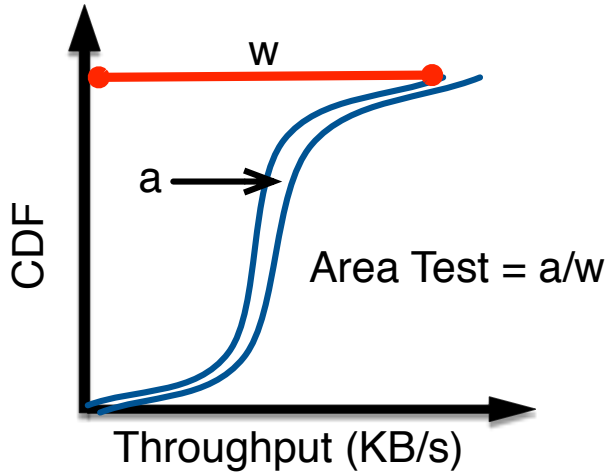

Mapping to QoE

- Currently measure throughput, loss, delay, jitter
  - Some clear mappings to video streaming **bitrate**
  - Fairly clear mapping to VoIP
  - … but unclear how it impacts user-perceived performance
- Key challenge: Applications are **adaptive**, servers **dynamic**
  - Users may not perceive impairment
- Current focus: Expose QoE metrics
  - YouTube does this, Netflix does not
  - What is the right format?
  - How do we know if users notice?

Summary

- Differentiation and its impact on QoE important to solve
- Detecting differentiation requires care
  - What you send (trigger it), how you send it (avoid it)
  - Detection approach should be resilient to noise
- Need a deeper understanding of impact on QoE
  - Adaptive applications pose a challenge
  - Potential approach is combination of
    - Exposing more QoE data from applications
    - Building better models to map QoS to QoE

Questions?

http://dd.meddle.mobi

Check out our app!