The bridge to possible

The New, Encrypted Protocol Stack & How to deal with it

Adding Real Value to (Mobile) Networks

Andreas Enotiadis, MIG Sales CTO
Bart Van de Velde, Sr. Director, Engineering, Networking CTO Office

cisco Live!

BRKSPM-2024

In memory of and based on the brilliant work of

Mark Gallagher

(14/09/1966-17/09/2021)

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated until February 24, 2023.

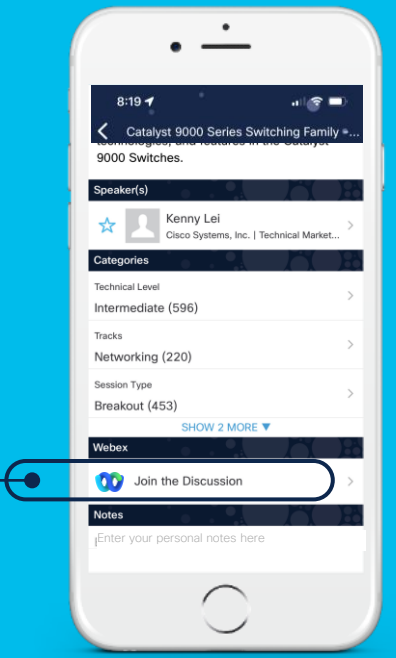

Agenda

- The New Internet
- The New IP protocol stack, and New Traffic Behaviour
- This is relevant to Service Providers
- Dealing with the new reality
  - Toolbox
  - Use cases

The New Internet

A bit of history…

In the beginning

- HTTP 1.1 (1999)

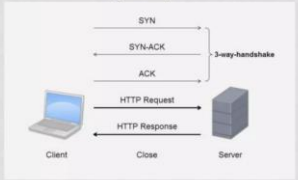

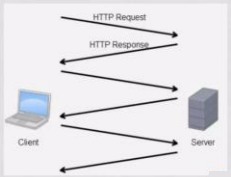

Need for speed

- Google SPDY (2009)
- Over TCP
- Multiplexing (requests not waiting for previous GETs to complete)
- Server Push
- Header compression

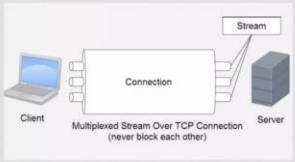

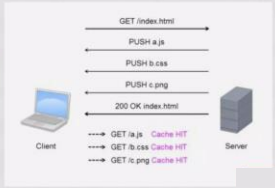

Finally

- QUIC (2013) = Build reliable layer on UDP
  - Widespread deployability in today's internet
    - Middlebox = tipically block anything other than TCP or UDP
    - Modification over TCP = Need kernel changes
    - TCP Extension = Over 10 years or more
  - Packet loss dosen't effects all streams
    - One stream for each request. Packets sent out of order
  - Less RTT
    - Ack dont needed. No handshake. 0RTT
  - Separate congestion window
    - One window for each stream.

QUIC

RFC 9000

The Internet Reality – circa 2020 – Major US Carrier

>90% of
Volume: encrypted

>70% of
Volume: to Cloud

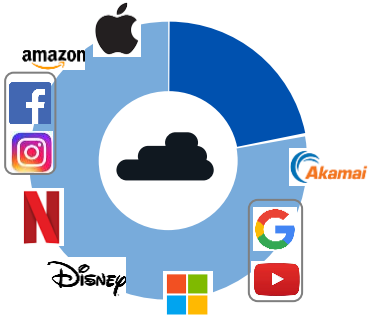

10 Cloud sites
“Elephant destinations”
not “Elephant flows”

~50% of
Flows: DNS

>20% of
Traffic: QUIC

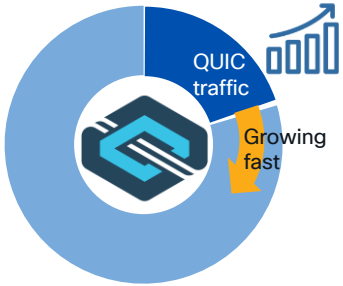

Many small flows
Micro-sessions

- Destination: all-encrypted world
- Cloud: concentrating the Internet

- Content: DNS is the load-balancer
- QUIC: Future Protocol of choice

The Internet is converging on a new normal

It's not one Internet anymore

Traffic Trends**

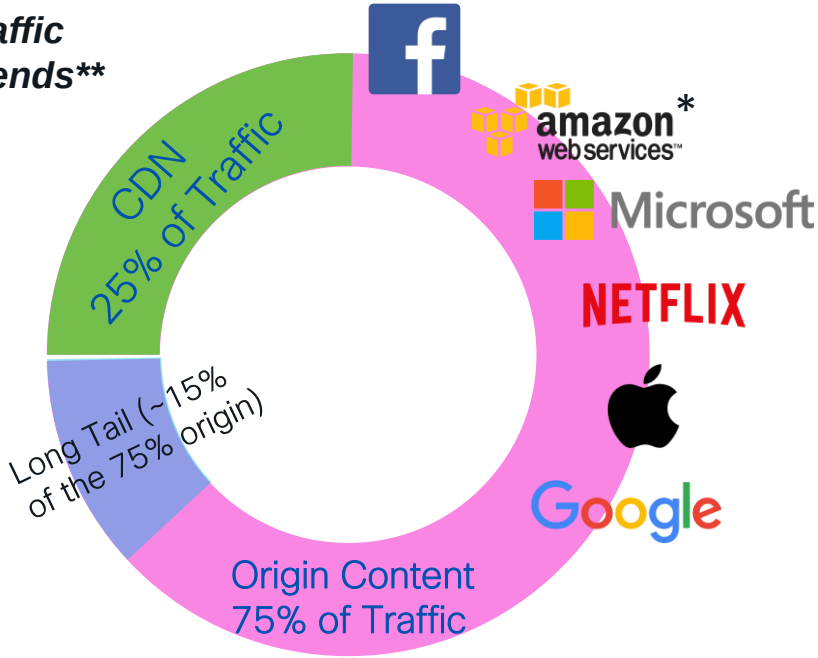

- **12 Cloud Domains = >80% of the Volume**
- **6 of 12 Cloud Origin Content Domains have their own CDNs and/or Secure DNS plans**
- **10 of 12 Cloud Domains Are implementing HTTP/3 + QUIC plans**

Widespread Impact :
Architecture, Network, Devices, Standards *and* Value-chain

**Consumer Traffic Analysis

* Amazon own ~1% of public IPv4

Fast forward 18 months - Tier-1 EU Mobile Carrier

Overall Volume

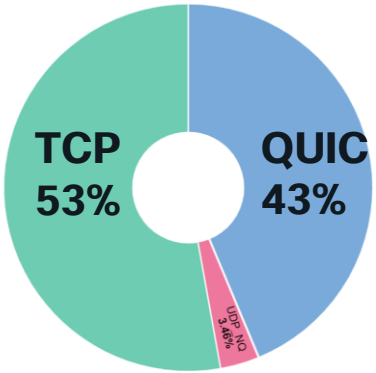

QUIC has doubled in 18 months

QUIC is 43% of total and rising

Volume

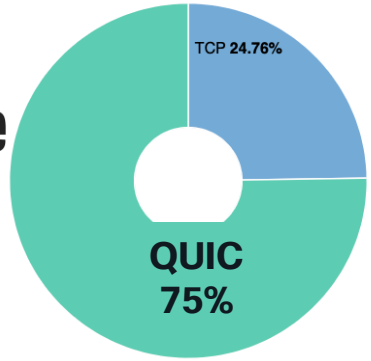

QUIC is "default"

Volume

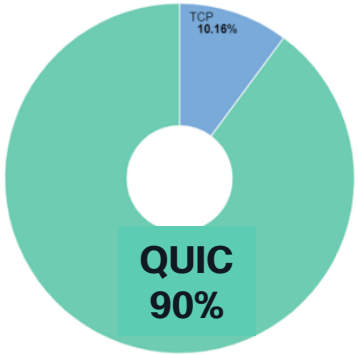

Meta has gone full QUIC

(snapshot 11/2/2022)

Top 5 Apps – QUIC is dominant 80/20 rule now

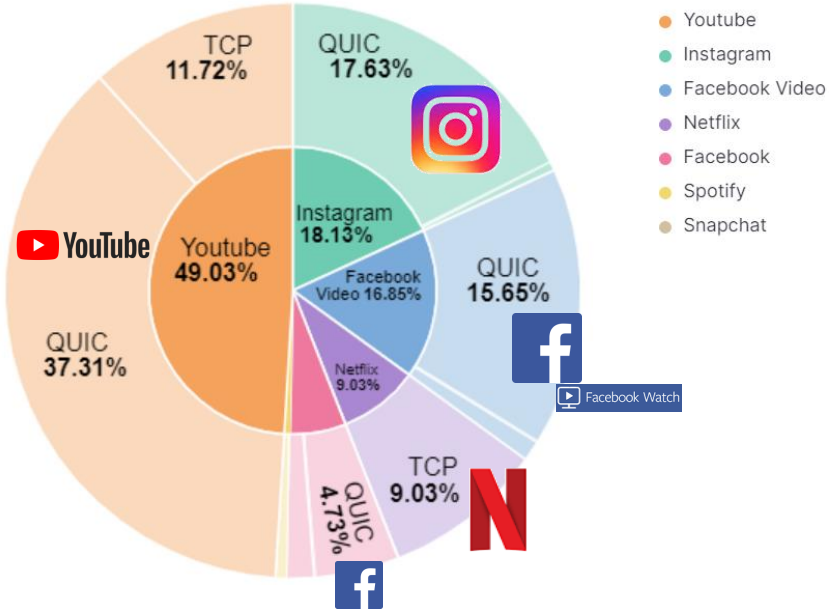

April 10 2022

Network Traffic by Volume and Flows

The big flows that matter are predominantly QUIC

Overall Volume by Apps

Big 5 is 48% of traffic
QUIC is 40% of traffic
“other traffic” still largely TCP, QUIC now visible (4.3%).

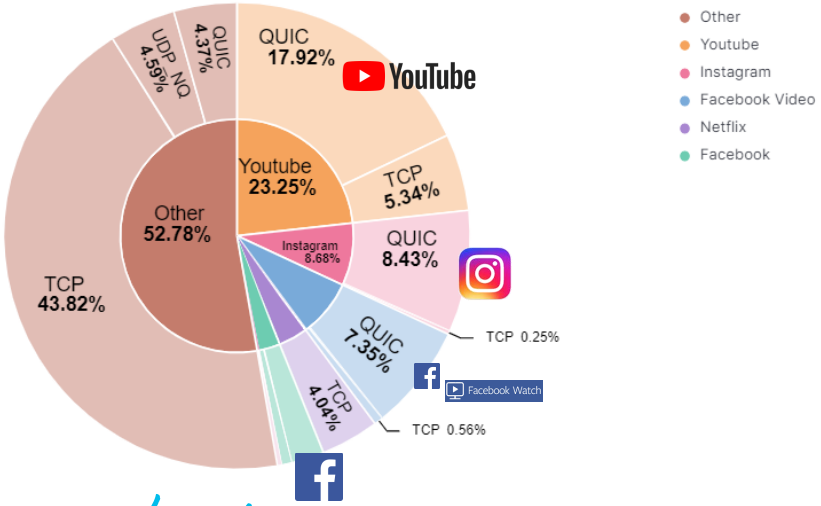

Total Flows by Apps

Lots of TCP sessions (likely IOT related, transactional related)
Big 5 APPs QUIC sessions are very targeted and high efficiency (video related behaviour); fewer but higher in volume

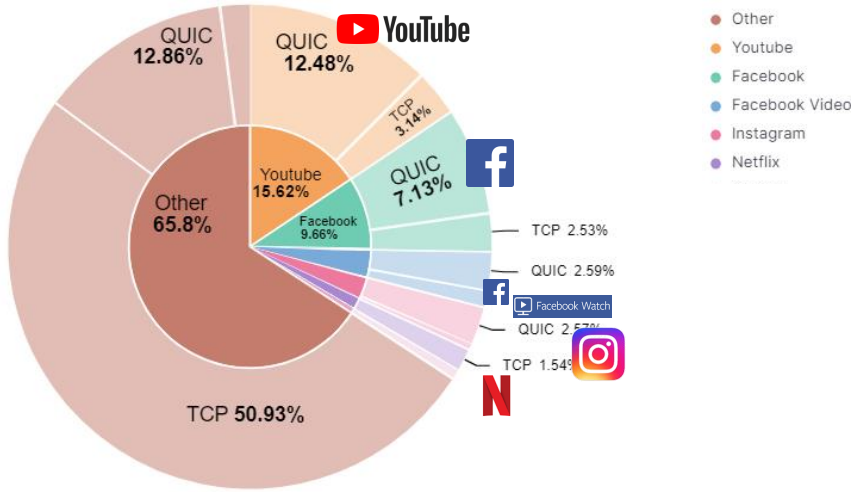

(snapshot 11/2/2022)

LATAM Status
QUIC in the lead now

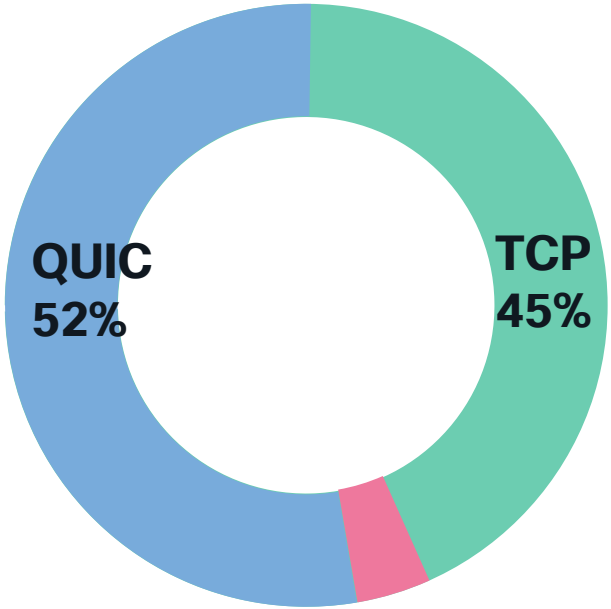

Volume

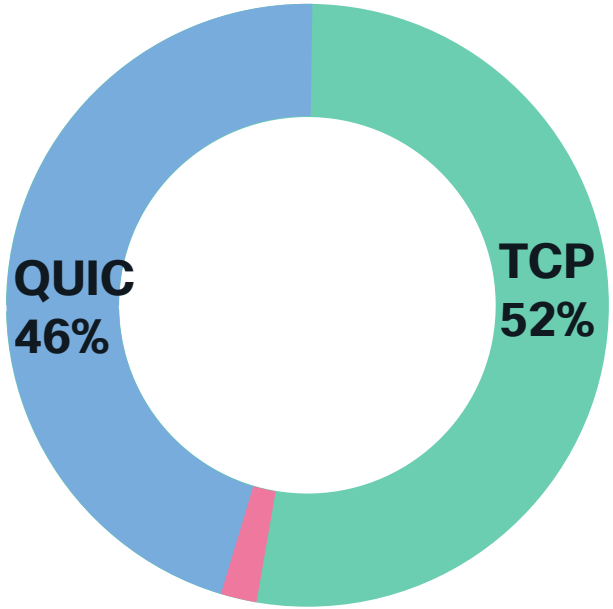

Session count

(snapshot 31/01/2023)

Fixed Broadband: It's not that different – May 2022
if different sources

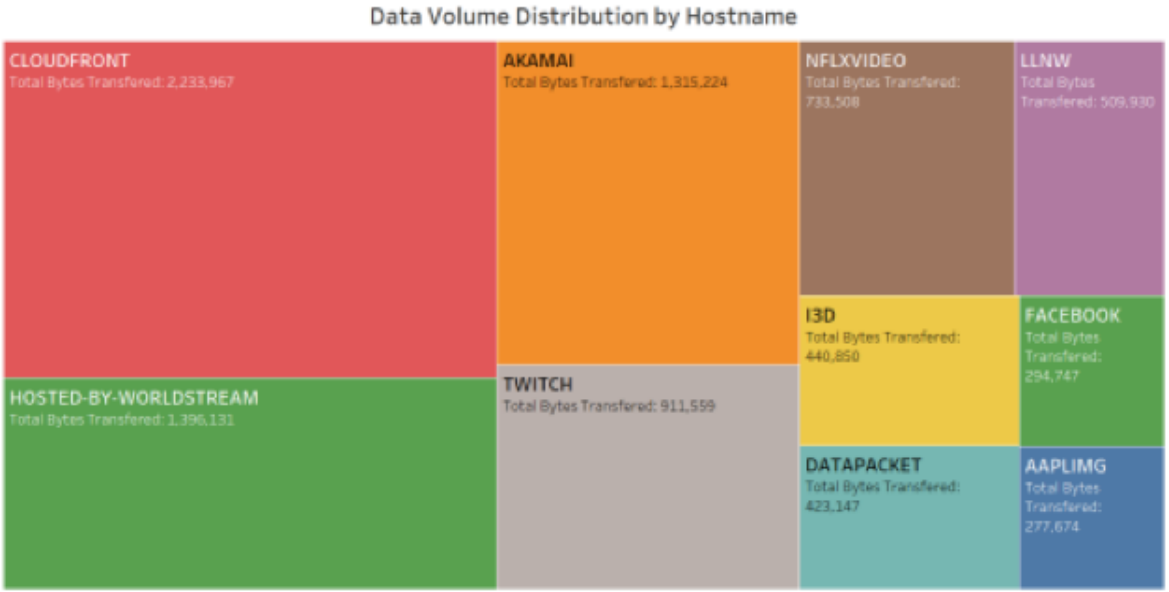

CDN

Hosting

Gaming

Video Streaming

Profile aligned with Fixed Broadband traffic (browser driven traffic)

QUIC : 41% TCP: 53% UDP (other): 6%

*source Tier 1 EU SP

Inversion of the **Internet?**

From connection first	**To** application first
TCP = 90% of Traffic	UDP = 90% of Traffic
100Million+ Important Sites	100's Important Sites*
Some encryption	All encrypted
Fixed Architecture First	Mobile & Cloud First

*Top 12 domains carry 80% of traffic

The old network design assumptions are challenged

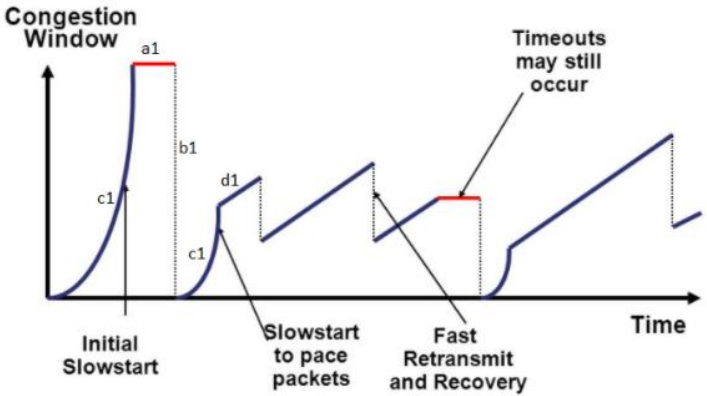

TCP goal is network fairness

Today IP Networks are architected with TCP behaviour as implicit assumption

So when packets are dropped TCP will take care of it at a higher layer

Scenario	Flow	Avg. throughput (std. dev.)
QUIC vs. TCP	QUIC	2.71 (0.46)
	TCP	1.62 (1.27)
QUIC vs. TCPx2	QUIC	2.8 (1.16)
	TCP 1	0.7 (0.21)
	TCP 2	0.96 (0.3)
QUIC vs. TCPx4	QUIC	2.75 (1.2)
	TCP 1	0.45 (0.14)
	TCP 2	0.36 (0.09)
	TCP 3	0.41 (0.11)
	TCP 4	0.45 (0.13)

*** Source : APNIC**

QUIC goal is "MY App" performance

Where are the IP Network Design assumptions wrt QUIC ?

The New IP stack

New Stack, New Behaviour

cisco Live!

An application driven global transition
HTTP/3 Stack = UDP+QUIC+TLS

Old App Stack

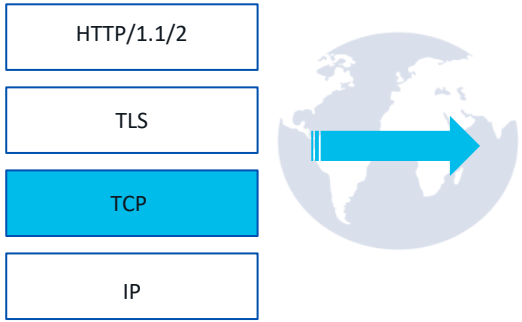

New App Stack

QUIC – RFC 9000
HTTP/3 – RFC9114

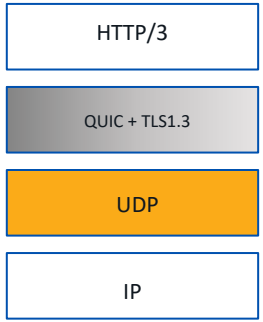

DoH

DoT – RFC7858
DoH – RFC8484

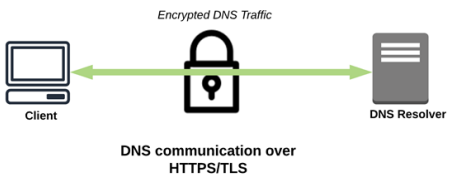

eSNI / ECH

RFC8744

Large Scale Adoption

IETF RFC 8446 – Underpins all others

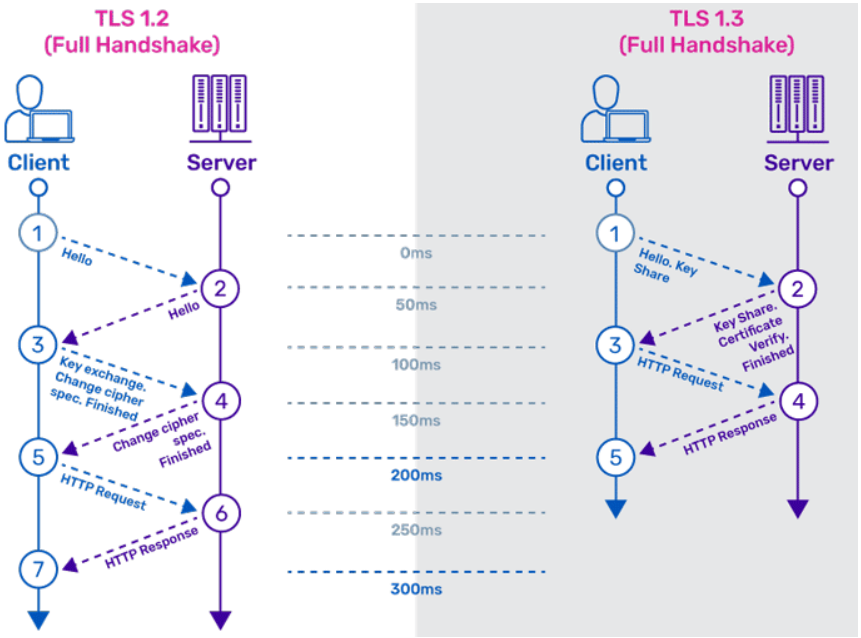

A Faster TLS Handshake

- Simpler, Stronger Cipher Suites
  - No Compromised algorithms
  - Only PFS (Perfect Forward Secrecy)
  - No Renegotiation
- 0-RTT (returning user to server)
  - Ideal for Mobile connections
  - No RSA (or other) Static Keys
  - PSK(pre-shared key) for 0-RTT – session resumption

IETF RFC 9000 – The new “TCP”
Optimised for RPC

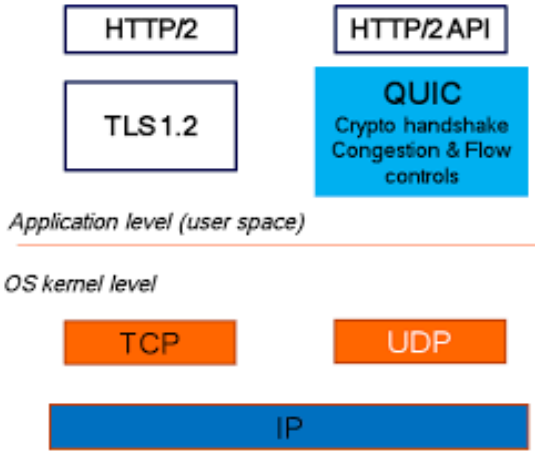

User Space
TLS 1.3 Encrypted

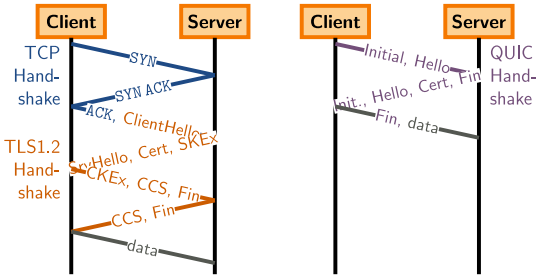

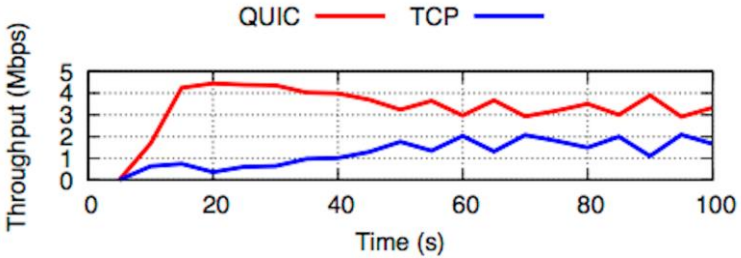

Fast of the blocks

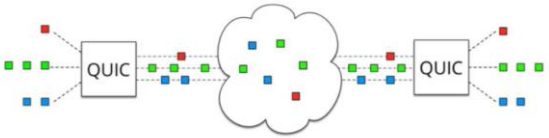

Deliver at all cost
(Multiplex, no-HOL)

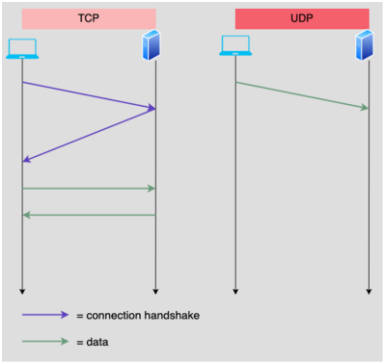

UDP is ”fire and forget”
App controls the rest

Moves Control of the User Experience to the App

Apps do not play nice – they will deliver over everyone else

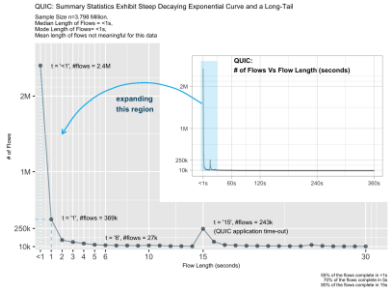

70% of interactions complete in <5s**

Scenario	Flow	Avg. throughput (std. dev.)
QUIC vs. TCP	QUIC	2.71 (0.46)
	TCP	1.62 (1.27)
QUIC vs. TCPx2	QUIC	2.8 (1.16)
	TCP 1	0.7 (0.21)
	TCP 2	0.96 (0.3)
QUIC vs. TCPx4	QUIC	2.75 (1.2)
	TCP 1	0.45 (0.14)
	TCP 2	0.36 (0.09)
	TCP 3	0.41 (0.11)
	TCP 4	0.45 (0.13)

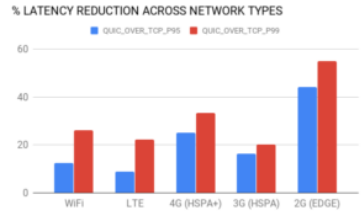

The poorer the network, the better the improvement*

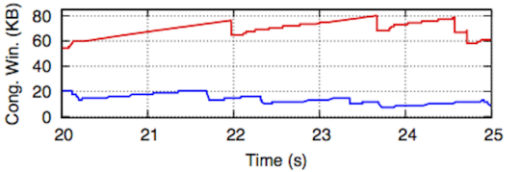

QUIC is "Unfair"***

Impacted Areas

(e.g. mobile)

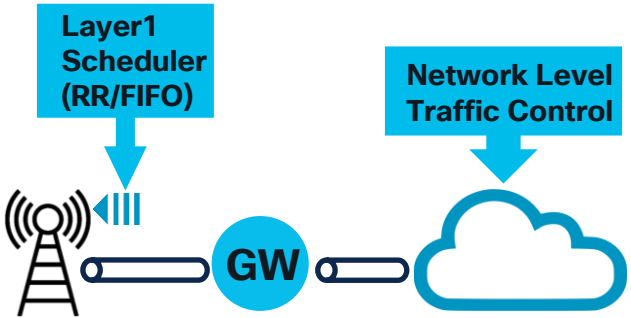

*uber engineering;**Cisco Analysis, cust.data;***APNIC study

Secure DNS – Domain lookup Privacy by default

DoH - RFC8484 is becoming mainstream

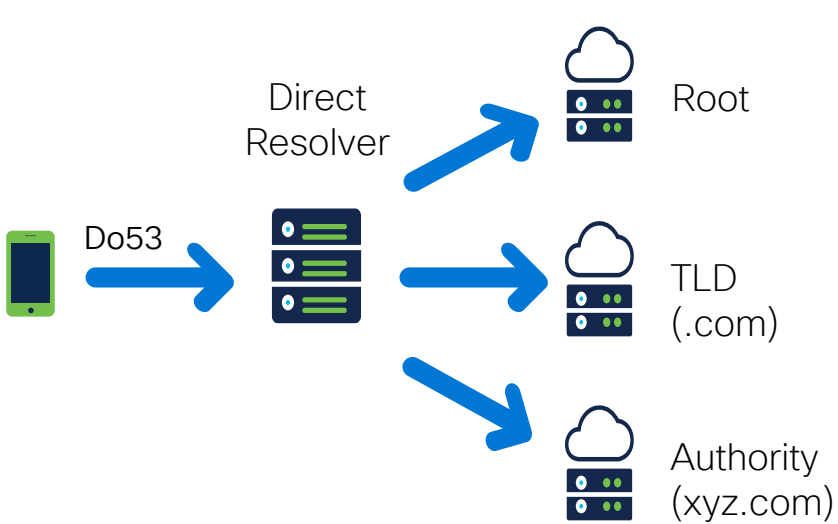

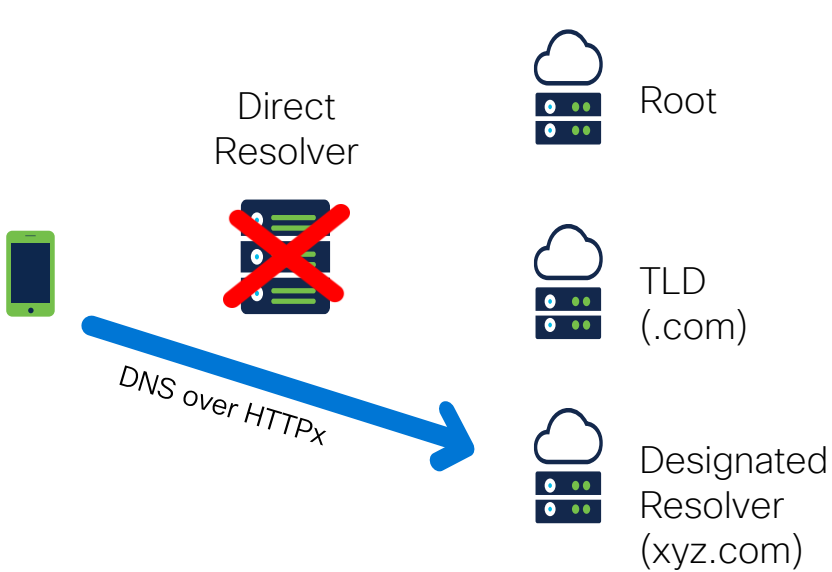

From: DNS Hierarchy + cleartext fields

To: DNS (direct) Connect + ciphered fields

+ DNS is controlled by Applications

Hiding the destination completely - eSNI & ECH

RFC6066

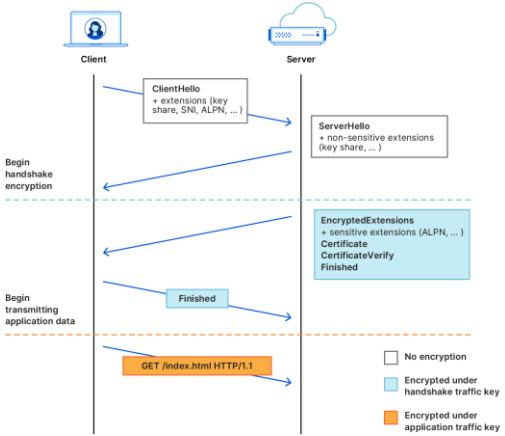

RFC8744

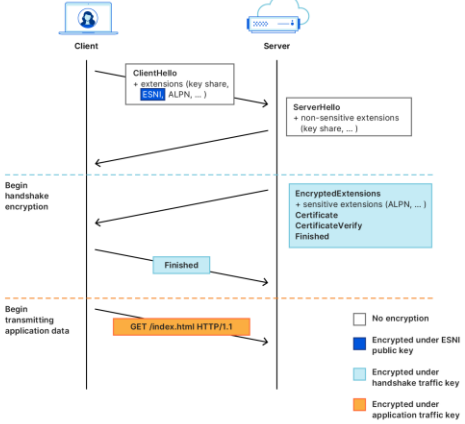

draft-ietf-tls-esni-14

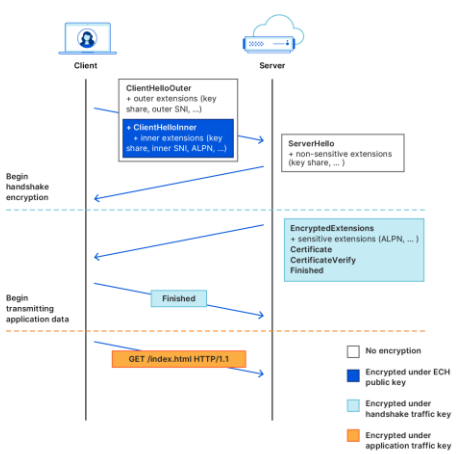

Classic SNI

Destination & Capabilities in the clear

*Application layer Protocol Negotiation

eSNI

Destination leaked via DNS

ALPN* still in the clear

ECH

Only CDN address visible – DoH

SNI & Capabilities fully encrypted

Bringing this all together

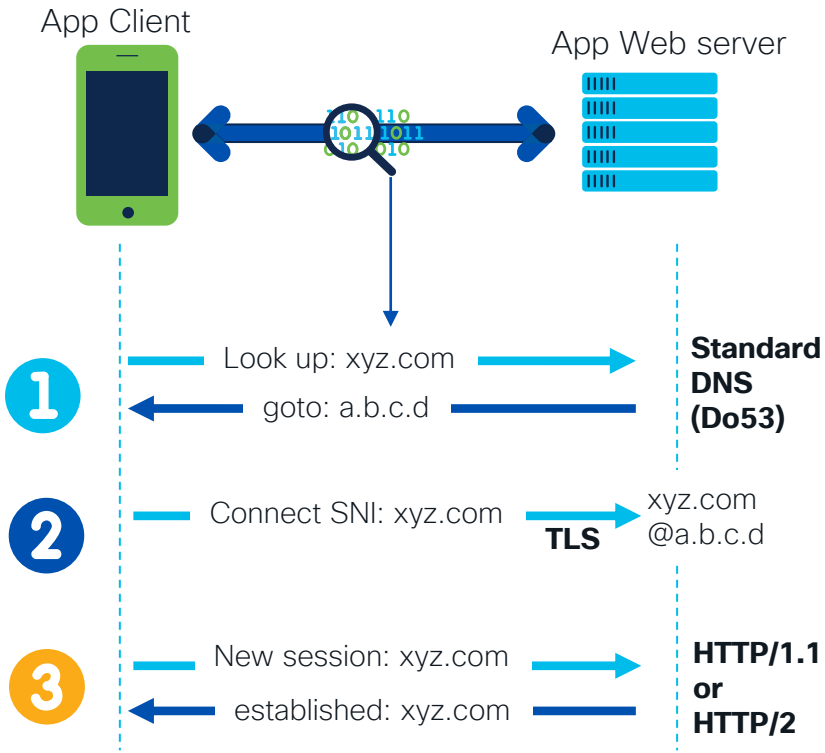

- ✅ Well-understood protocol stack
- ✅ Foundation of **all** web traffic
- ✅ Adopted by Applications
- ✅ Globally scaled

First, apply secure DNS

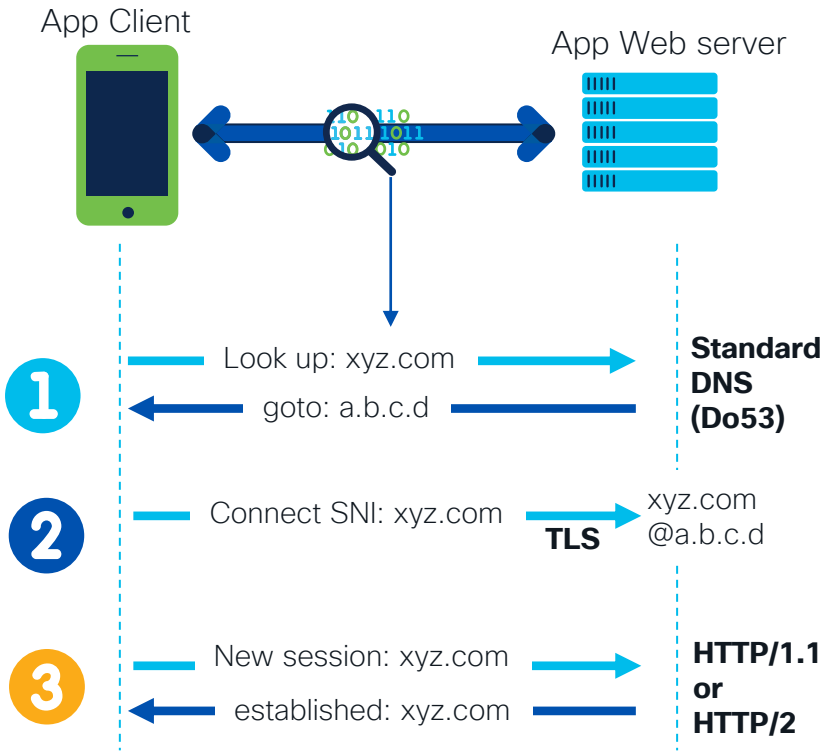

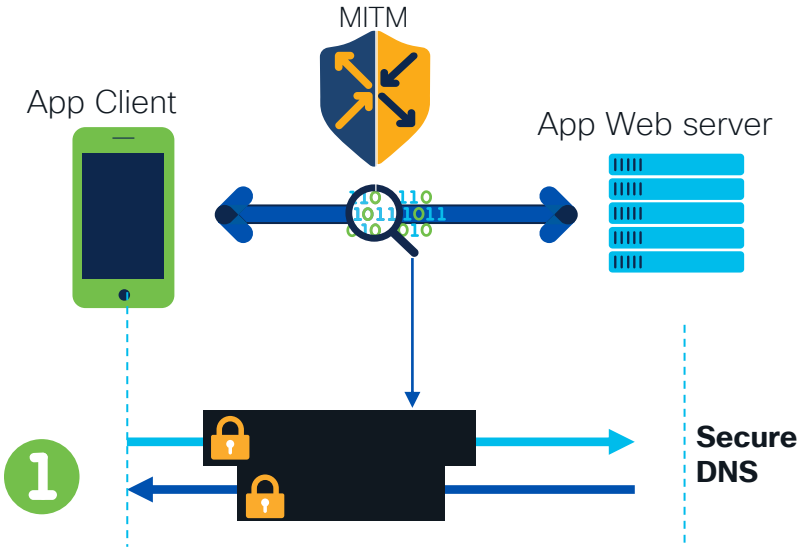

Need TLS and HTTP inspection to recover information

Second, TLS 1.3 Features used

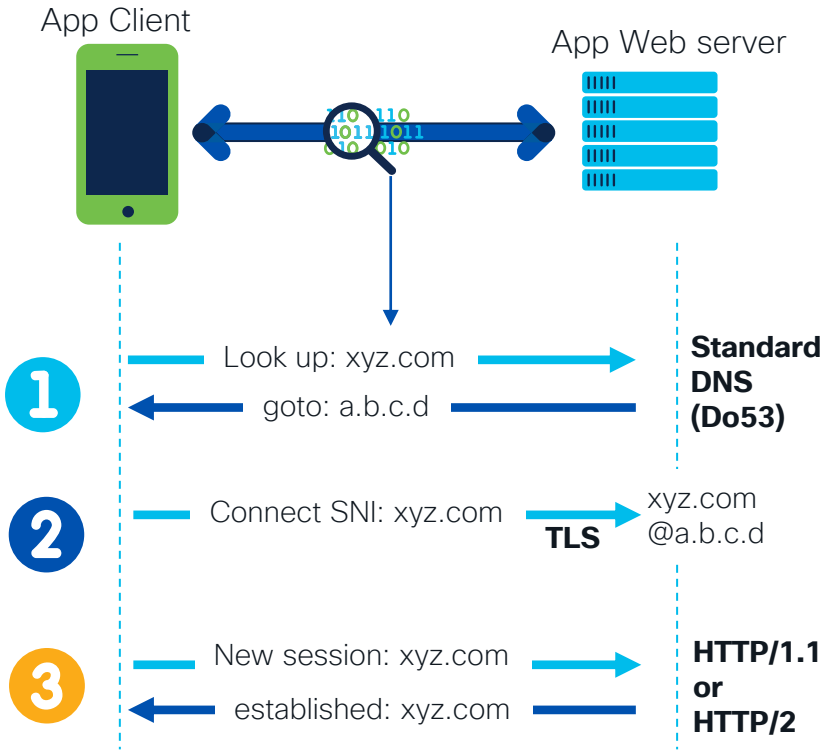

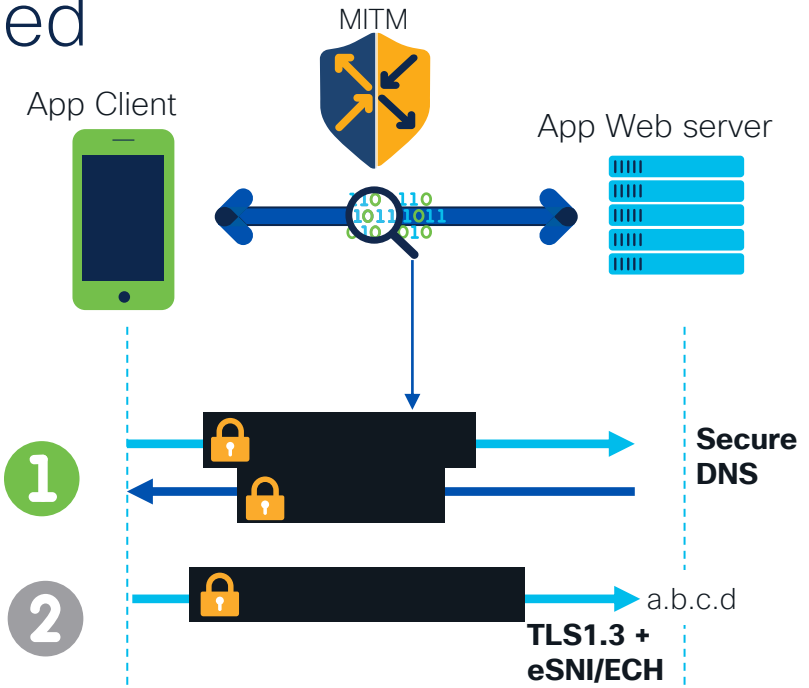

Session setup encrypted but session patterns can provide insights

Finally, HTTP/3 and QUIC multi-session technology

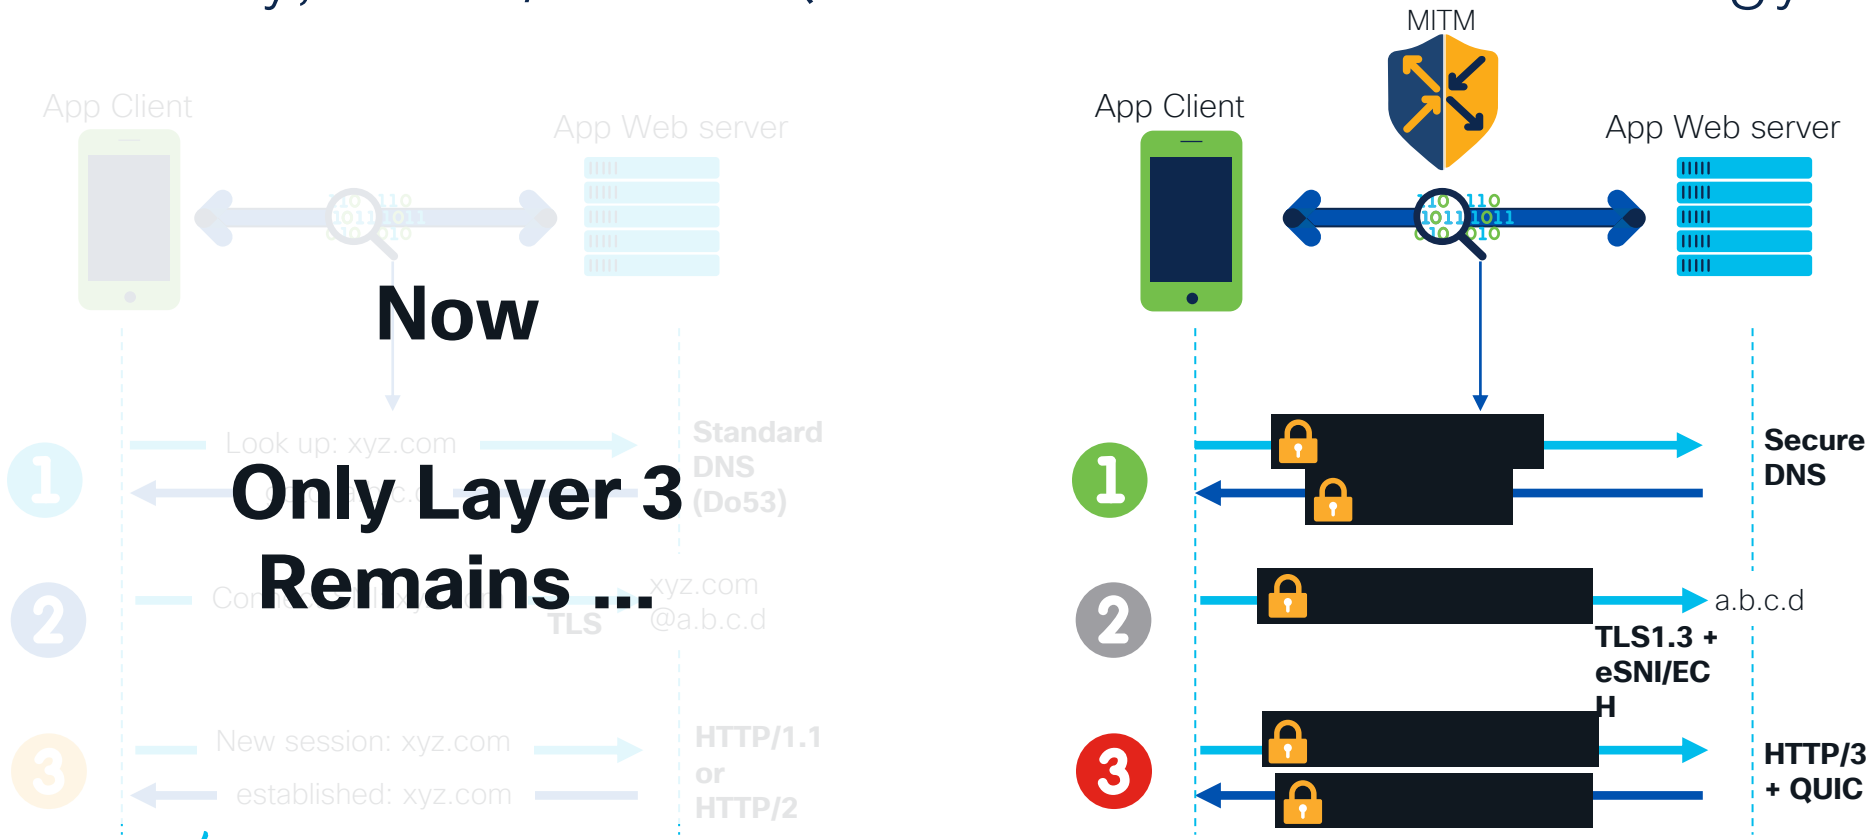

Visibility is lost

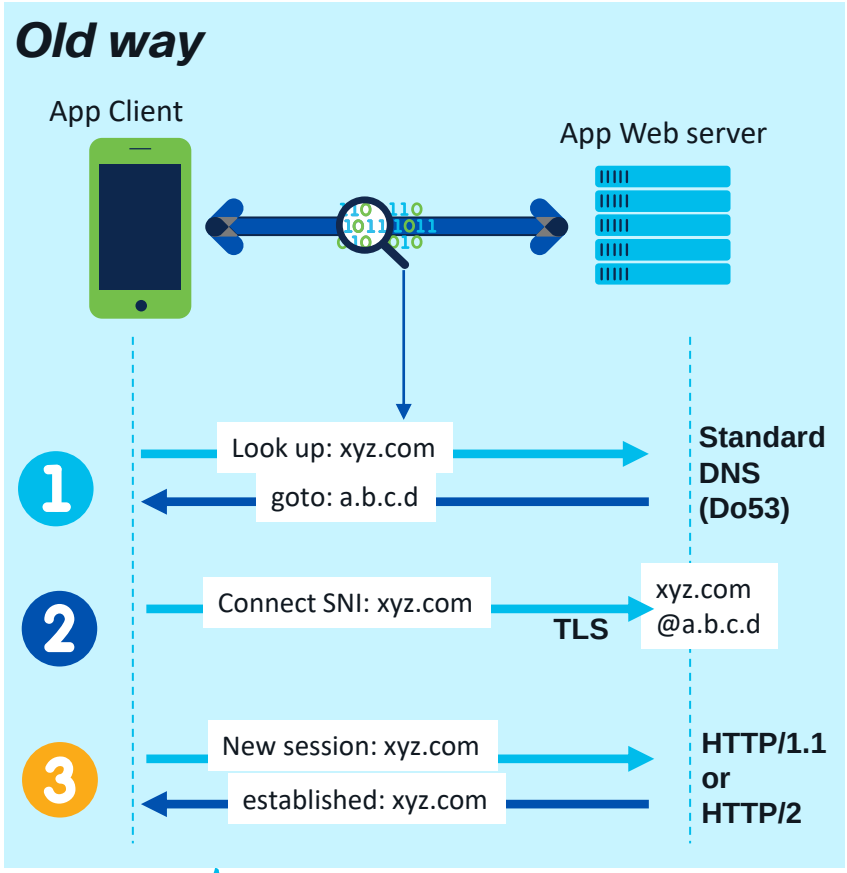

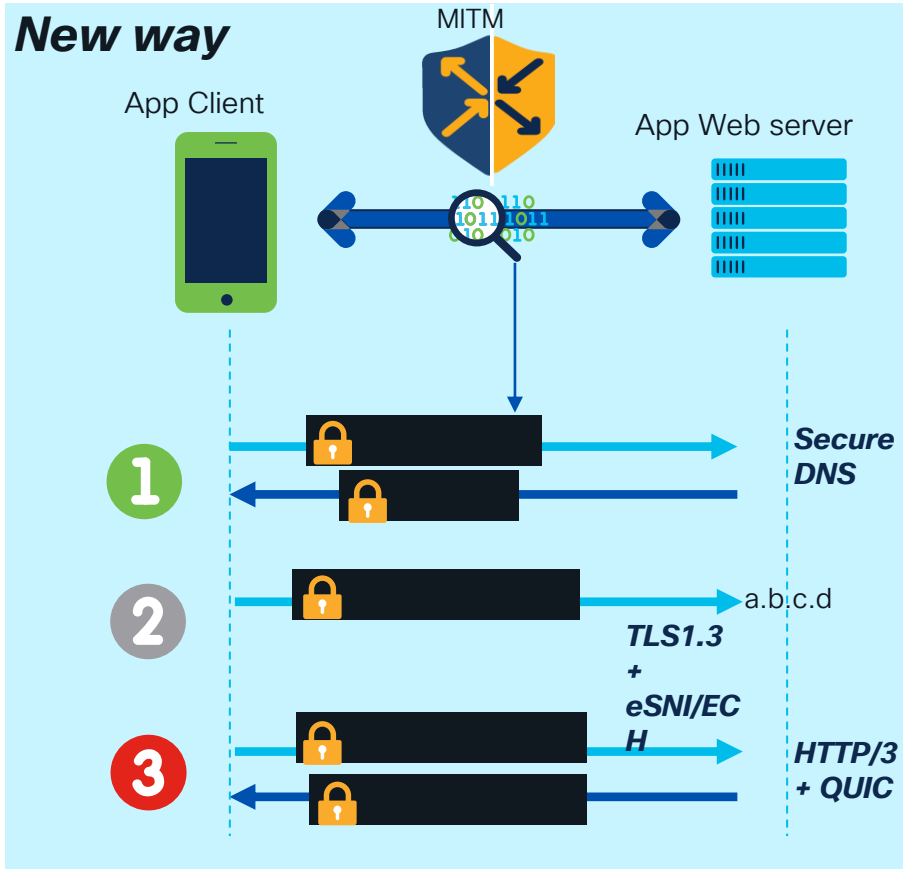

An application driven global transition
HTTP/3 Stack = UDP+QUIC+TLS

Old App Stack

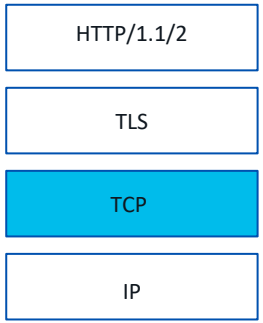

New App Stack + DoH + eSNI / ECH

QUIC – RFC 9000
HTTP/3 – RFC9114

DoT – RFC7858
DoH – RFC8484

RFC8744

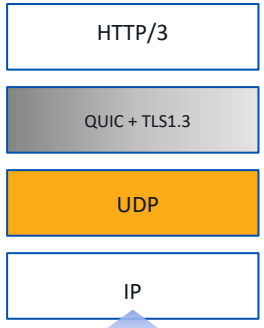

- *Improved Security*
- *Multi-session*
- *Improved QoE*
- *APP friendly design*

Application Controlled DNS
DNS Traffic not observable

Target Domain is opaque / unobservable

Google & CloudFlare serve 50% of global DNS requests
Both support DoH
All major OSs & Browsers support DoH (Firefox Defaults for US to CloudFlare)

DPI Ineffective

including alternative hints e.g. DNS or SNI analysis

Large Scale Adoption

But there's more:
New Stack, New Behaviour

The Old way:
Network forwards and TCP does the rest

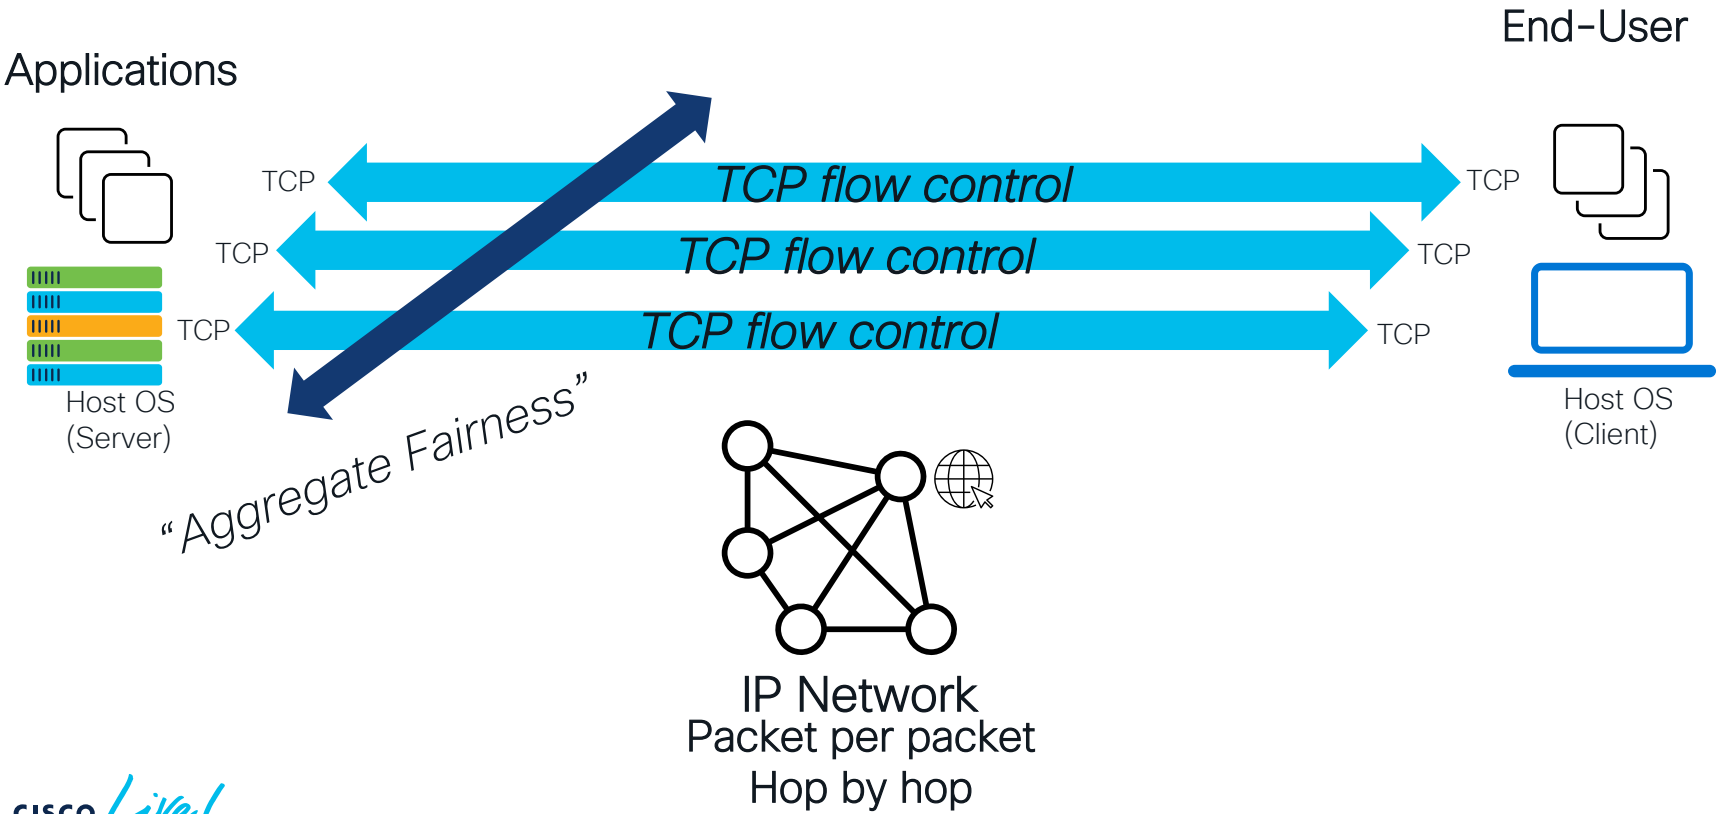

The New Way
QUIC Protocol inside the APP

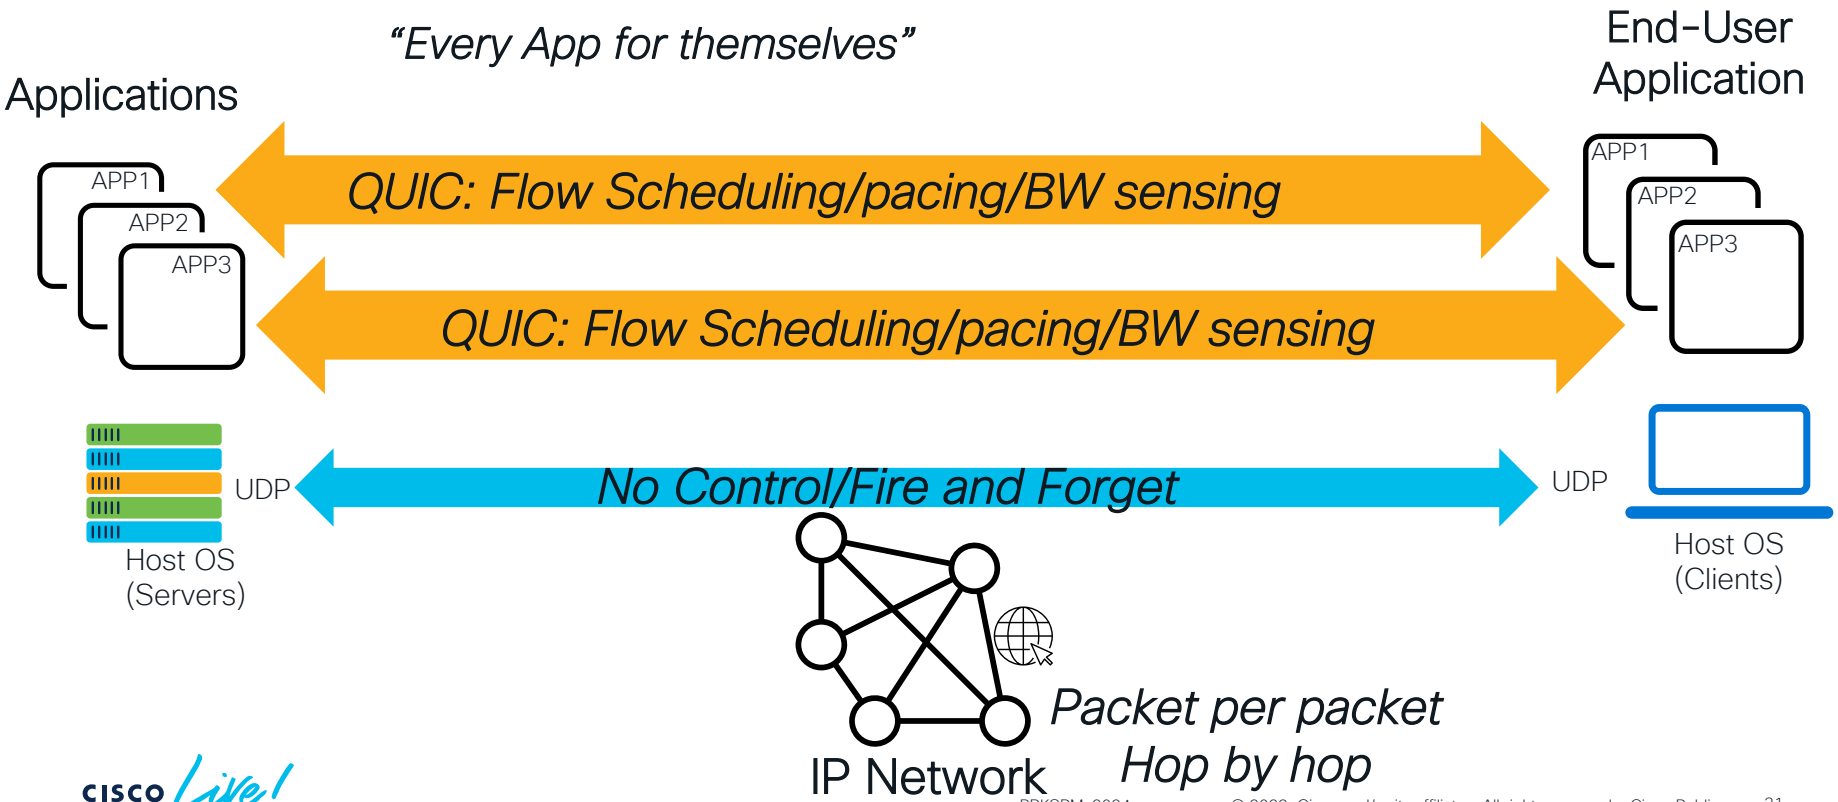

Dropping packets means back-off?

Our queuing design assumptions are losing validity

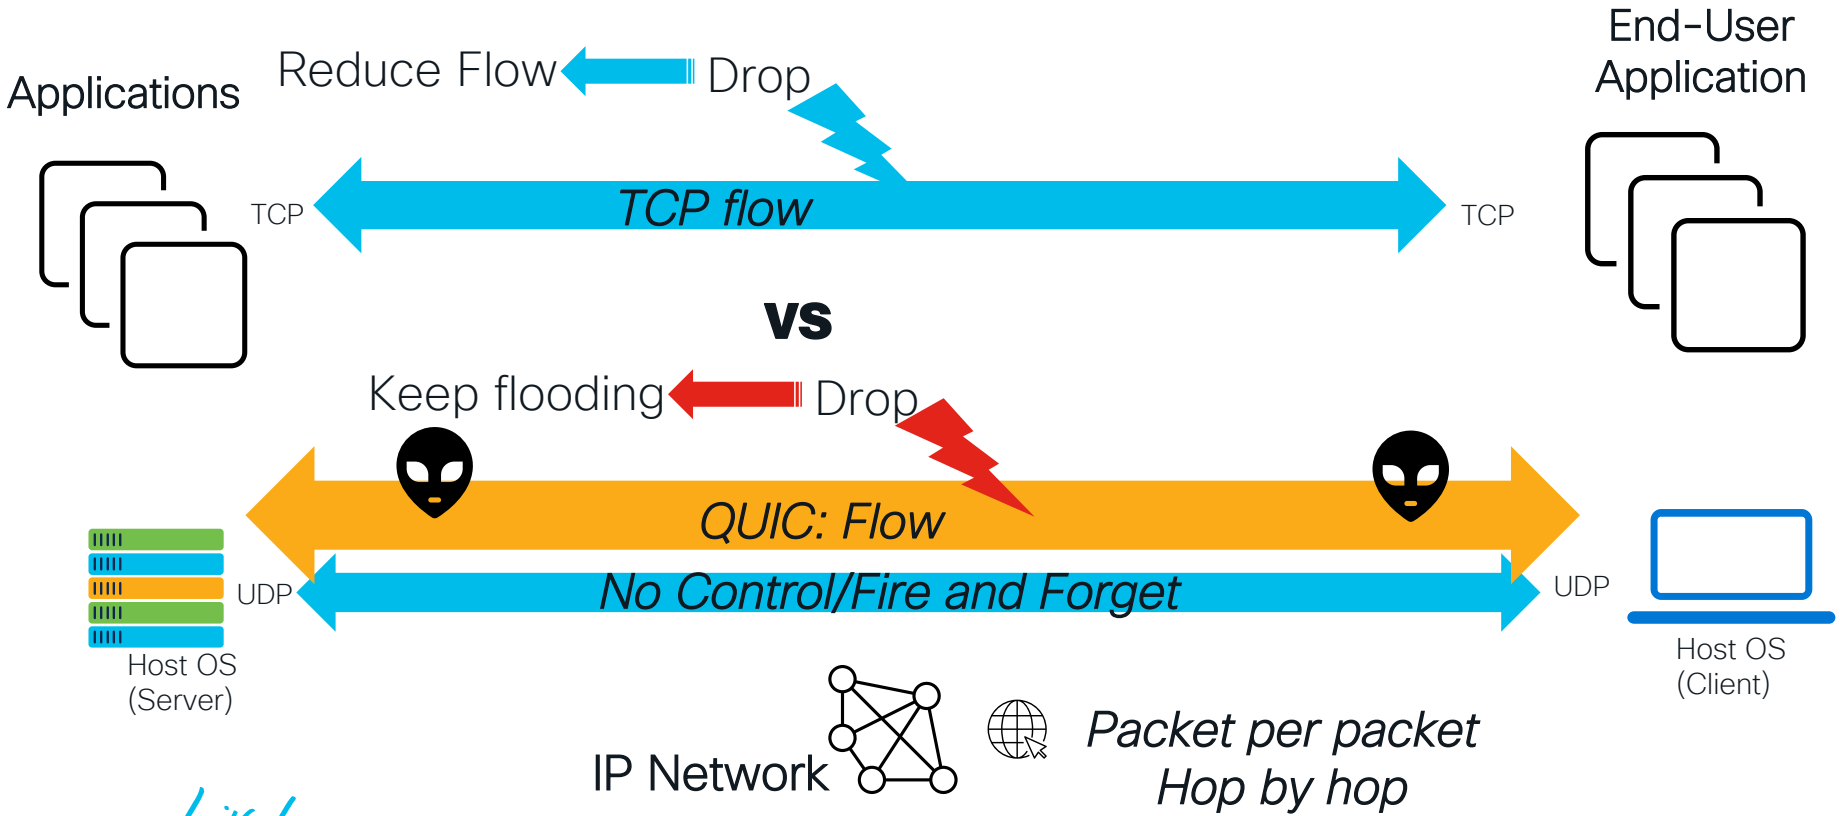

QUIC/H3/DoH stack is in business

Content Delivery **Security** **Privacy** **Loadbalancing** **App Infrastructure** **App Experience**

Standards are there
Continuous standard adoption

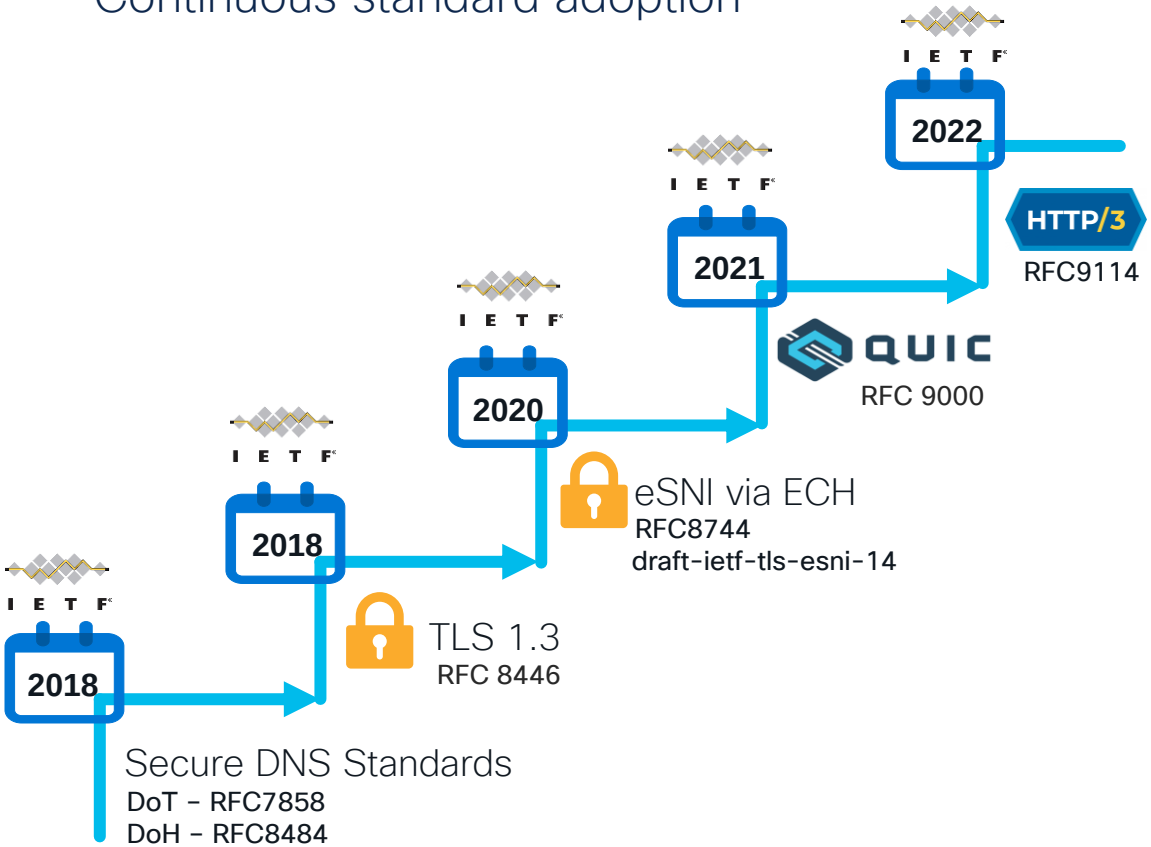

- ✔ At scale, in production
- ✔ Client — chromium, Microsoft Edge, Mozilla
- ✔ Application
- ✔ Cloud

cisco Live!

In the new, fully encrypted world

- The only certain datasets are:
  - UE & user data
  - CDN (GCP, AWS, Apple, CloudFlare….)
  - Temporal flow behaviour
- The information we can glean is
  - Type of flow (Video Stream, download, …)
  - Form of content (short form Video – e.g. Snap , mid form Video – e.g. YouTube, Long form Video – e.g. Netflix
- Everything else (e.g. app read from eSNI) is a short term bonus that will go away

- Most use cases and enablers that rely on higher level protocol or specific application recognition are affected at some level
  - Traffic Management
  - Security
  - Traffic Optimisation
  - Traffic Reporting
- Some use cases (esp. Optimization) can be rethought for the encrypted world
- 80-20 rule will apply

This is relevant to Service Providers

The consumers will observe benefits

QoE Drives QUIC Adoption

facebook

1.8B DAU – 3B MAU
QUIC and H/3 are protocols of choice*

Uber

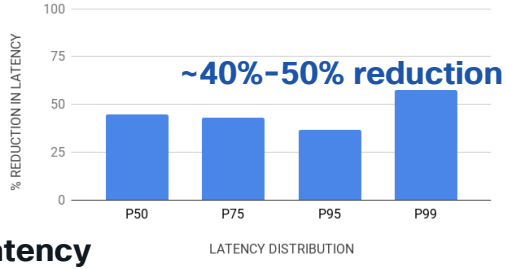

Latency reduced significantly**

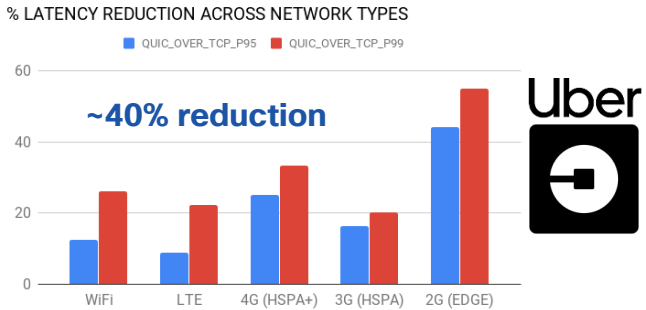

Uber

The more fragile the network, the more QUIC excels**

*source Facebook engineering

** source Uber engineering

Desintermediation of the SP?

Who has the control points to generate the services?

Diffferentiated billing

Differentiated traffic management

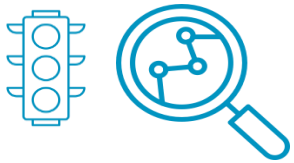

Traffic optimisation

Content delivery & Caching

How QUIC Flows e.g. impair eNodeB Performance

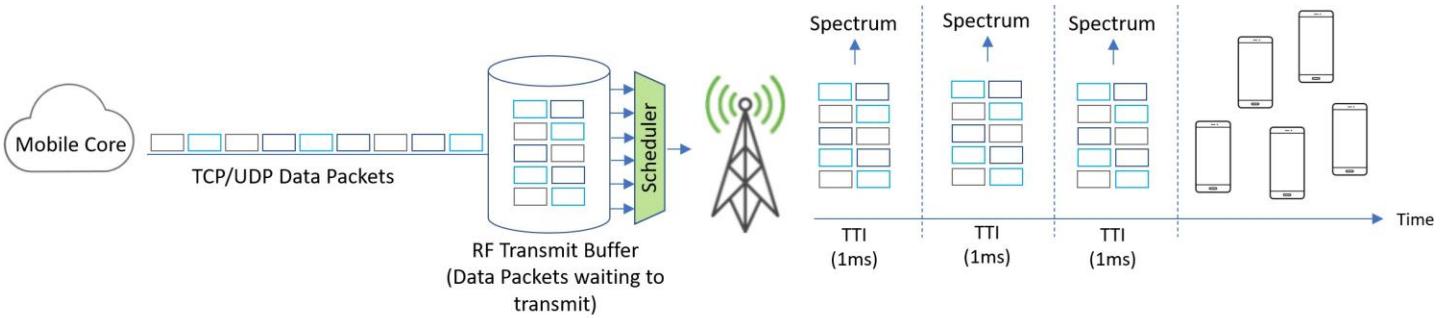

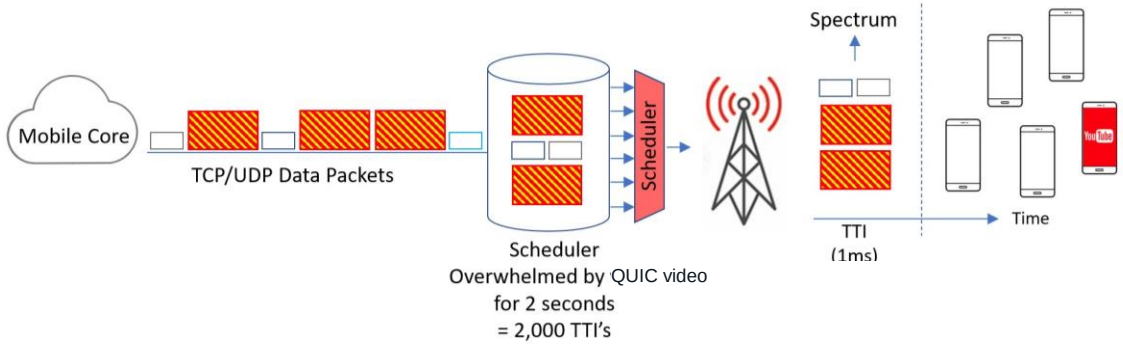

1)When buffer overflows happen at the radios scheduler level the scheduler will drop PDUs

2)The system assumes the majority of the traffic is TCP so the congested flow will back off

3)QUIC flows (as seen previously) ignore it and swamp the link trying for delivery of THEIR app

4)Everyone's user experience (and the cell 'goodput')is impaired…

Multiple parallel internets

Google Global Cache

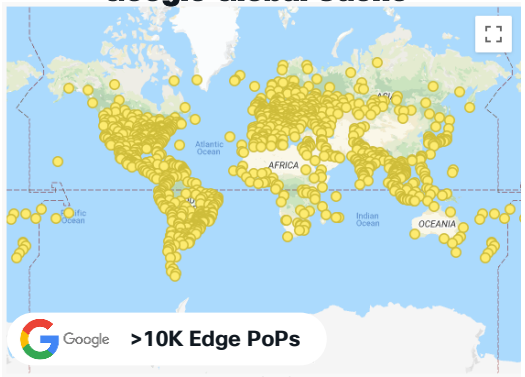

android chrome OS ~3B active

Apple Private Relay

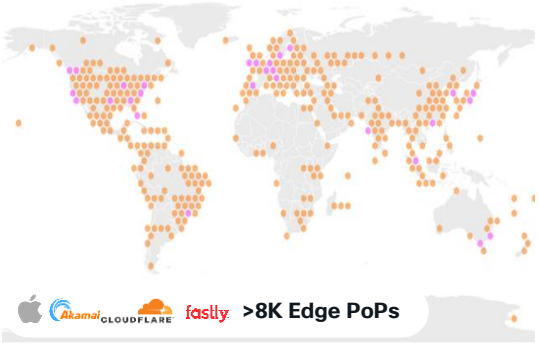

iOS macOS ~1.6B active

Azure Front Door

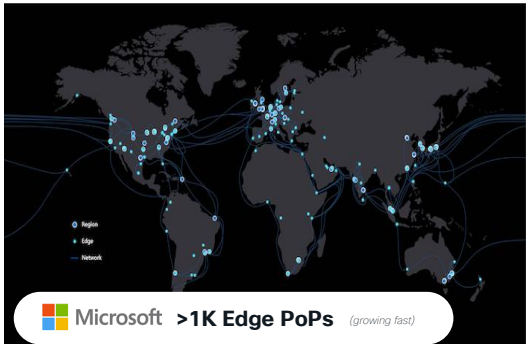

Windows Microsoft 365 ~1.3B active

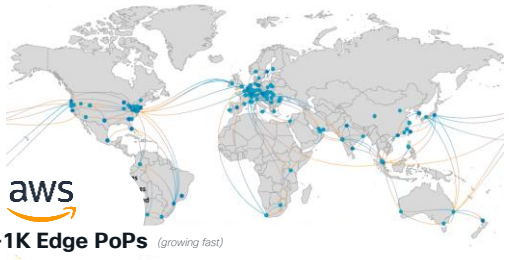

AWS CDN

AWS Cloudfront
AWS Wavelength
AWS Local Zones
AWS Outposts

Builder Community / AWS Sidewalk

Apple Private Relay (iCloud+ Service)

Targets Safari and some App traffic in Phase 1

Example

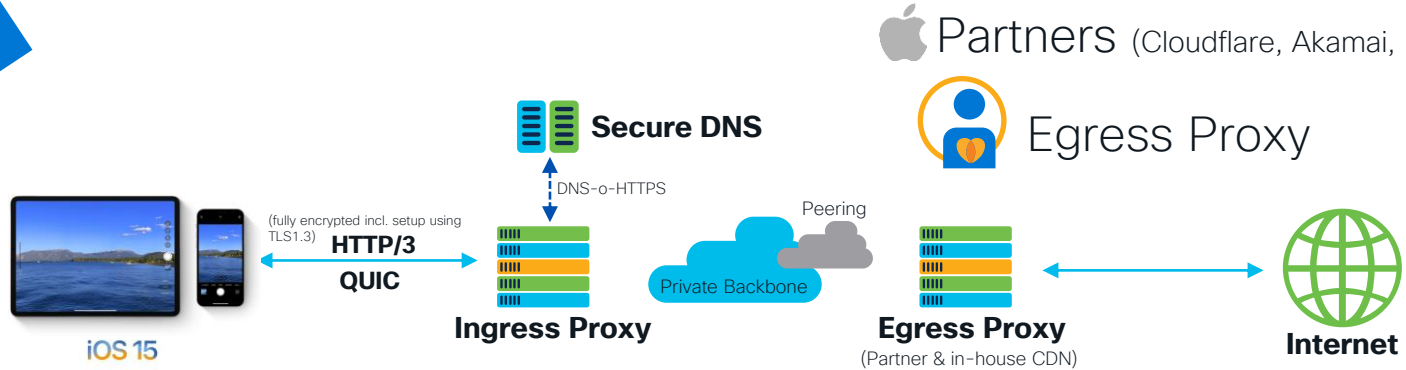

Apple Private Relay
"Monetizing Privacy"

- Ingress Proxy
- Secure DNS
- Private Backbone & Peering
- Native Client Stack & MDM

* From press reports, Not Public

SP Domain has less insights on traffic

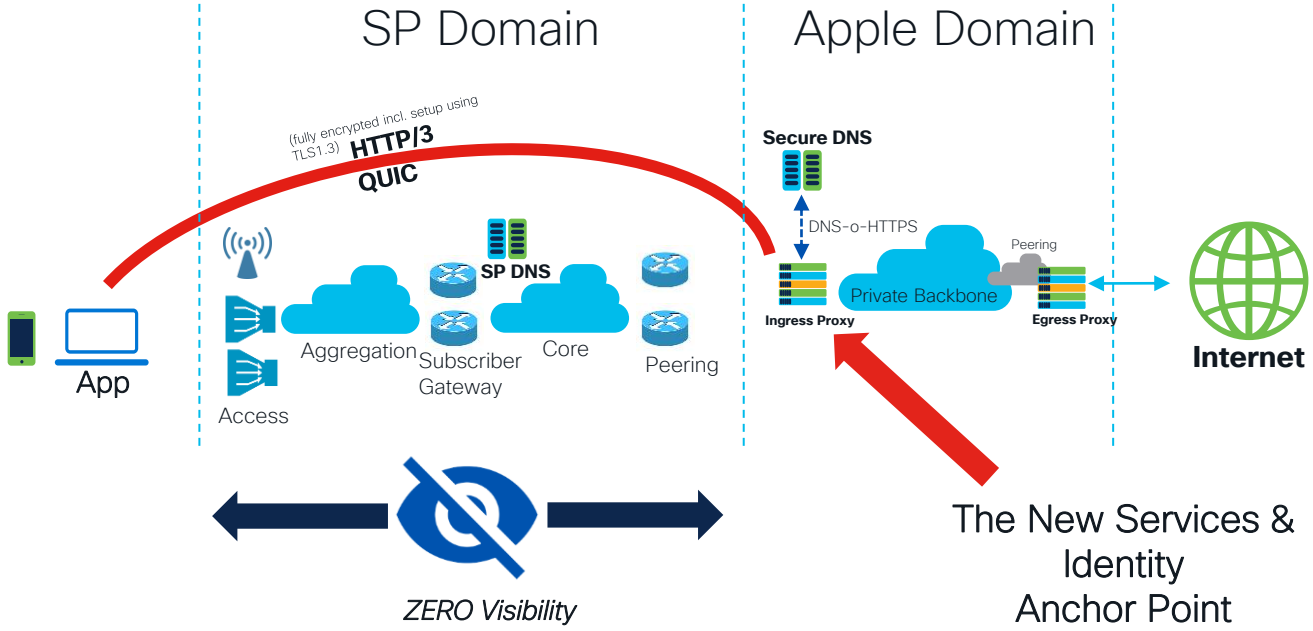

SP Services 'under pressure'

Regulated Services
Site blocking

Differentiated Billing
Zero rated Apps

Traffic Engineering
Peering

Traffic Engineering
SP Edge

Geo-location in the usual way

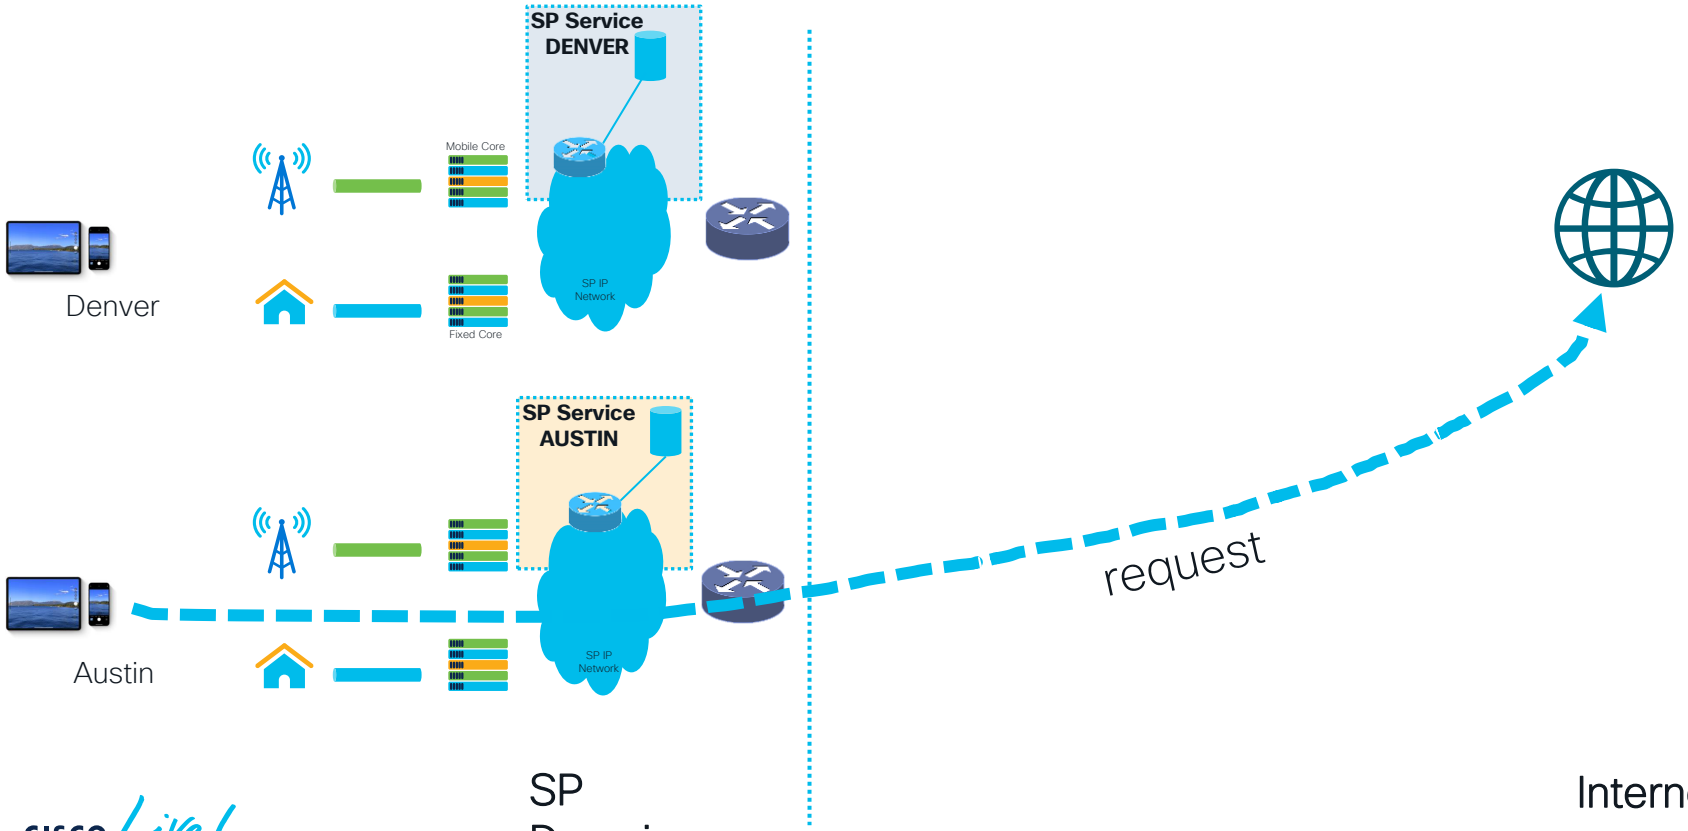

Geo-location in the usual way

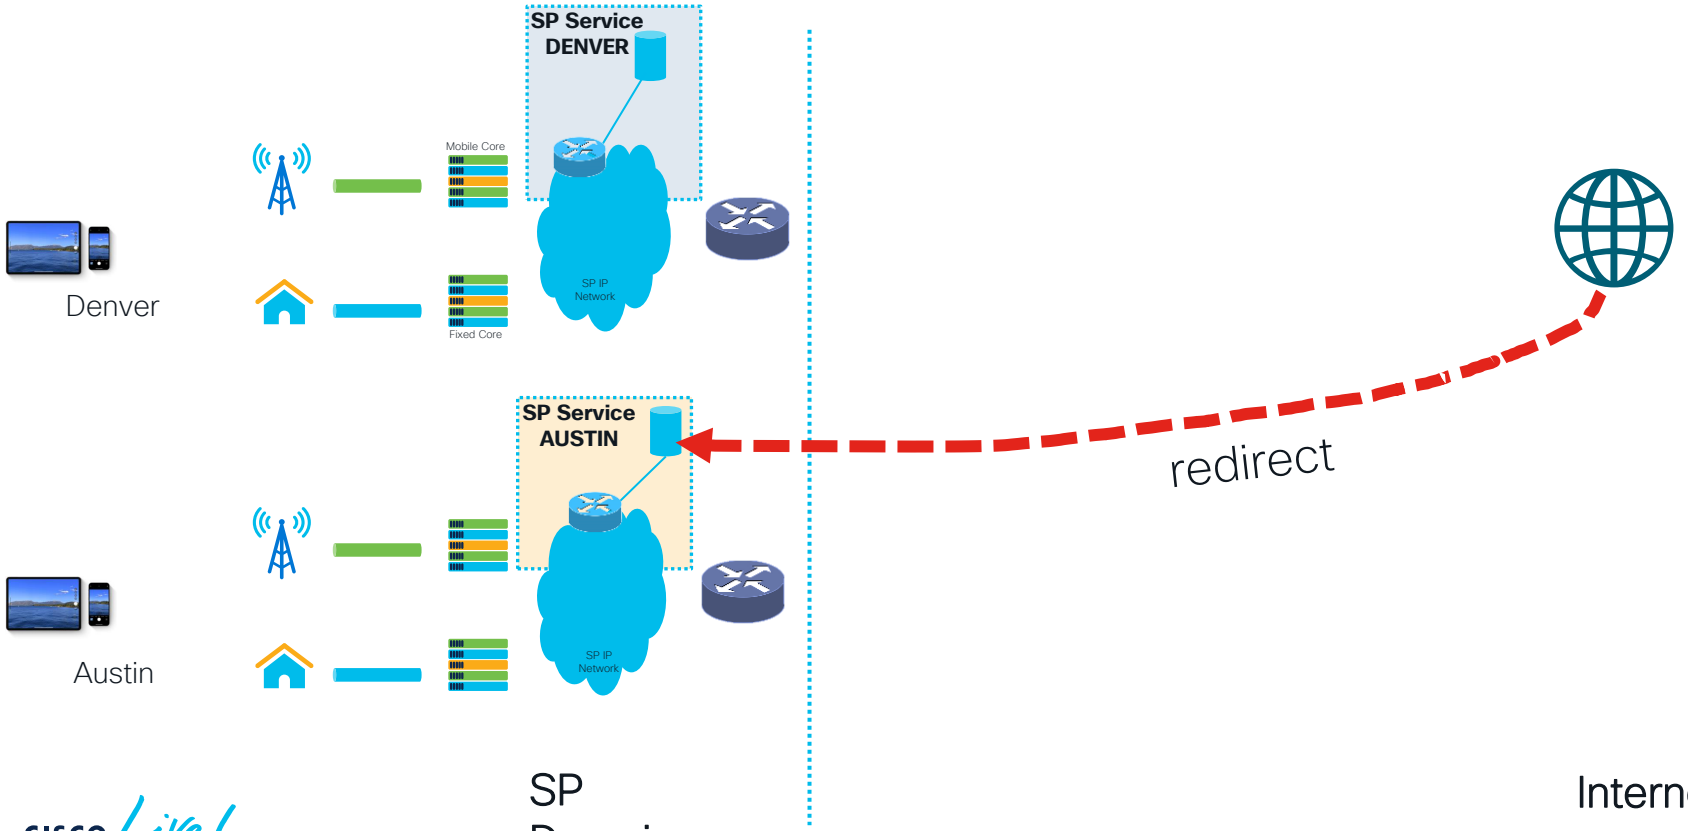

Geo-location in the usual way

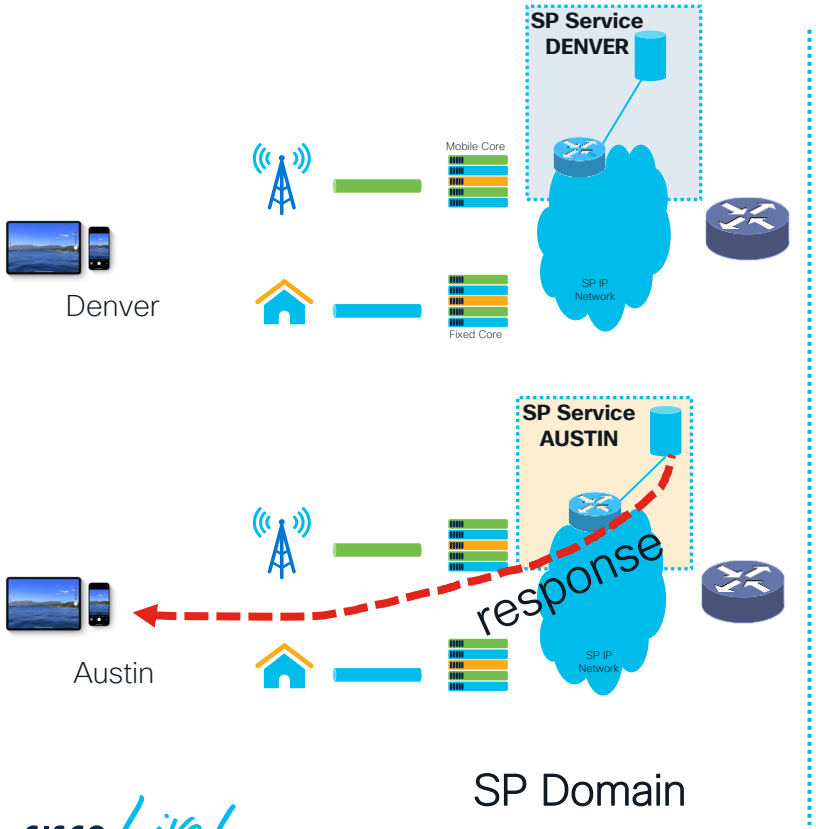

With Apple Private Relay

Tunnel to Apple Network

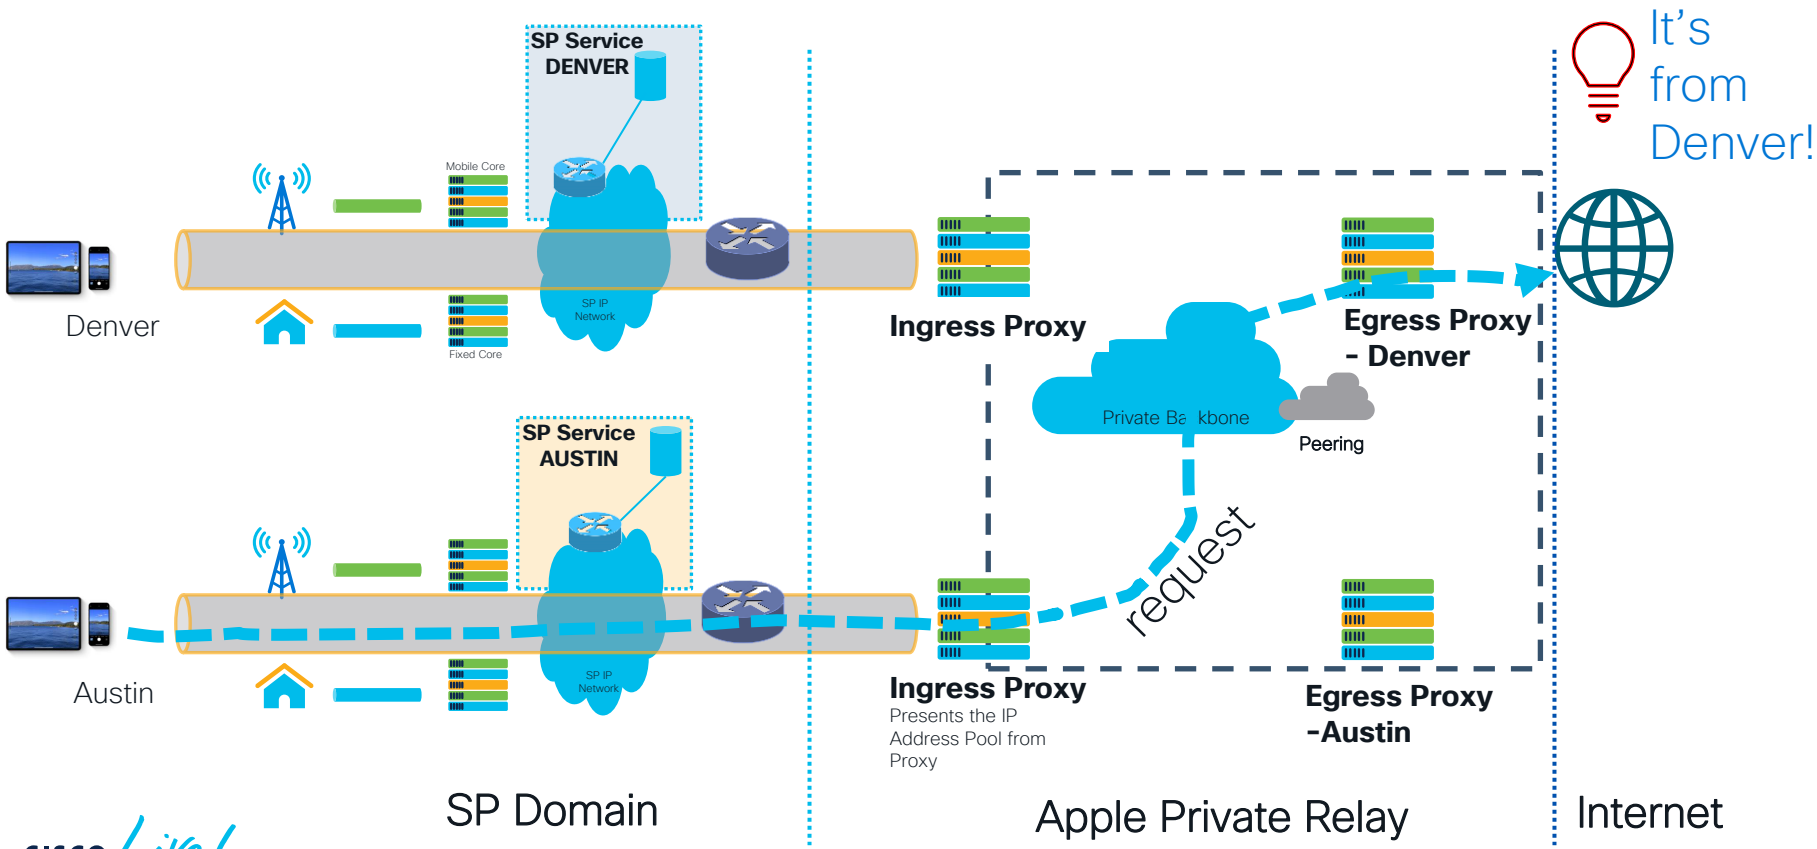

With Apple Private Relay

Location determined by Apple

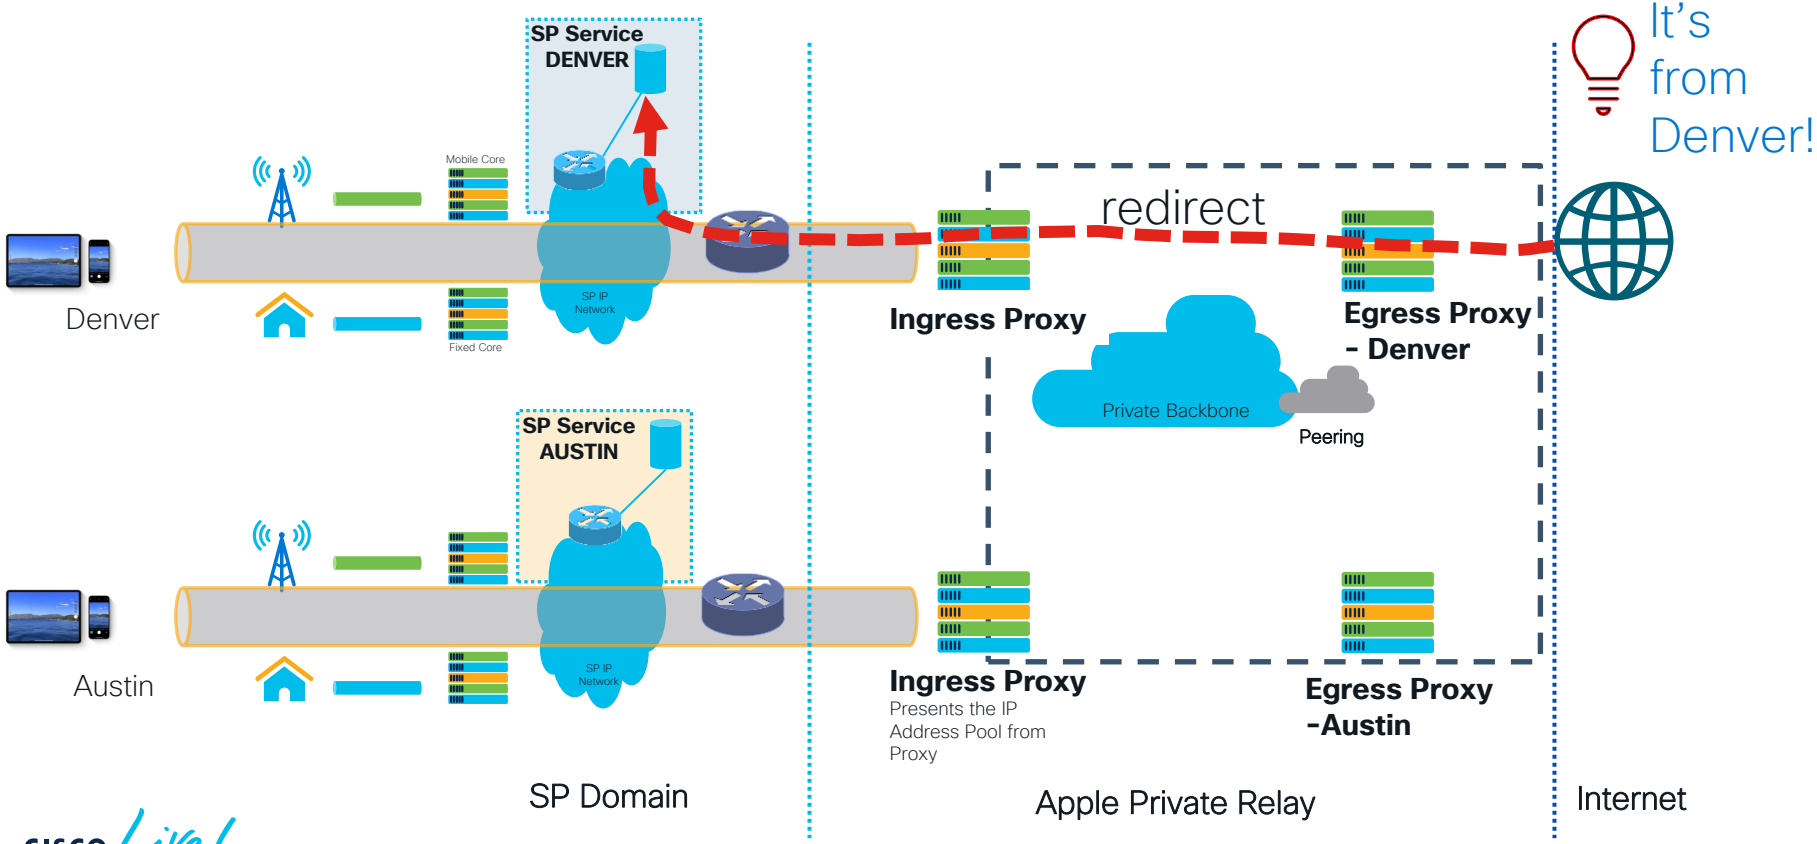

With Apple Private Relay

SP Edge services blown away

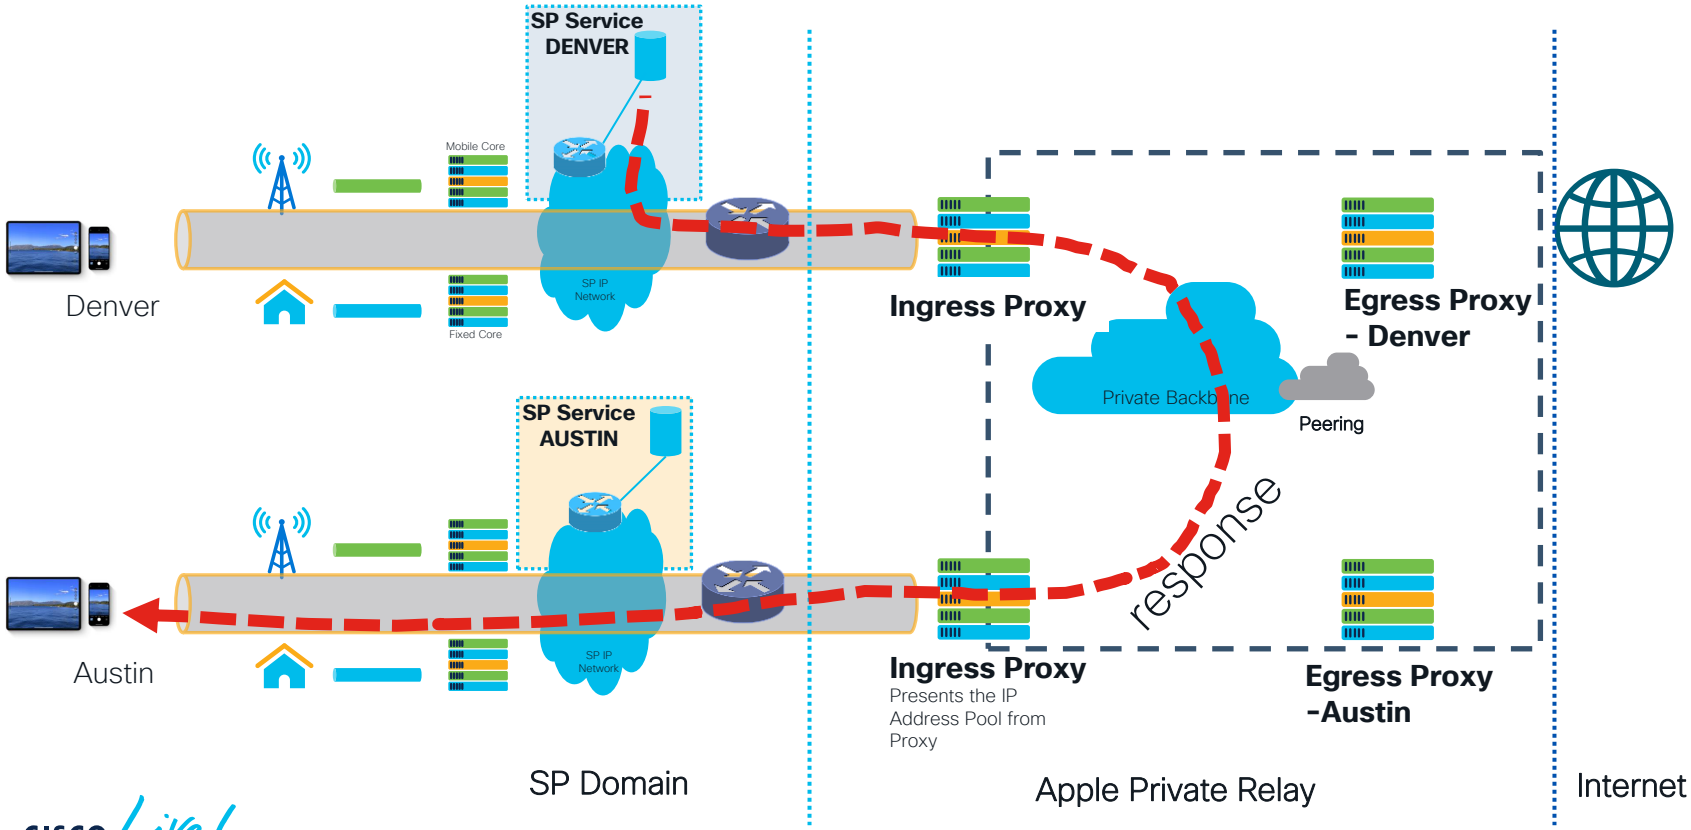

Where this might lead to …
(speculative view, not reality today - yet)

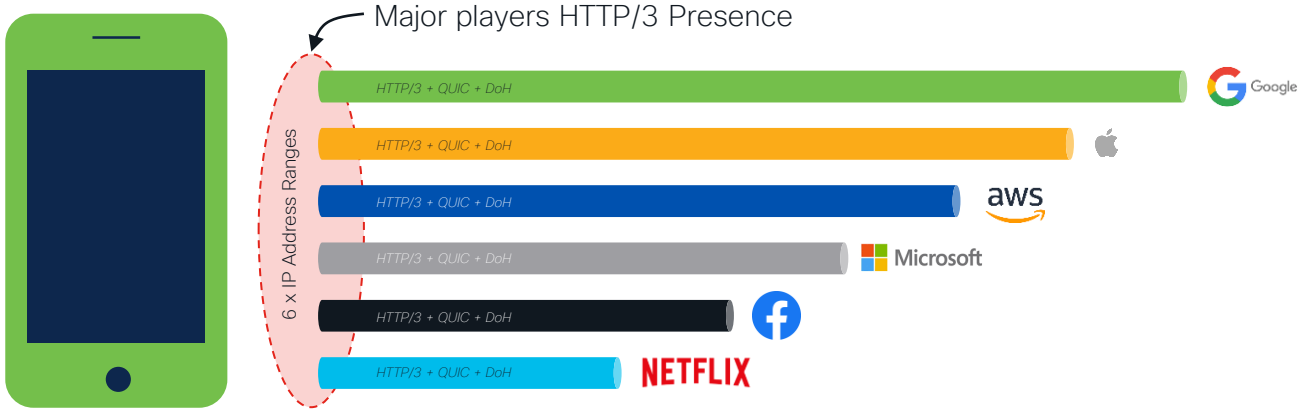

Plus a very, very long tail ….

Observe these trends
- Popularity in App stores
- Browser Preferences
- Cloud architectures
- Client-side OS

SP Services Portfolio needs assessment
(non-exhaustive list)

Differentiated Billing

- *Zero rated Apps*
- *App aware service*

Regulated Services

- *Site blocking*
- *Traffic intercept*

Traffic Management

- *Peering*
- *Optimal interconnect*

non-exhaustive list

Areas of focus

1

Client – Server Architecture

3

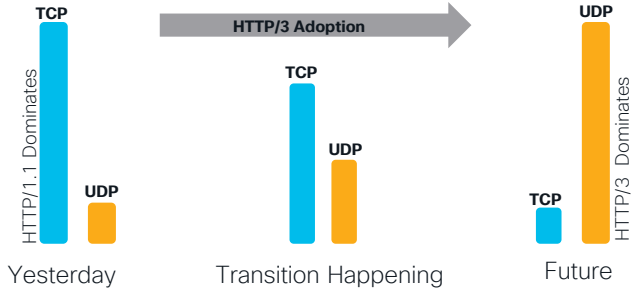

Rise of QUIC

2

Parallel Internet

1 2 3

Lead to:

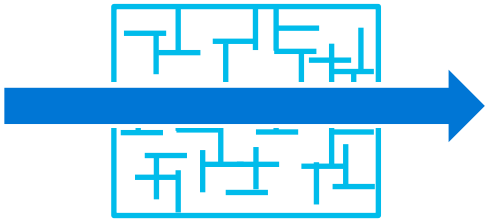

Mass Simplification

Dealing with the new reality: Toolbox & Use Cases

cisco Live!

Customers are looking for solutions
Example Use Cases Asked

Manage video downloads vs video streaming, downloads being the priority

DPI won't work anymore in QUIC

Recognise type of flow and act accordingly

Manage Snap video vs Snap apps

Same problem

Account for encrypted traffic in terms of source/destination

More generically: Identify and manage QUIC flows; mitigate impact on Radio; optimise against industry metrics; future-proof network smarts

App (e.g. Video) Behavior varies by protocol and use case

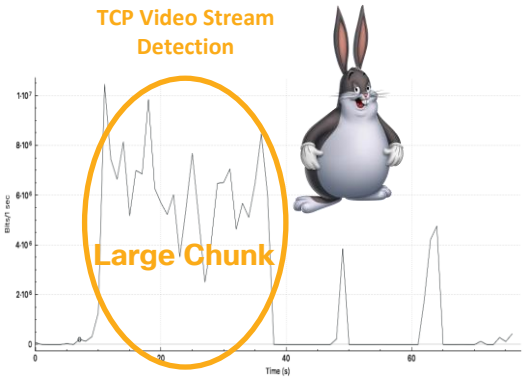

TCP based ABR video players prefer **larger, sustained downloads** due to high cost of establishing the TCP session and reducing time spent in TCP slow start.
Often use HTTP/2 connection. (DASH/HLS) to fix HOL.

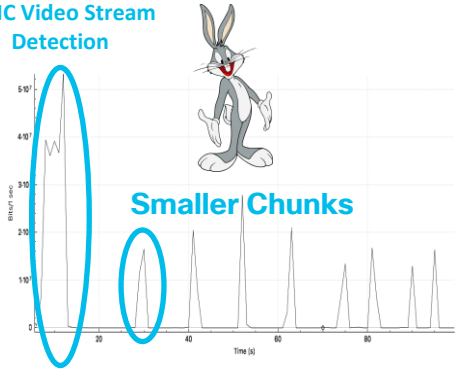

QUIC based ABR video players prefer requesting **video in smaller chunks**.

Multiple QUIC Streams in many cases to (different) servers

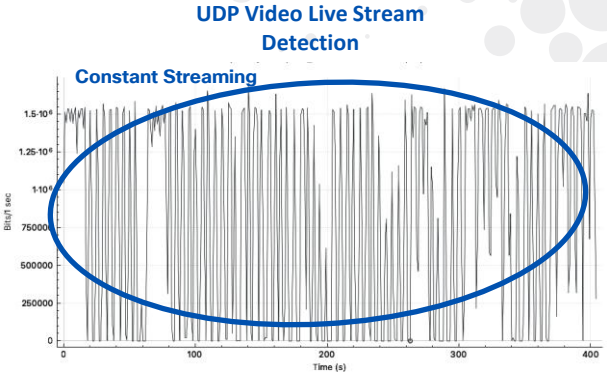

UDP based video players are extremely reliant on consistent network performance. Small buffer, sustained T'put
Applications: YouTube Live, WebEx, Microsoft Teams, Zoom

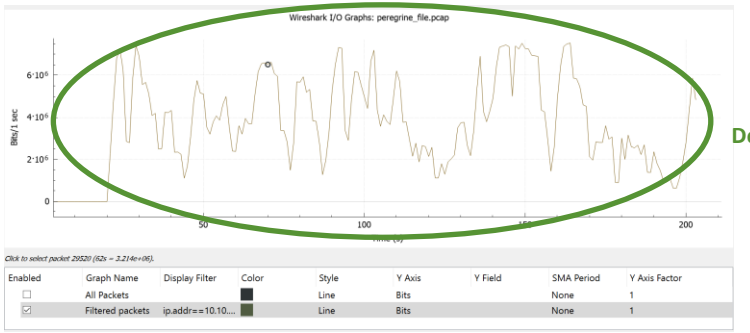

Download Stream Detection

Premium

NETFLIX

YouTube TV

webex by CISCO

zoom

Time Domain Flow recognition

- Observe all flows
- Profile per flow (Time domain matched)
- The resulting profile will allow to distinguish the nature of the flow
  - Content Download
  - (x-Form) Streaming content
  - Real time 2 way communication
  - Video/non-video
  - Short lived flows

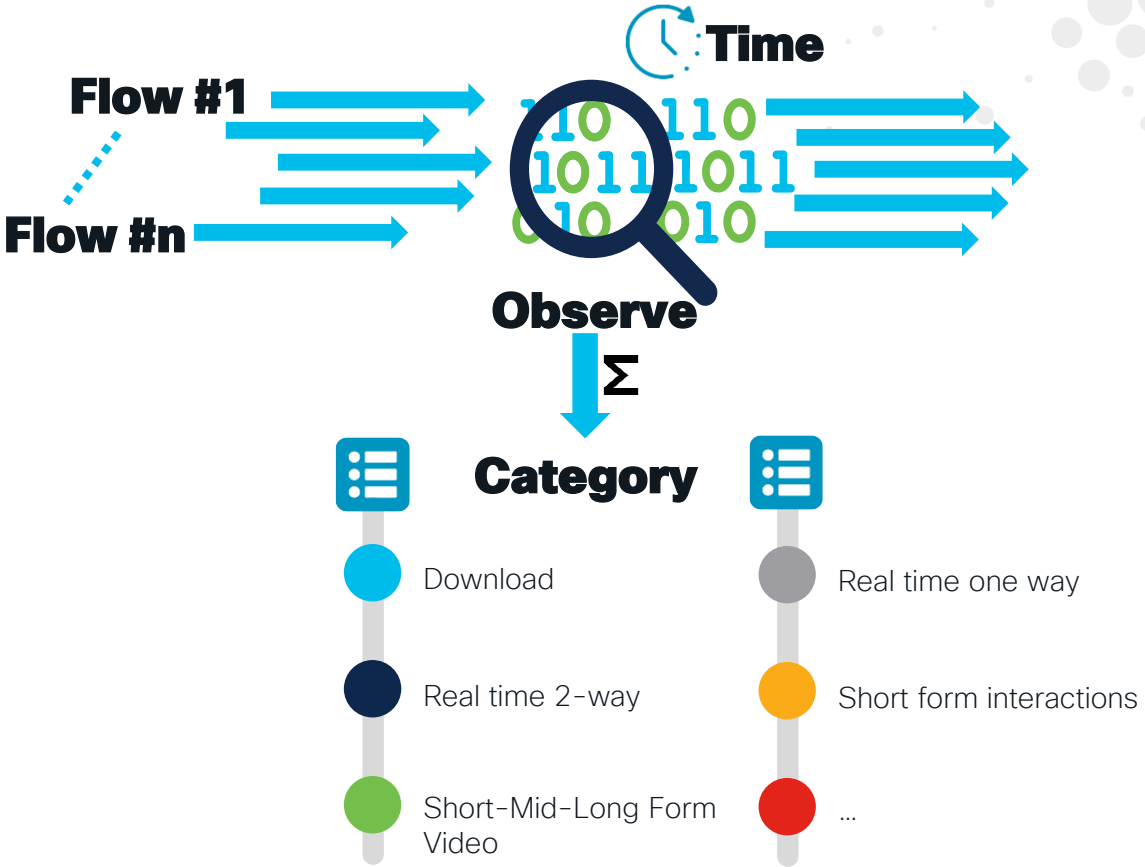

Inferring congestion

- Different congestion algo's have different behaviour
- Time-domain observation + anomaly detection -> congestion inference

Reno vs CUBIC vs BBR behaviour*

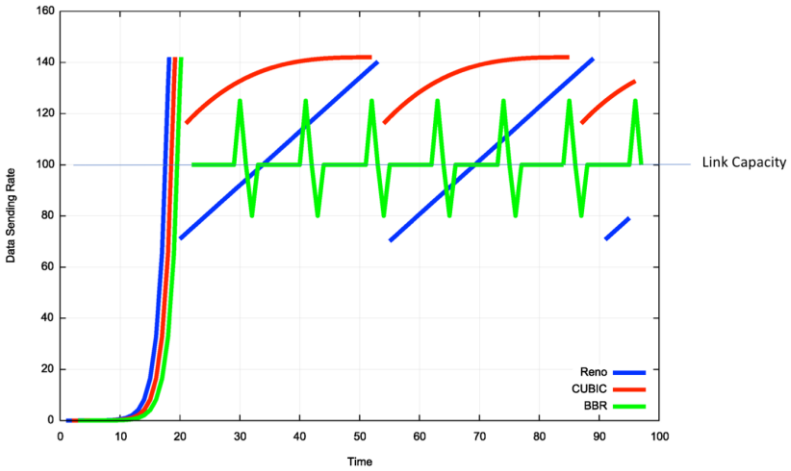

- Assessment of various flows in parallel
- Understand Protocol behaviour: congested or not
- This serves as input for Policy Application

* https://blog.apnic.net/2017/05/09/bbr-new-kid-tcp-block/

Programmable Traffic Management

- Traffic can be controlled in various ways.
  - Buffer
  - Discard
  - Flow control
  - …
- e.g. CUTO is a pre-compiled example where the parameters are implicitly configured

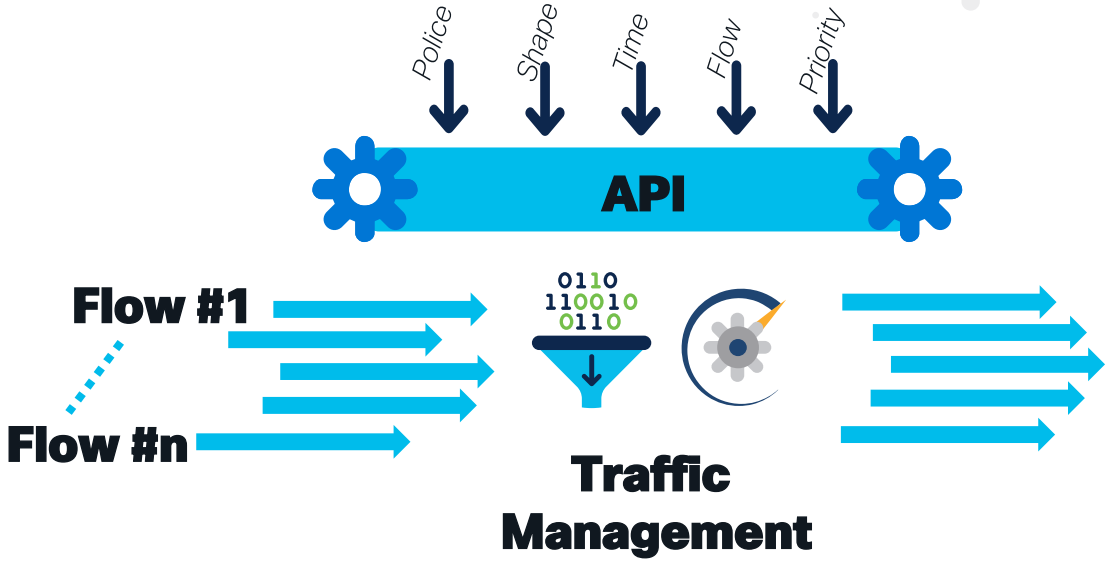

Overall System Logic

Basis for building use cases

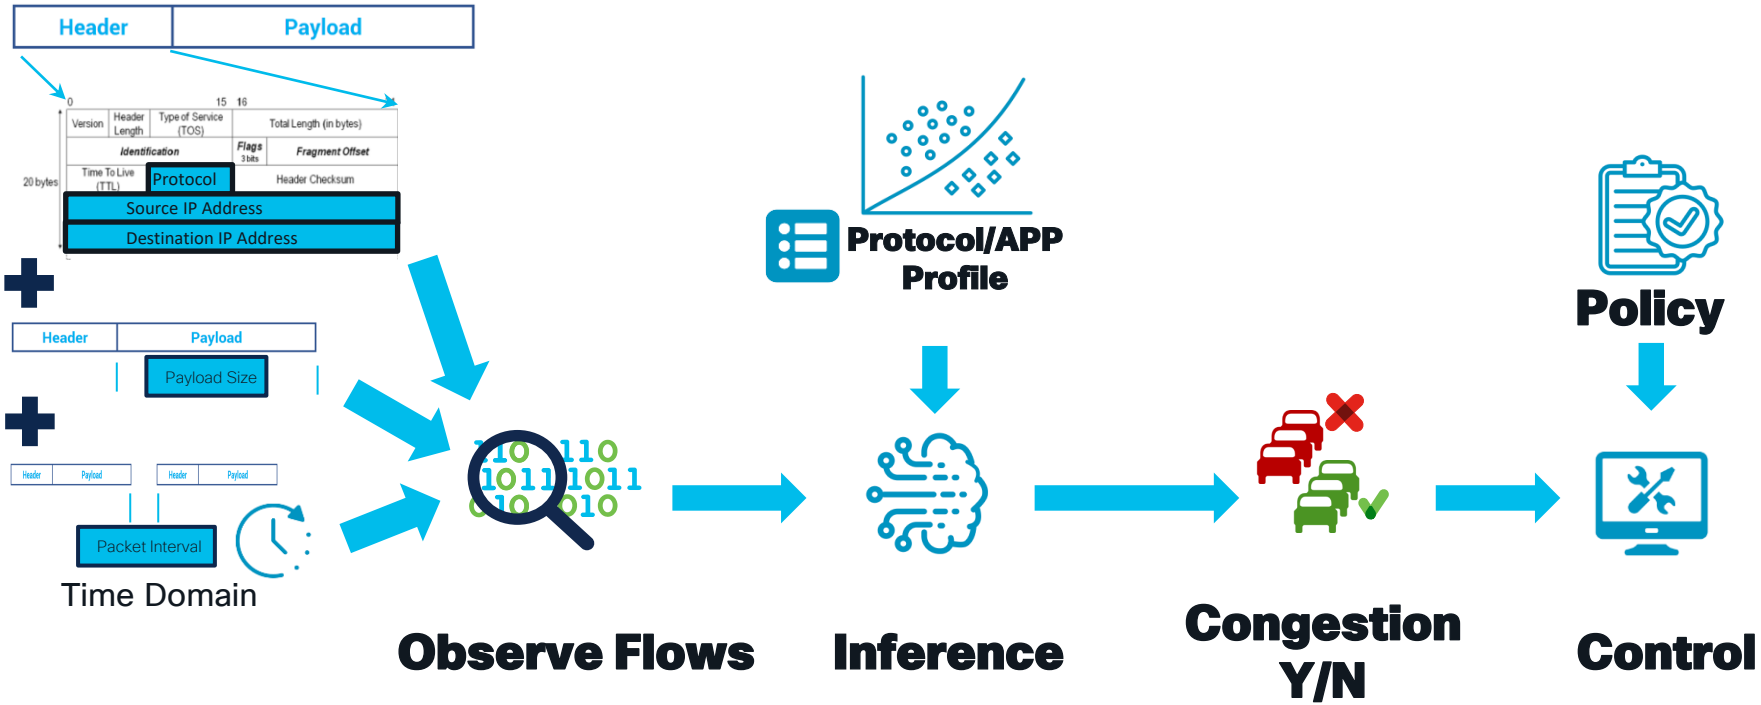

Use Case : Monitoring and analytics

Network Traffic by Volume and Flows

Overall Volume by Apps

Big 5 is 48% of traffic
QUIC is 40% of traffic
"other traffic" still largely TCP, QUIC now visible (4.3%).

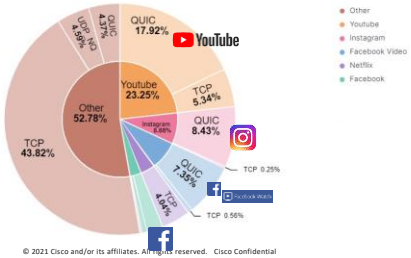

Total Flows by Apps

Lots of TCP sessions (likely IOT related, transactional related)
Big 5 QUIC sessions are very targetted and high efficiency
(video related behaviour)

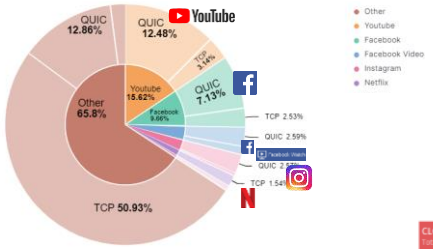

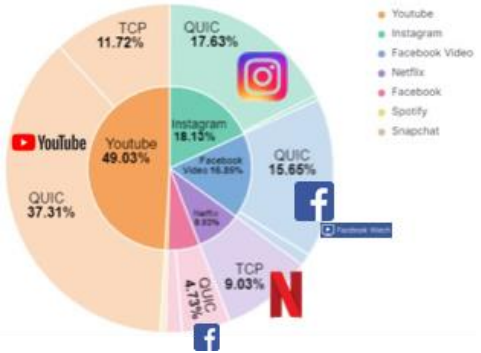

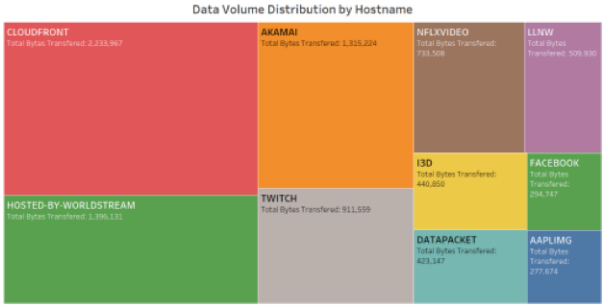

- Monitor all flows
- Infer information for Source (DNS, SNI/eSNI), CDN (ECH), Flow Type (Time domain behaviour)
- ELK (elastic Search, Logstash, Kibana) analytics engine
- Extensible to enriched CDR production

CDN

Hosting

Gaming

Video Streaming

Profile aligned with Fixed Broadband traffic (browser driven traffic)

QUIC : 41%

TCP: 53%

UDP (other): 6%

CUTO
User Experience optimisation under congestion

Congestion inference determines which links are congested and which flows are impacted
Elephant Flow Detection identifies which (QUIC or not) Flows can be managed.
Then Machine Learning determines if that Flow is being delivered during congestion (red circle) and require Flow Control or not (blue circle)

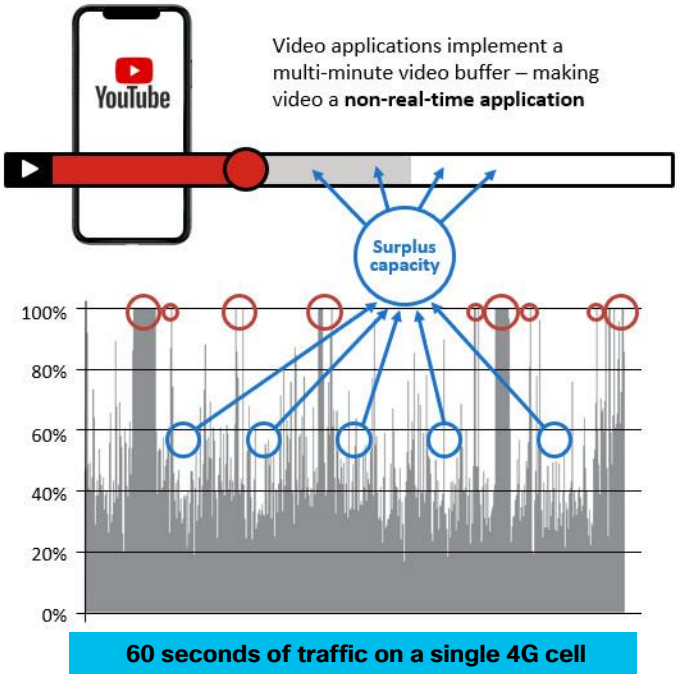

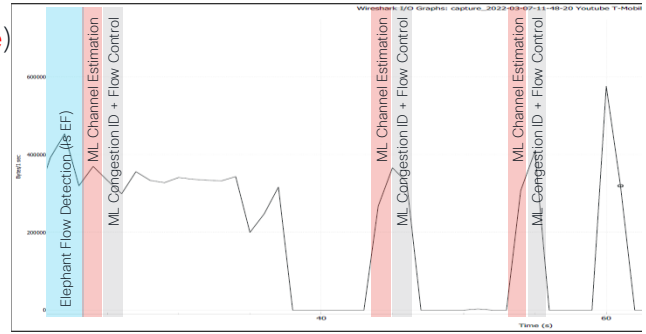

Real World outcome: Tier-1 EU Operator – CUTO use case

Date	Congestion - Carriers					
	Congested Carriers Count			Congested Hours Count		
	EFO Off	EFO On	Percent Change	EFO Off	EFO On	Percent Change
1/17/2022	17	9	-47%	40	16	-60%
1/18/2022	21	10	-52%	57	26	-54%
1/19/2022	27	11	-59%	74	33	-55%
1/20/2022	22	15	-32%	72	46	-36%
1/21/2022	18	11	-39%	68	30	-56%
1/22/2022	23	11	-52%	70	36	-49%
1/23/2022	28	16	-43%	110	57	-48%
Average	**22**	**12**	**-47%**	**70**	**35**	**-50%**

-47%
Frequency Congestion

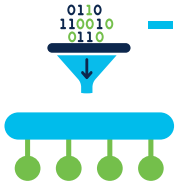

-50%
Network Congestion

90% **Speed Test Improvement**

OOKLA SPEEDTEST

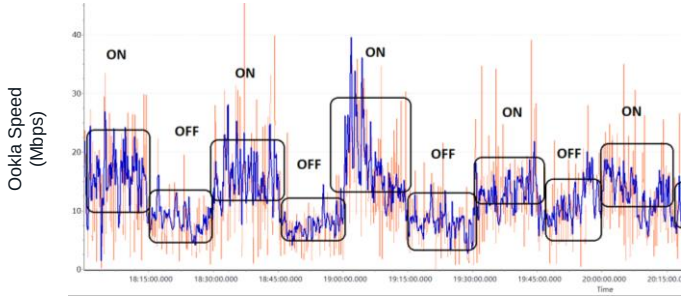

CUTO* Wireline
Enhanced User Experience within SLA Boundaries

Situation

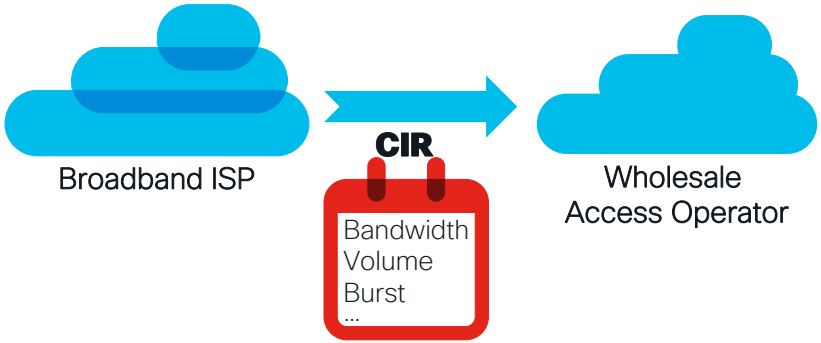

Conform to SLA results in predictable cost ✓

Violate SLA results in additional cost ✗

Indiscriminate Policing leads to bad user experience

Solution

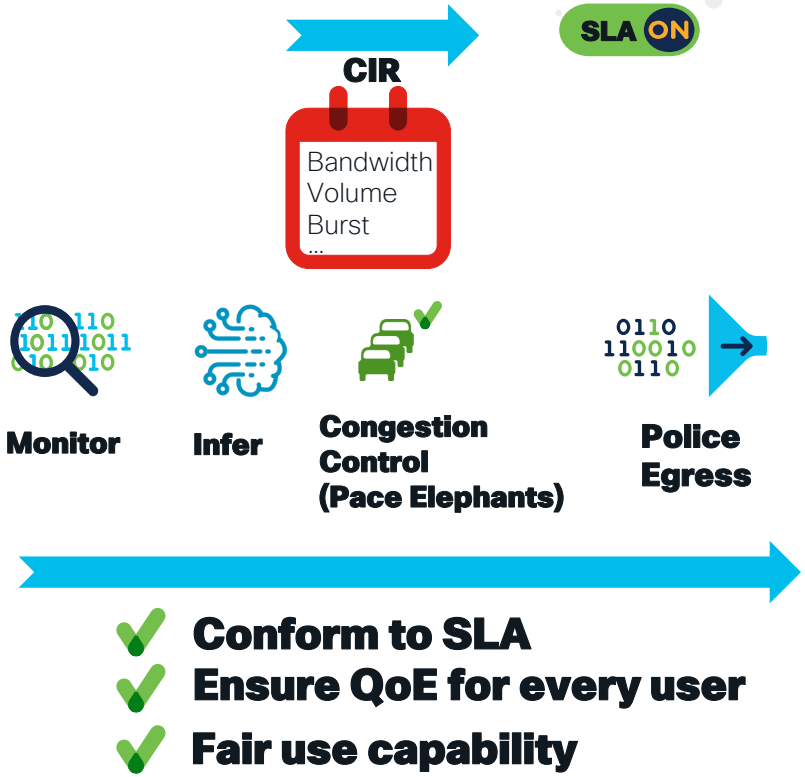

Conform to SLA

Ensure QoE for every user

Fair use capability

*Cisco Ultra Traffic Optimization

Custom Policy Enforcement

e.g. Differentiate between ”download” and “streaming” (within same app)

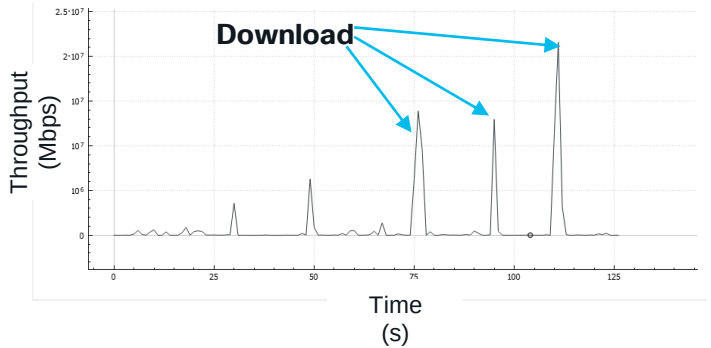

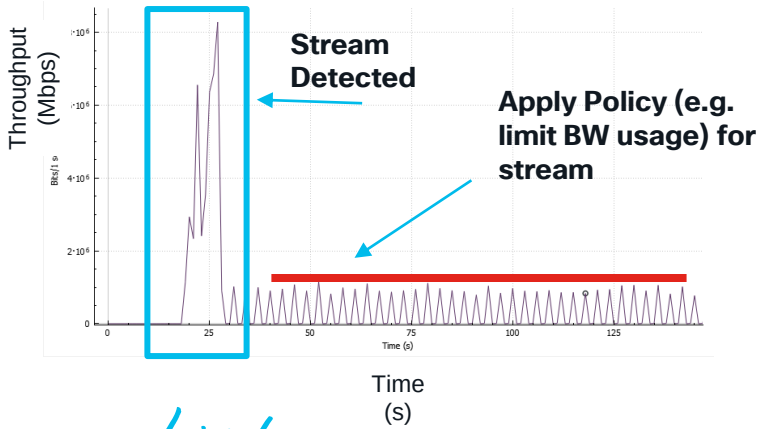

- Same Source/Destination Address
- Differentiate between download versus streaming *on the same SA/DA*
- **Apply Policy per flow type, e.g.**
  - **Download Policy: no action**
  - **Streaming Policy: Limit to set BW profile (police/buffer/…)**

Use Case: Dynamic TCP Optimization
Contextualized TCP Management

Policy Enforcement Engine

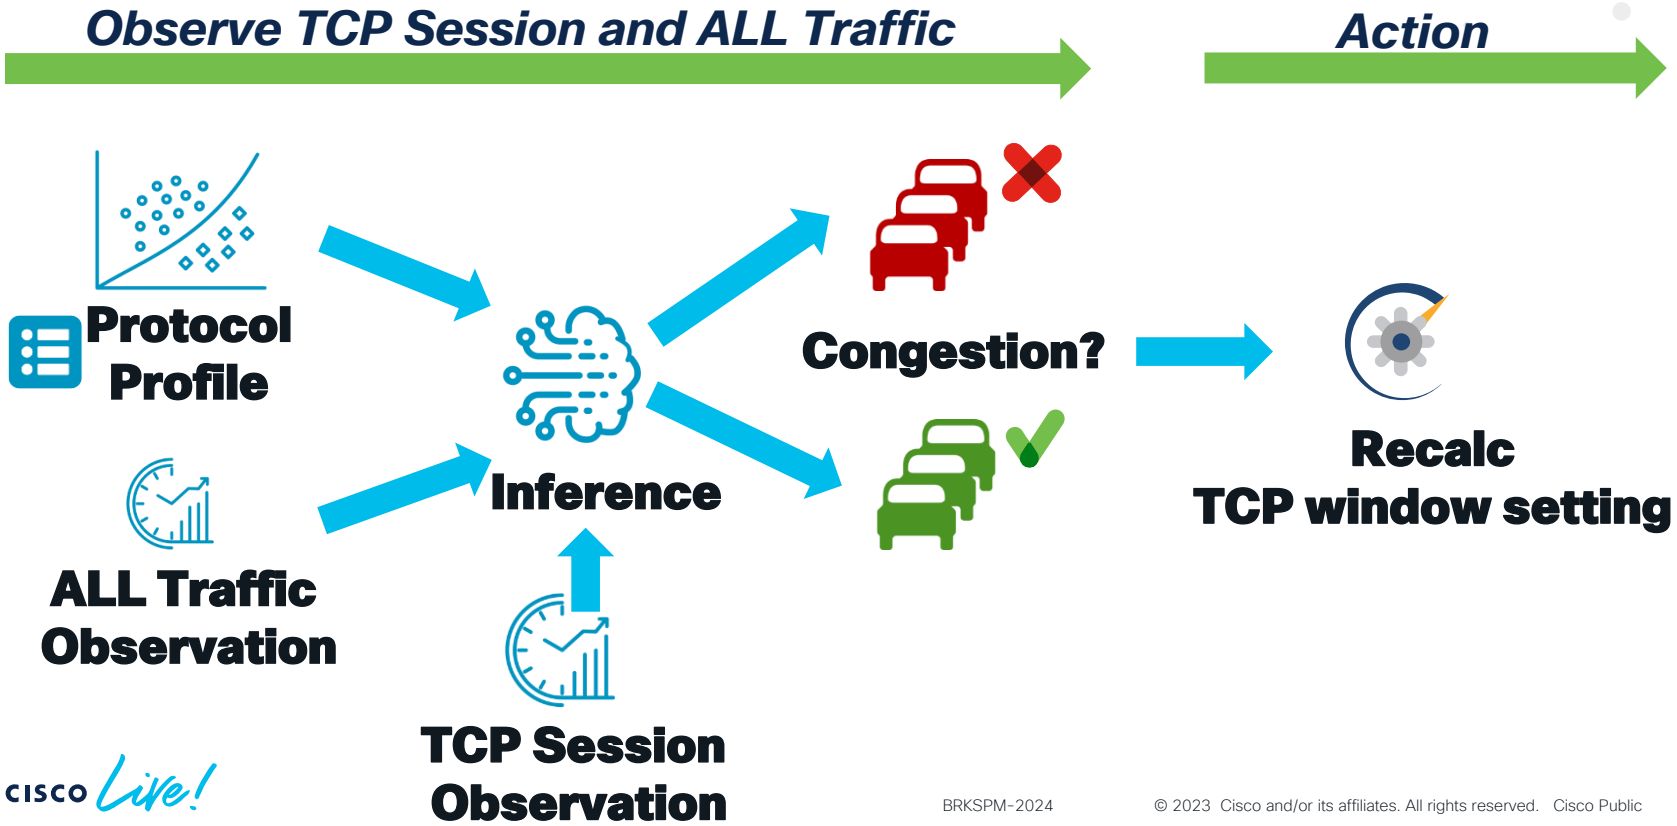

Use Case : Protecting Real-time Traffic

Observe traffic, detect videoconferencing stream, measure steady state Bandwith usage of video conf stream, shape traffic to (total-videoconf BW)

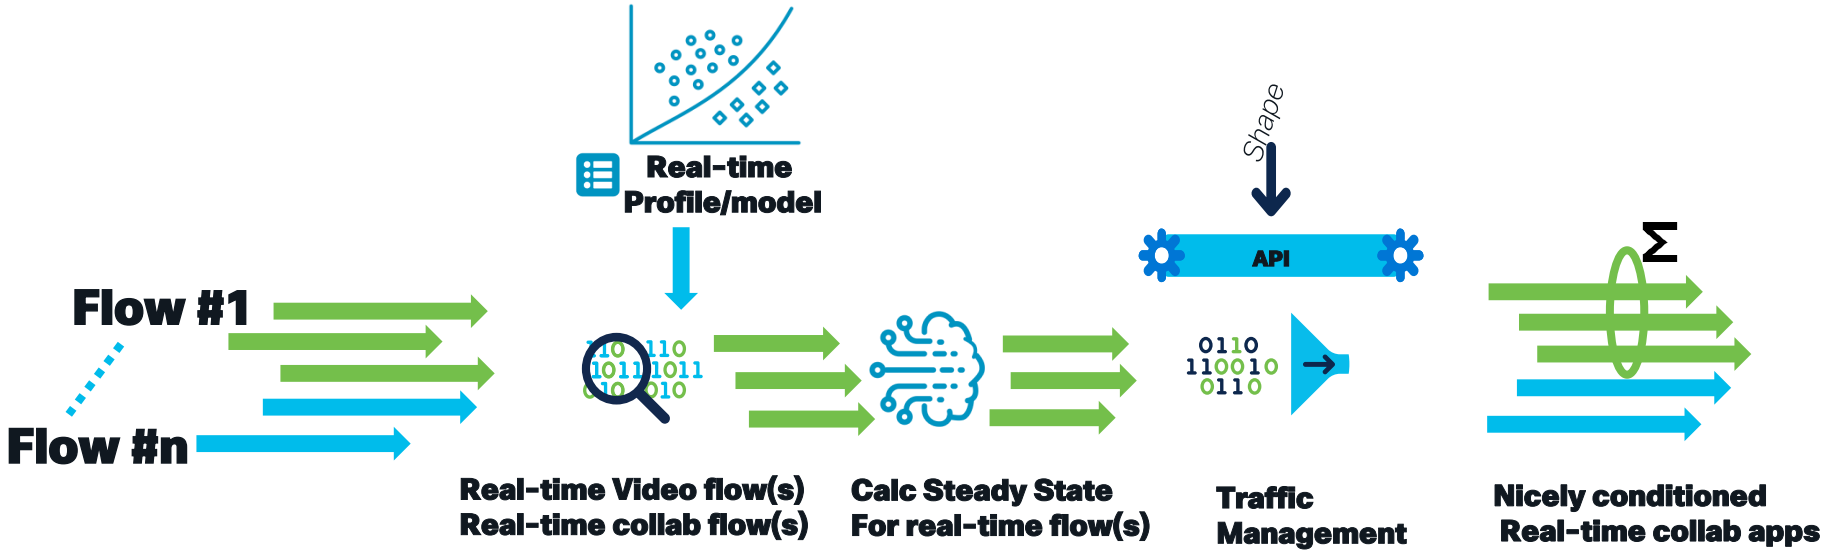

Why does this scale

Simple

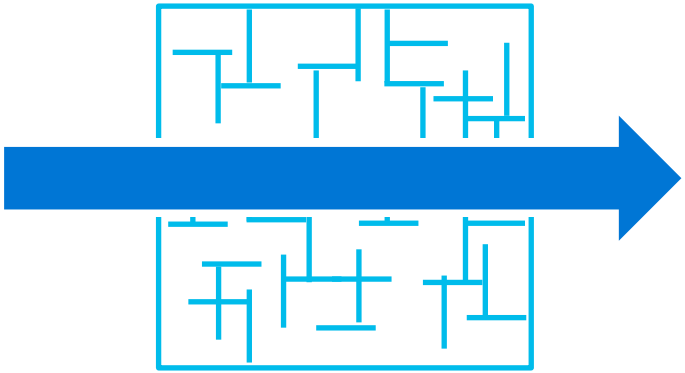

- I only use state on the important/interesting stuff
  - 20% of the flows generate 80% of the volume

Smart

- I only use state if I need it
  - when there is a reason e.g. congestion

Summary

- Traffic is encrypted, application controlled, and obfuscated
- Traditional DPI approaches (w)(d)on’t work
- This evolution will affect Service Provider consumer offering policy
- An IP centric approach is feasible and addresses several use cases

Complete your Session Survey

- Please complete your session survey after each session. Your feedback is important.
- Complete a minimum of 4 session surveys and the Overall Conference survey (open from Thursday) to receive your Cisco Live t-shirt.

- All surveys can be taken in the Cisco Events Mobile App or by logging in to the Session Catalog and clicking the "Attendee Dashboard" at https://www.ciscolive.com/emea/learn/sessions/session-catalog.html

Continue Your Education

Visit the Cisco Showcase for related demos.

Book your one-on-one Meet the Engineer meeting.

Attend any of the related sessions at the DevNet, Capture the Flag, and Walk-in Labs zones.

Visit the On-Demand Library for more sessions at ciscolive.com/on-demand.

The bridge to possible

Thank you

For your time and attention !


CISCO Live!
ALL IN