
cisco Live!
Let's go

The bridge to possible

The New, Encrypted Protocol Stack

and how to deal with it

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

Agenda

- The New Internet
- The New IP protocol stack & New Traffic Behaviour
- What's left?

In memory of
and based on the
brilliant work of
Mark Gallagher

14/09/1966-17/09/2021

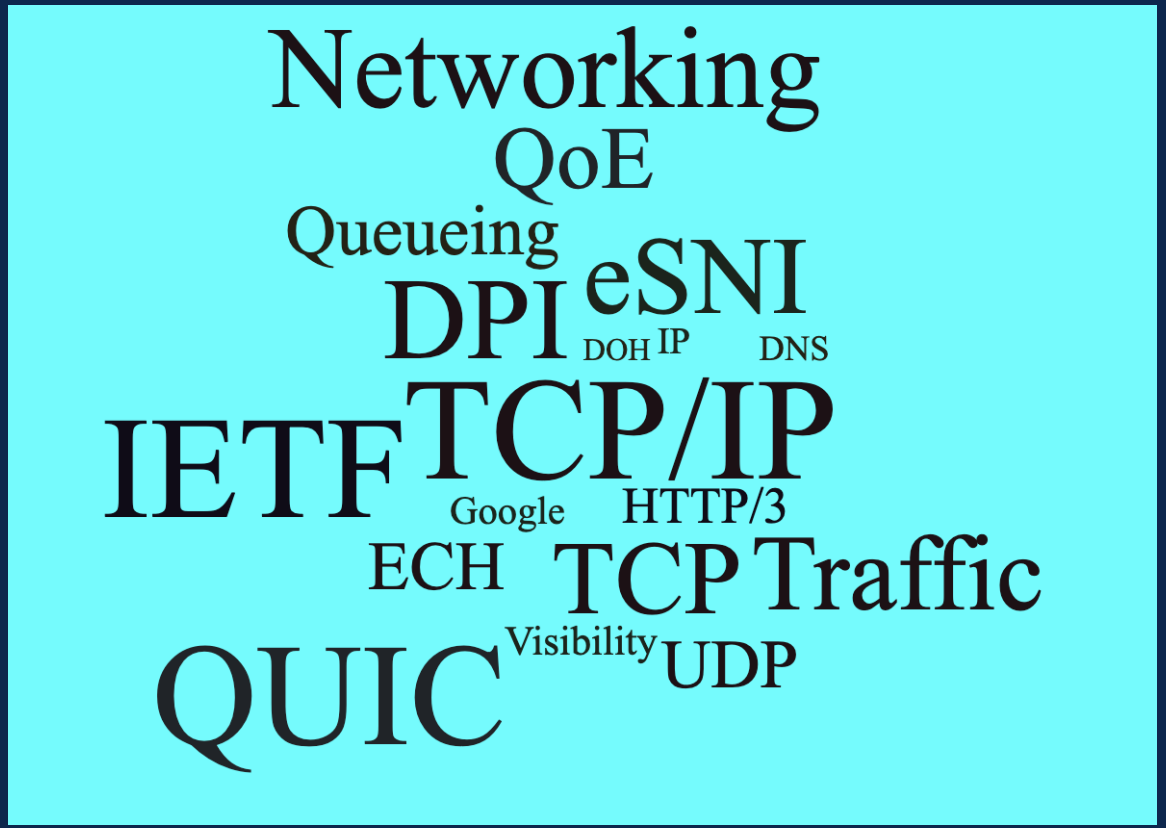

The New Internet

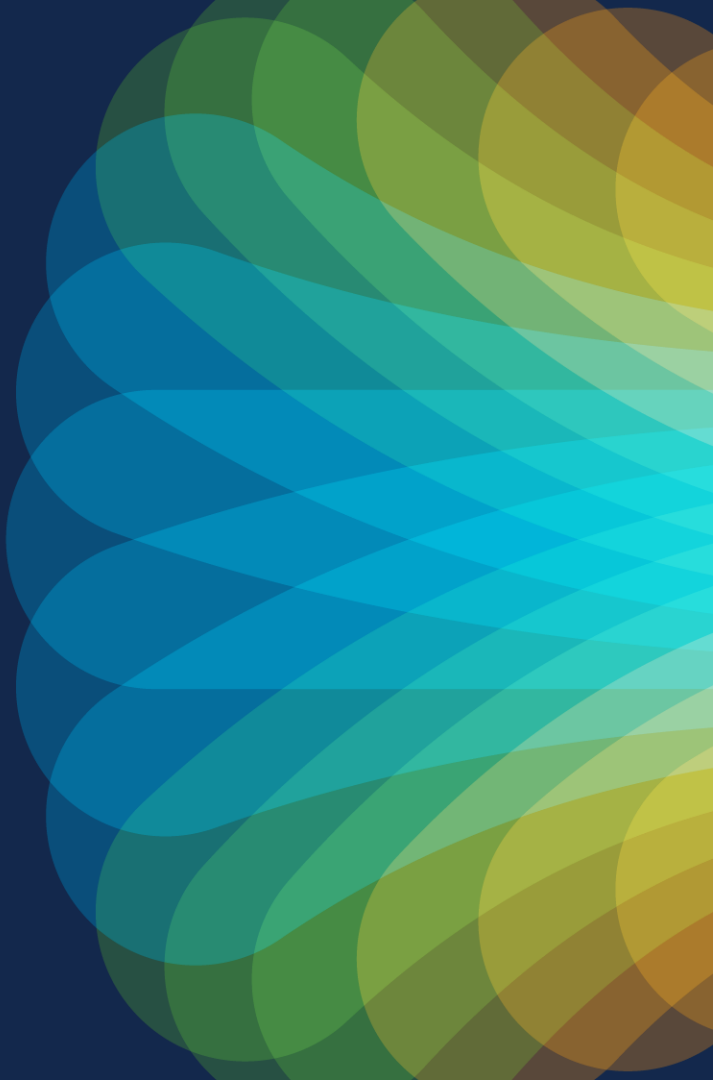

The Internet Reality – circa 2020 – Major US Carrier

>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

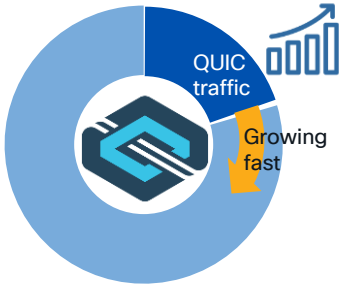

Many small flows
Micro-sessions

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

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

Fast forward 18 months – Tier-1 EU Mobile Carrier

Overall Volume

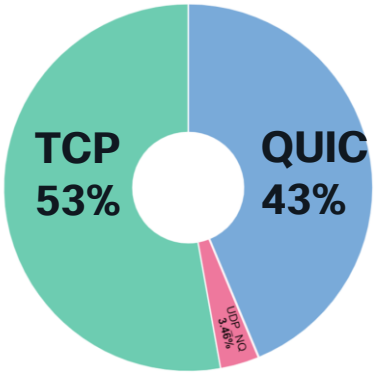

QUIC has doubled in 18 months

QUIC is 43% of total and rising

YouTube

Volume

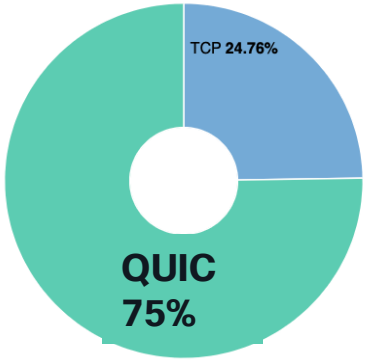

QUIC is “default”

Meta

Volume

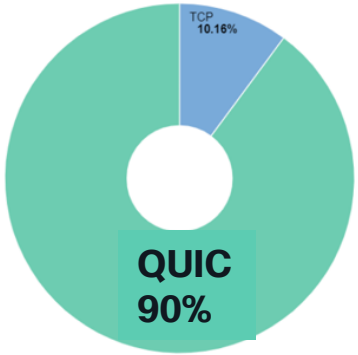

Meta has gone full QUIC

(snapshot 11/2/2022)

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%).

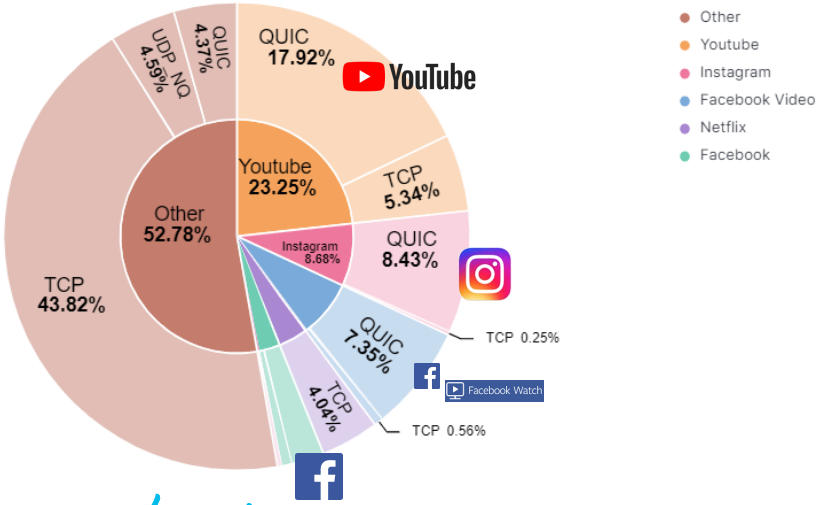

Total Flows by Apps

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

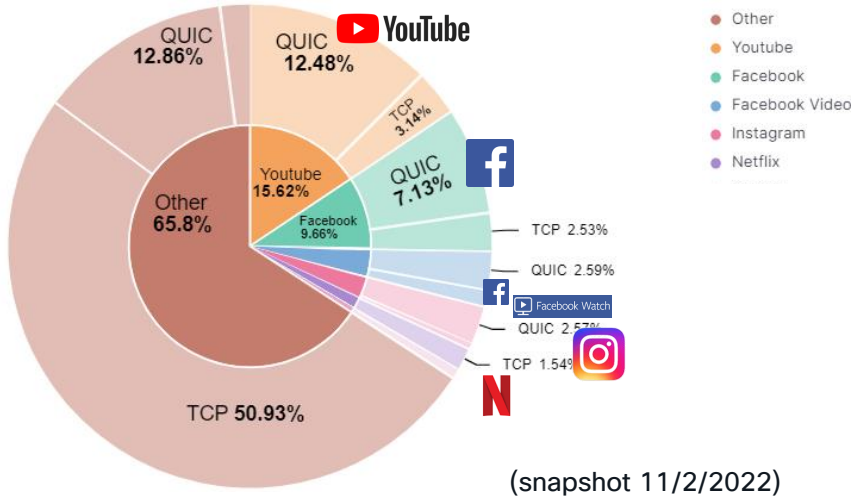

(snapshot 11/2/2022)

The pattern persists worldwide into 2023

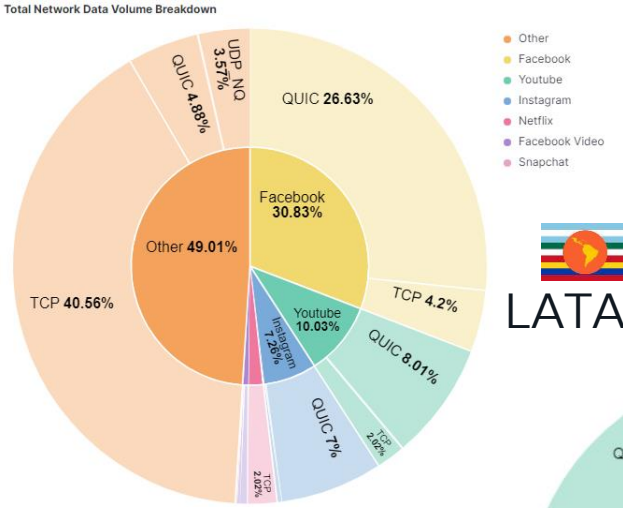

LATAM

QUIC: 46.52%

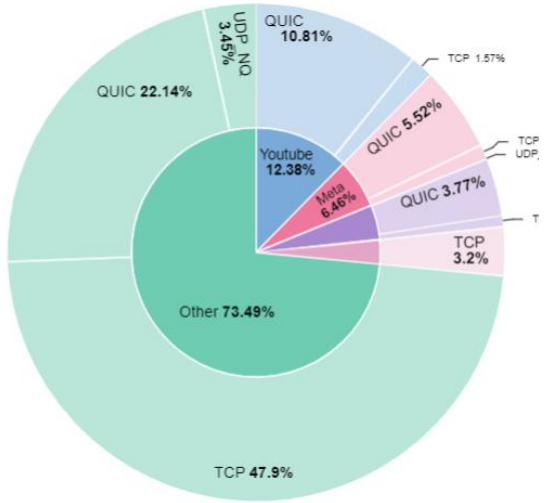

EU

QUIC: 42.24%

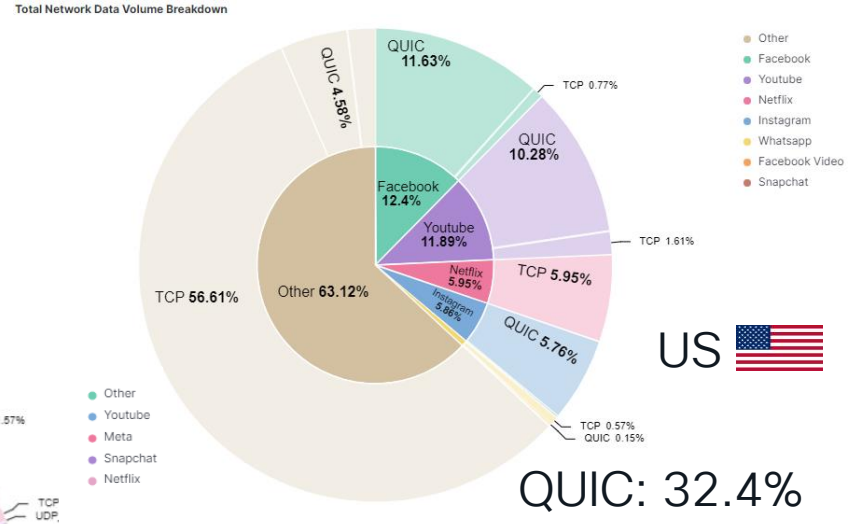

US

QUIC: 32.4%

Early 2024 Data: QUIC still going strong

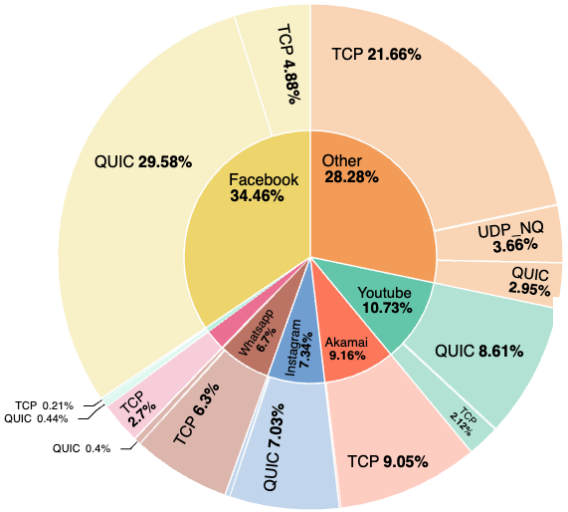

LATAM

QUIC: 47.31%

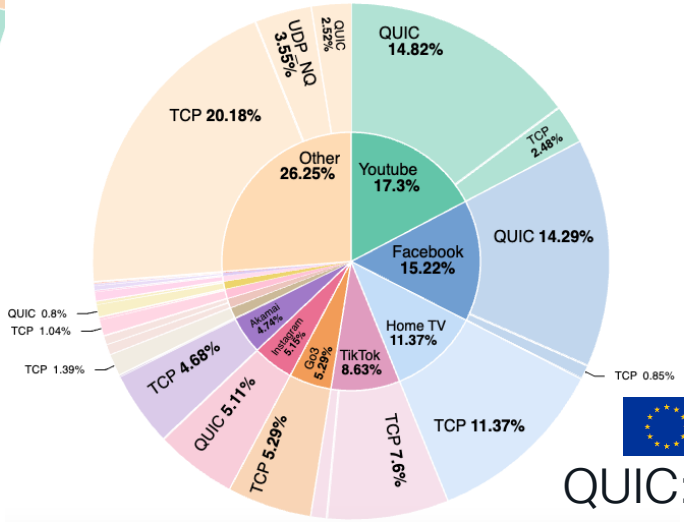

QUIC: 41.98%

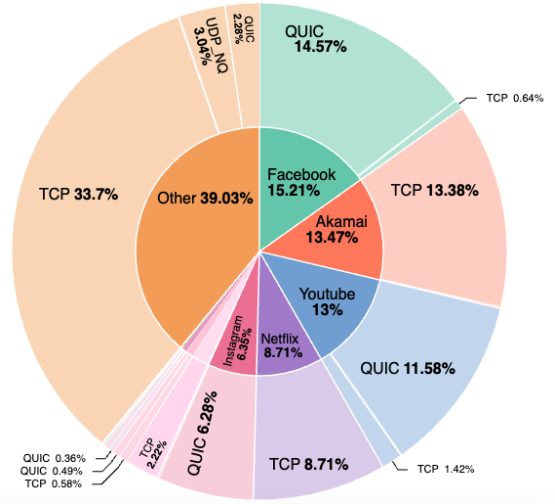

QUIC: 41.5%

2024 APJC New Measurements

QUIC Volume:
57.97%

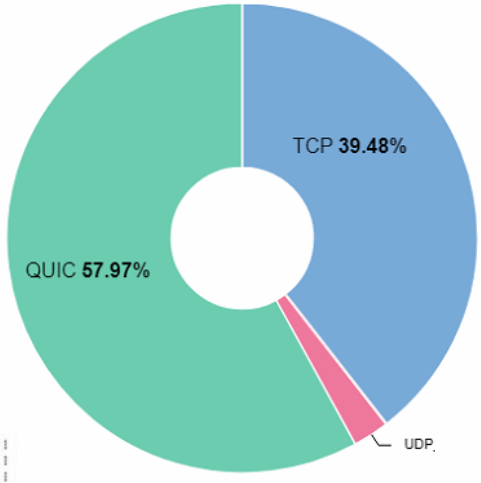

Flows:

QUIC: 46.12%
TCP: 52.7%

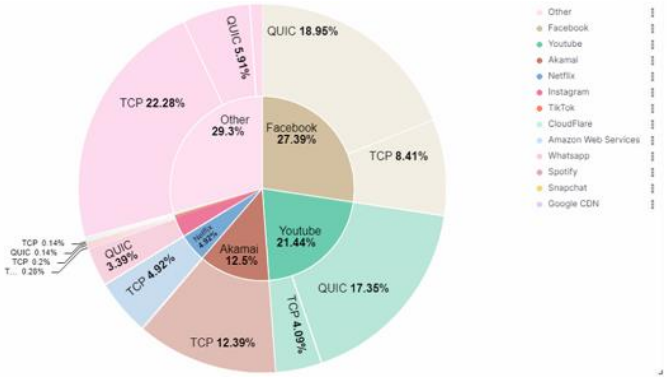

#Flows

37% is Meta
“Meta Addiction”
Volume

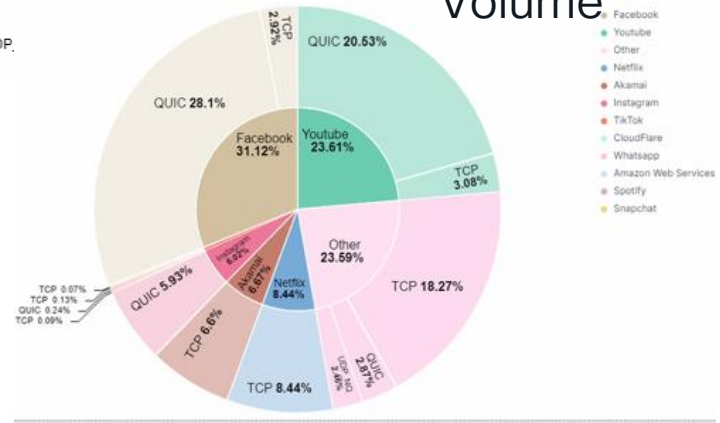

QUIC is growing across the world

various snapshots – Approaching 50% WW

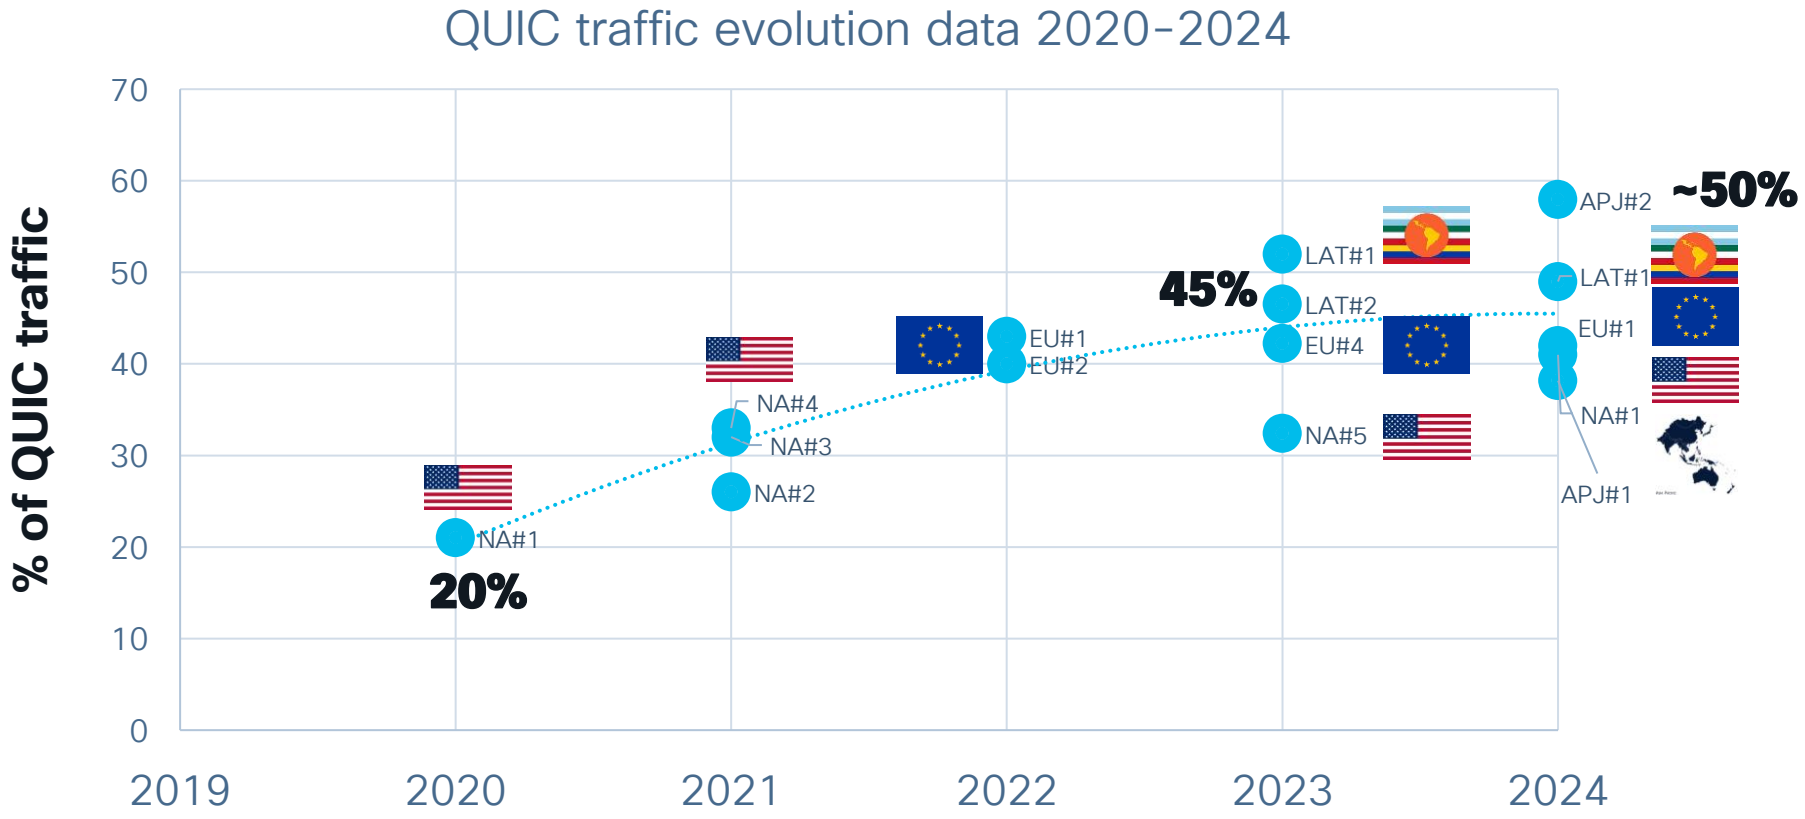

Top 5 Apps – QUIC is dominant
80/20 rule now

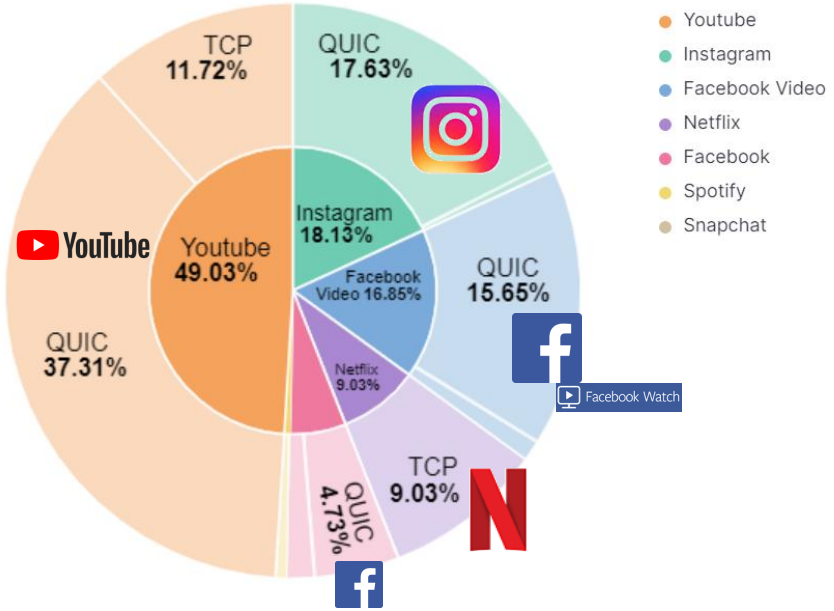

April 10 2022

The New IP stack

New Stack, New Behaviour

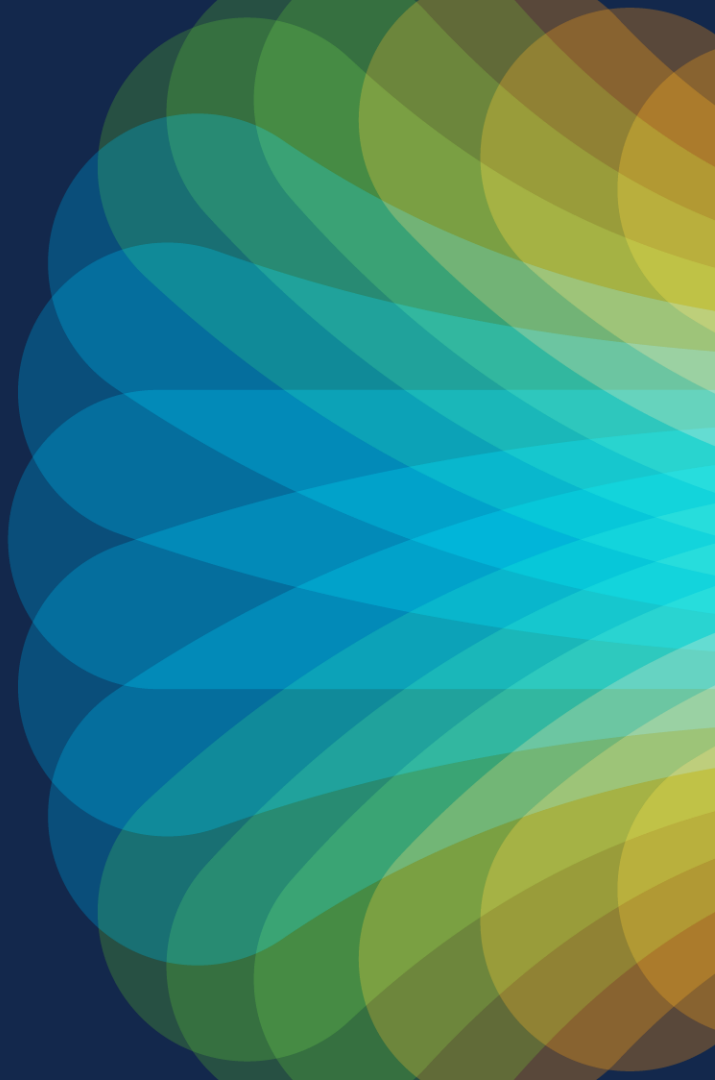

cisco Live!

An application driven global transition

HTTP/3 Stack = UDP+QUIC+TLS

Old App Stack

- HTTP/1.1/2
- TLS
- TCP
- IP

New App Stack

QUIC – RFC 9000
HTTP/3 – RFC9114

- HTTP/3
- QUIC + TLS1.3
- UDP
- IP

DoH

DoT – RFC7858
DoH – RFC8484

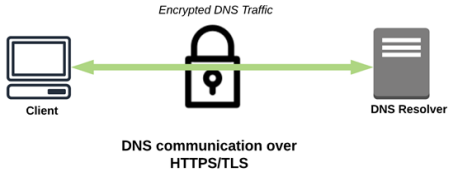

eSNI / ECH

RFC8744

Large Scale Adoption

IETF RFC 8446 – Underpins all others

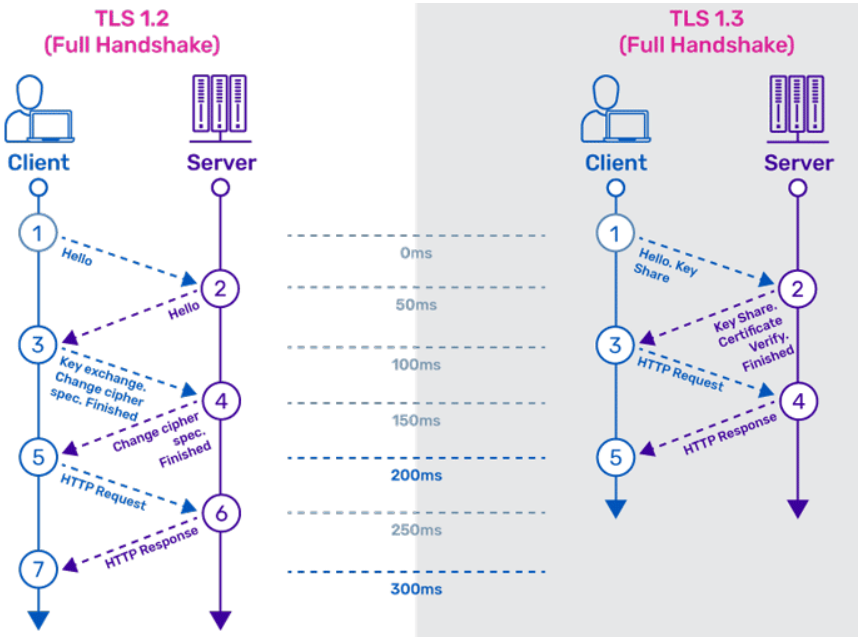

A Faster TLS Handshake

- 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

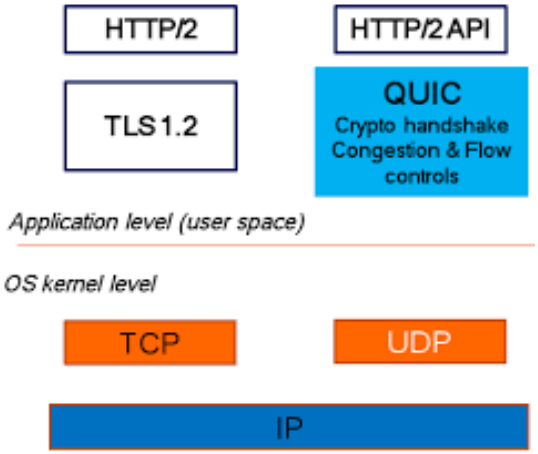

User Space
TLS 1.3 Encrypted

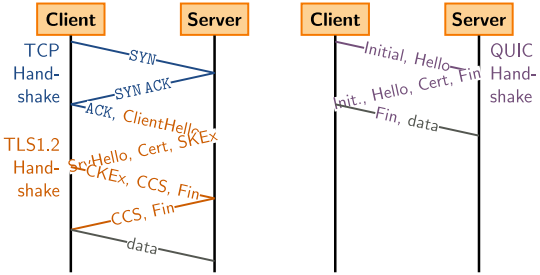

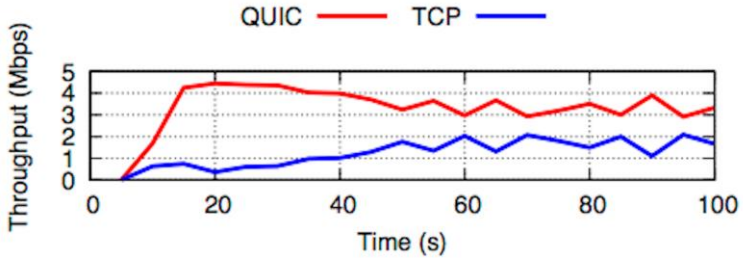

Fast of the blocks

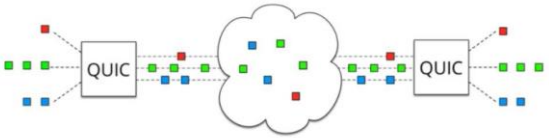

Deliver at all cost
(Multiplex, no-HOL)

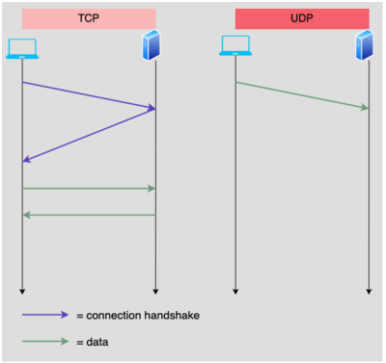

UDP is ”fire and forget”
App controls the rest

Moves Control of the User Experience to the App

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

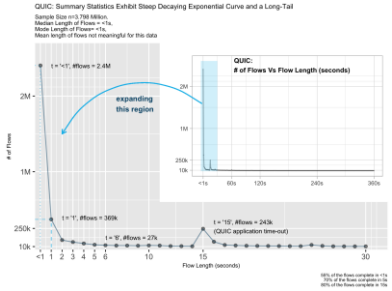

70% of interactions complete in <5s**

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)

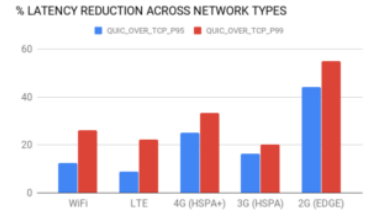

The poorer the network, the better the improvement*

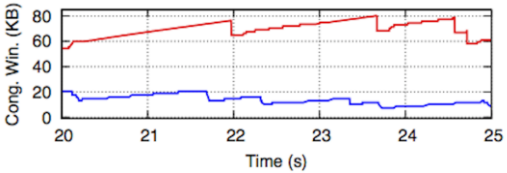

QUIC is "Unfair"***

Impacted Areas

(e.g. wireless access)

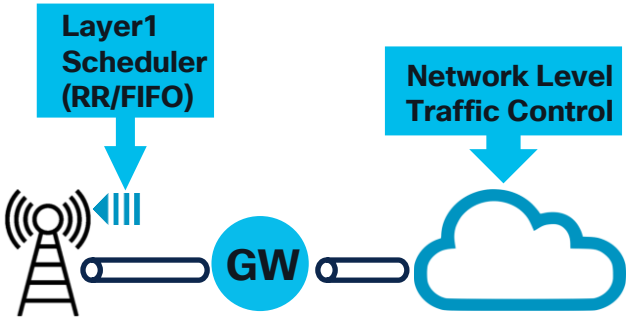

*uber engineering;**Cisco Analysis, cust.data;***APNIC study

Secure DNS – Domain lookup Privacy by default

DoH - RFC8484 is becoming mainstream

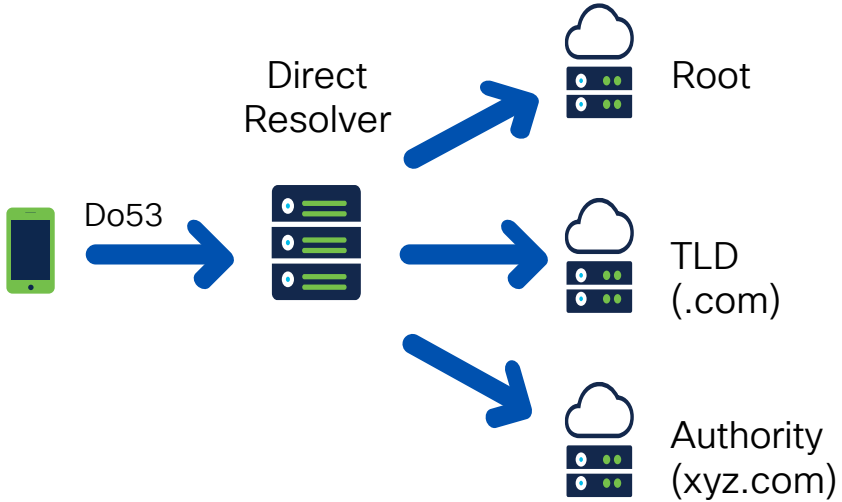

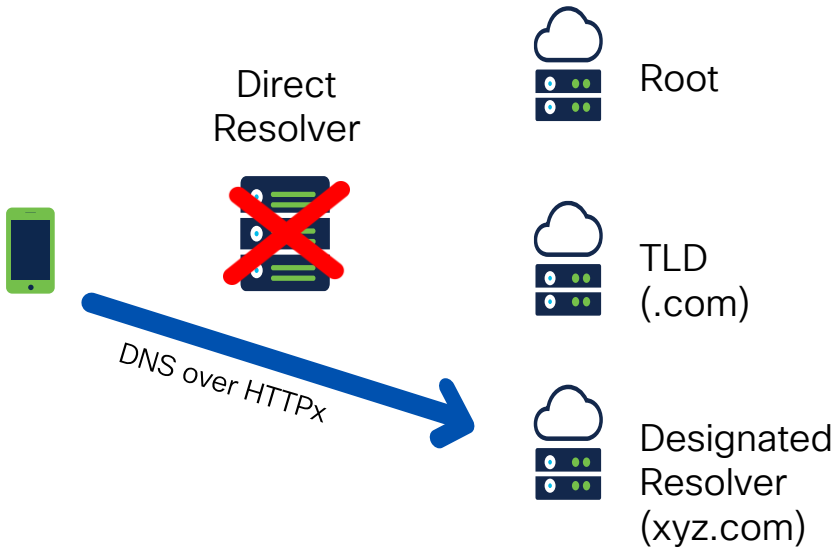

From: DNS Hierarchy + cleartext fields

To: DNS (direct) Connect + ciphered fields

+ DNS is controlled by Applications

Hiding the destination completely - eSNI & ECH

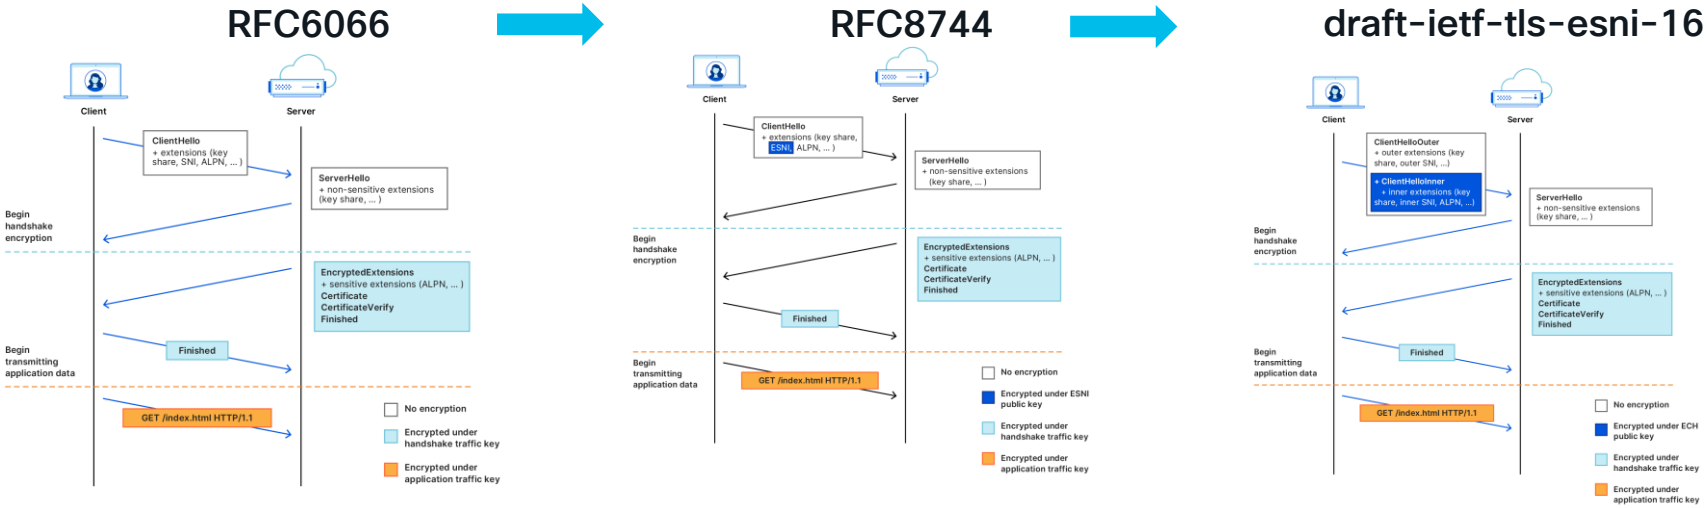

Classic SNI
Destination & Capabilities in the clear

eSNI
Destination leaked via DNS
ALPN still in the clear*

ECH
Only CDN address visible – DoH
SNI & Capabilities fully encrypted

*Application layer Potocol Negotiation

Bringing this all together

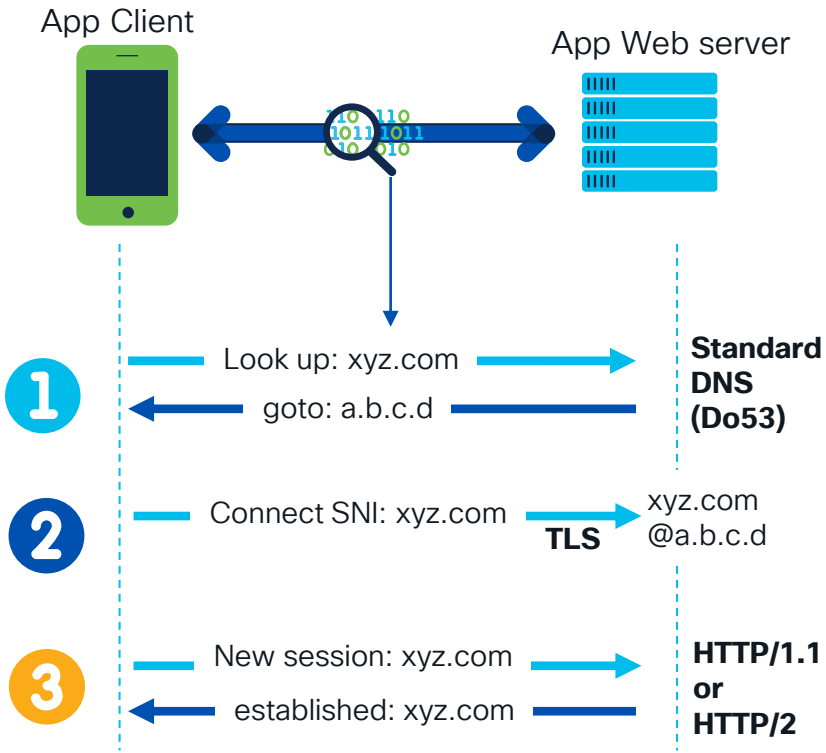

- ✅ Well-understood protocol stack
- ✅ Foundation of **all** web traffic
- ✅ Adopted by Applications
- ✅ Globally scaled

First, encrypt DNS

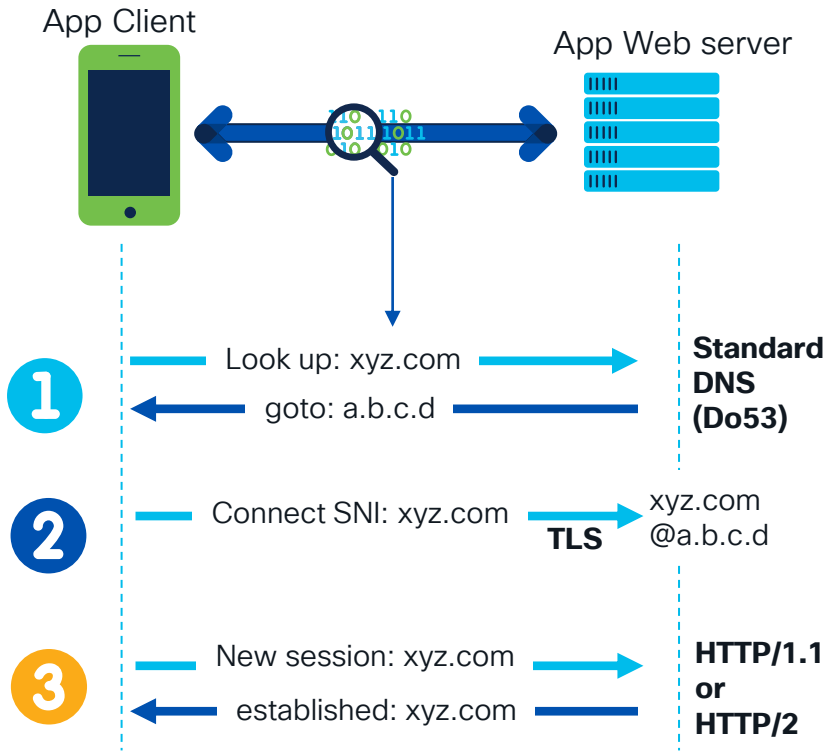

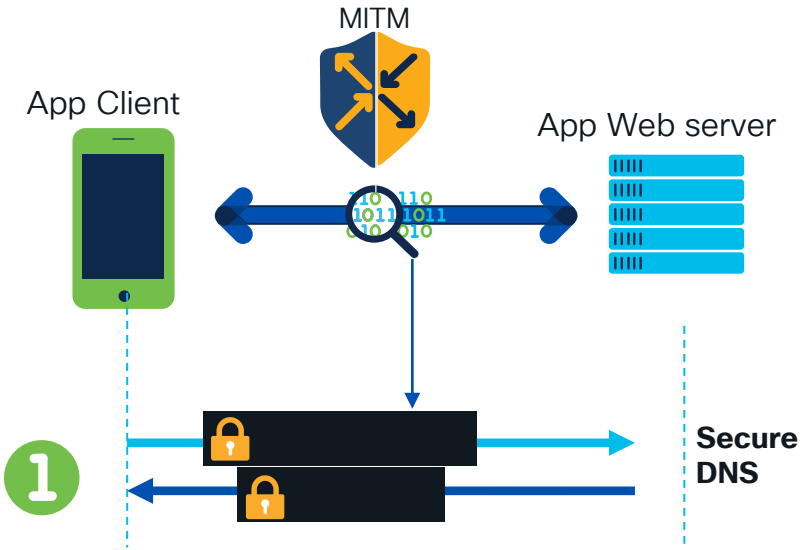

Need TLS and HTTP inspection to recover information

Second, use TLS 1.3 Features

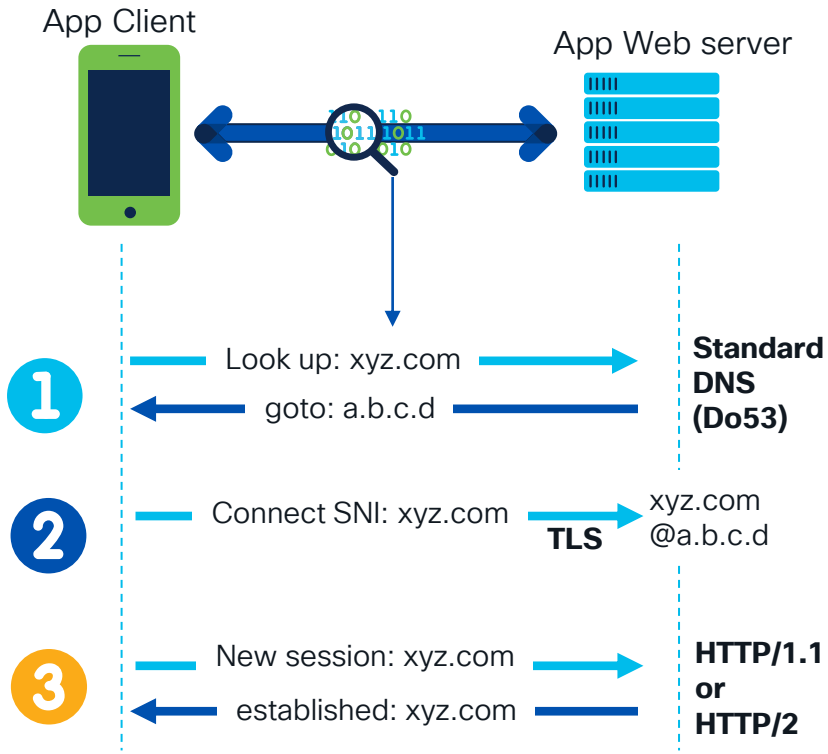

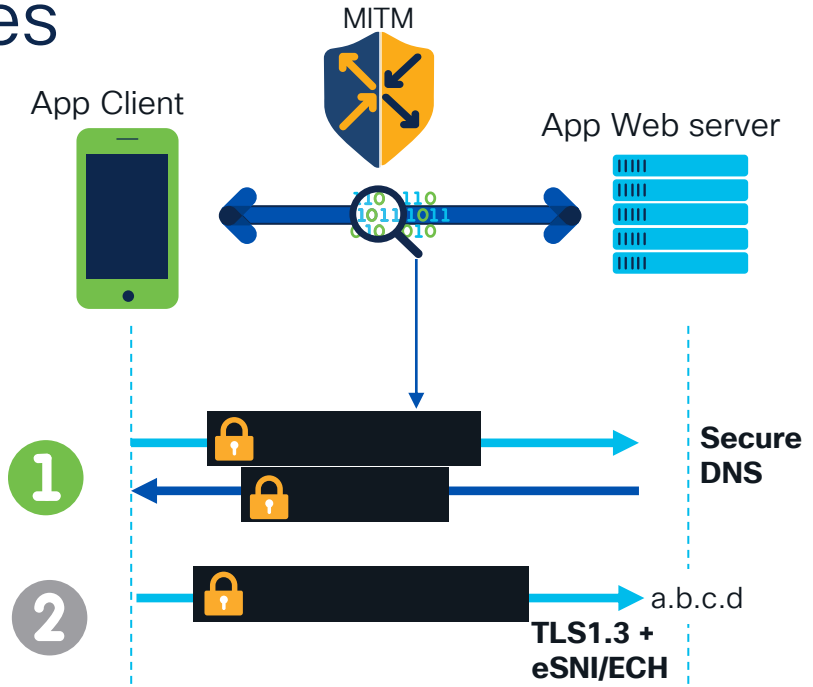

Session setup encrypted but session patterns can provide insights

Finally, HTTP/3 and QUIC multi-session technology

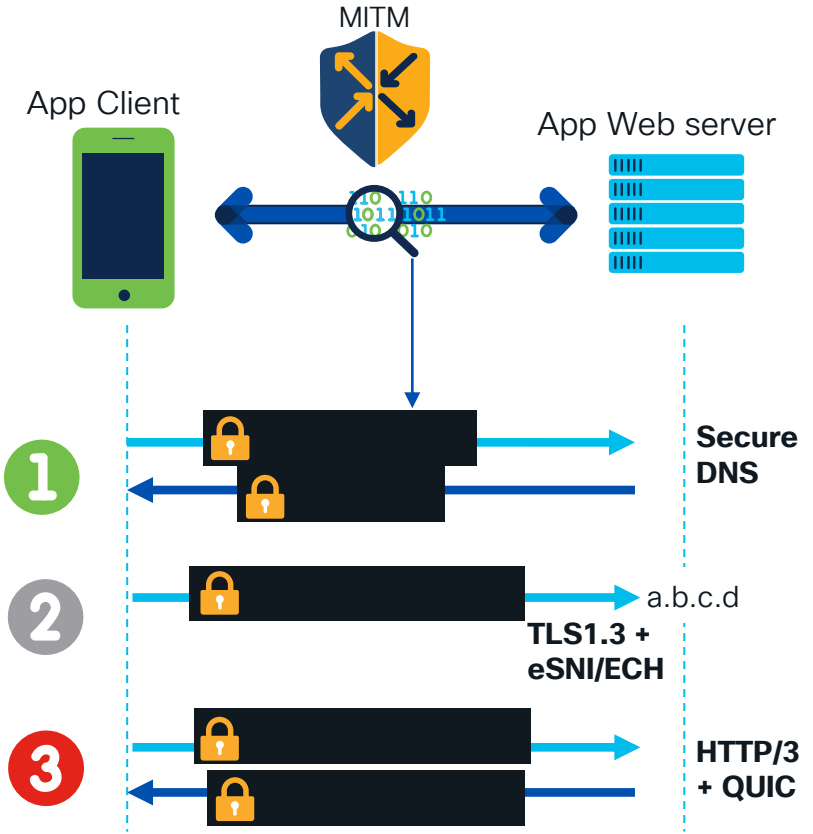

Visibility is lost

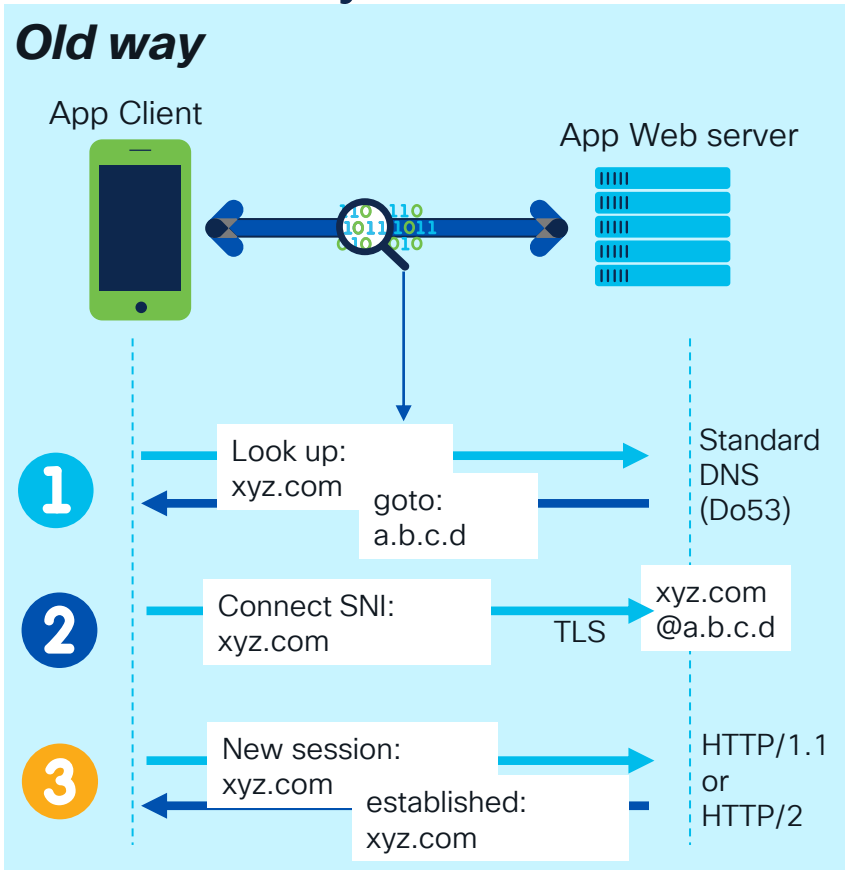

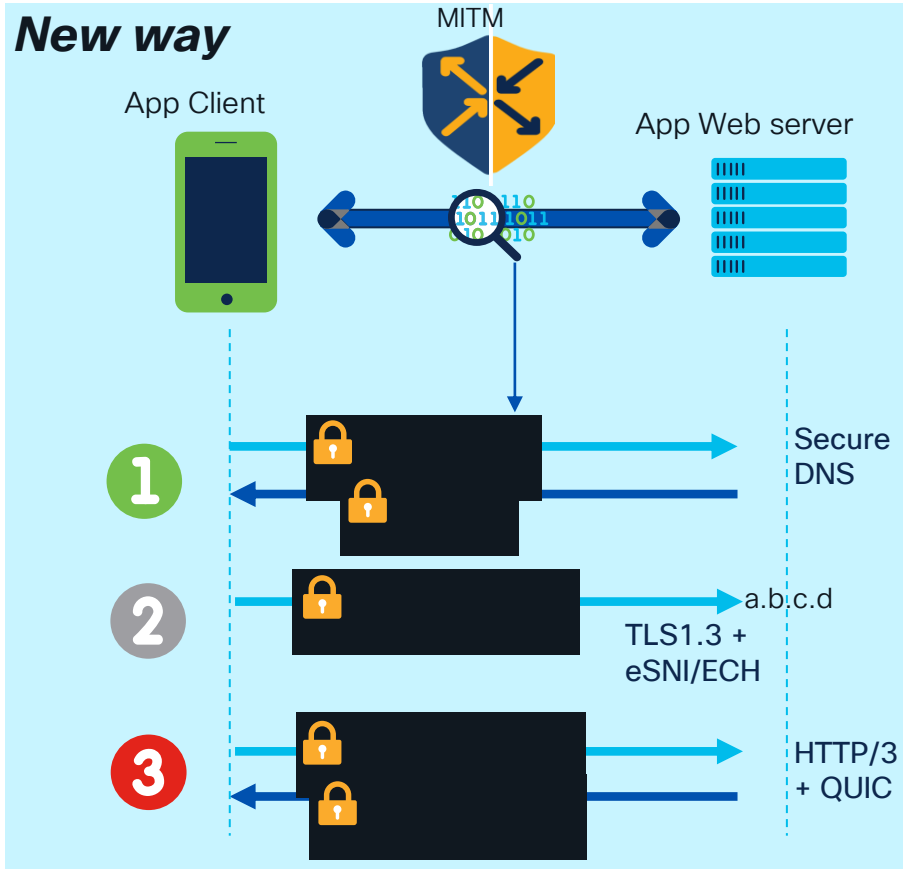

DPI is gone

HTTP/3 Stack = UDP+QUIC+TLS+H3+DoH+eSNI/ECH

Old App Stack

- HTTP/1.1/2
- TLS
- TCP
- IP

New App Stack

QUIC – RFC 9000
HTTP/3 – RFC9114

- HTTP/3
- QUIC + TLS1.3
- UDP
- IP

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

DoH

DoT – RFC7858
DoH – RFC8484

Application Controlled DNS
DNS Traffic not observable

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

eSNI / ECH

RFC8744

Target Domain is opaque / unobservable

DPI Ineffective

including alternative hints e.g. DNS or SNI analysis

Large Scale Adoption

TikTok Uber

cisco Live!

But there is more...

Tunneling is the new normal for APPs

also: welcome to a new threat vector (exfiltration tool?)

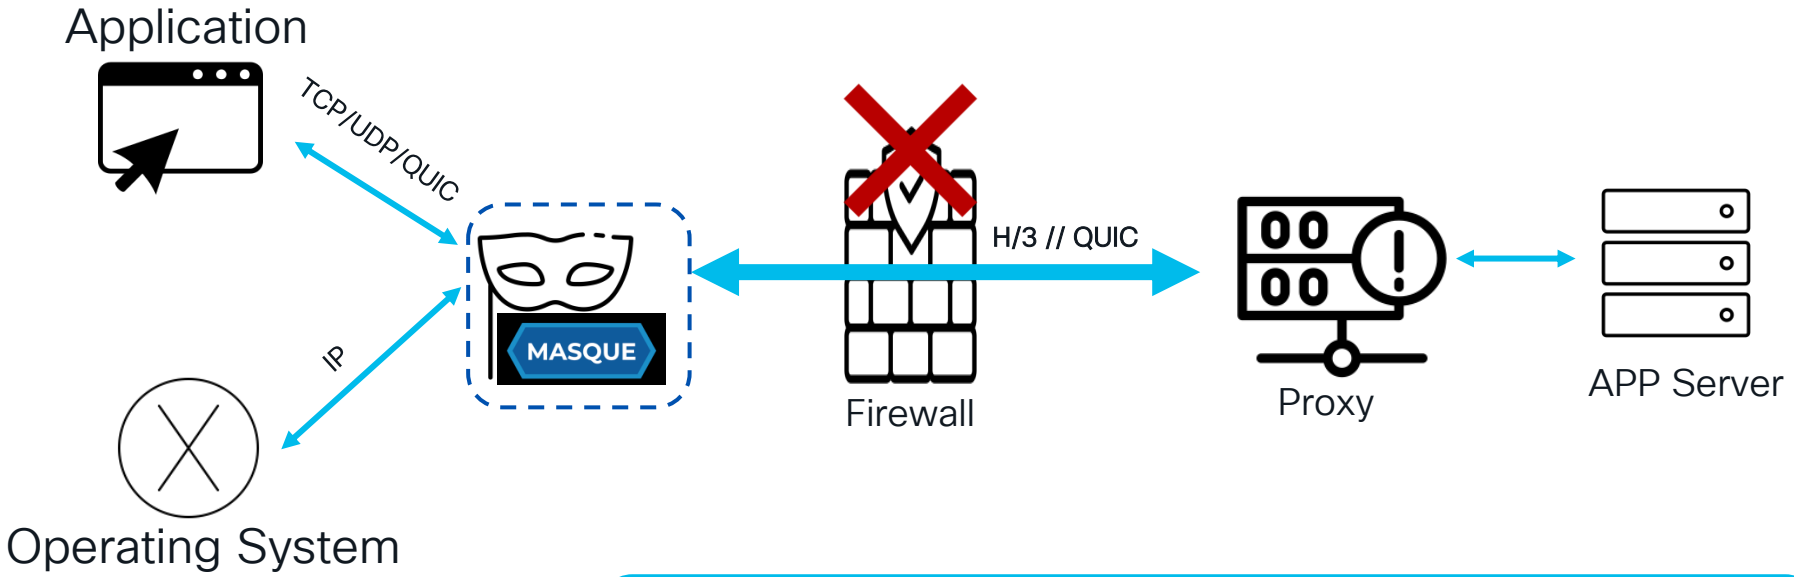

MASQUE

Multiplexed Application Substrate over QUIC Encryption

Goal is to develop mechanism(s) that allow configuring and concurrently running multiple proxied stream- and datagram-based flows inside an HTTP connection.

Options for Masque

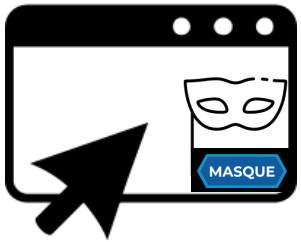

Inside the App

Inside the O/S

Client to O/S

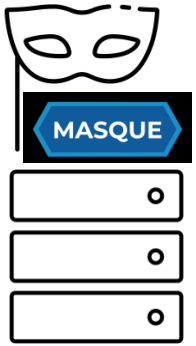

Network Appliance (tunnel IP)

Example: CLOUDFLARE 1.1.1.1

Apple Private Relay: Dual Hop Masque

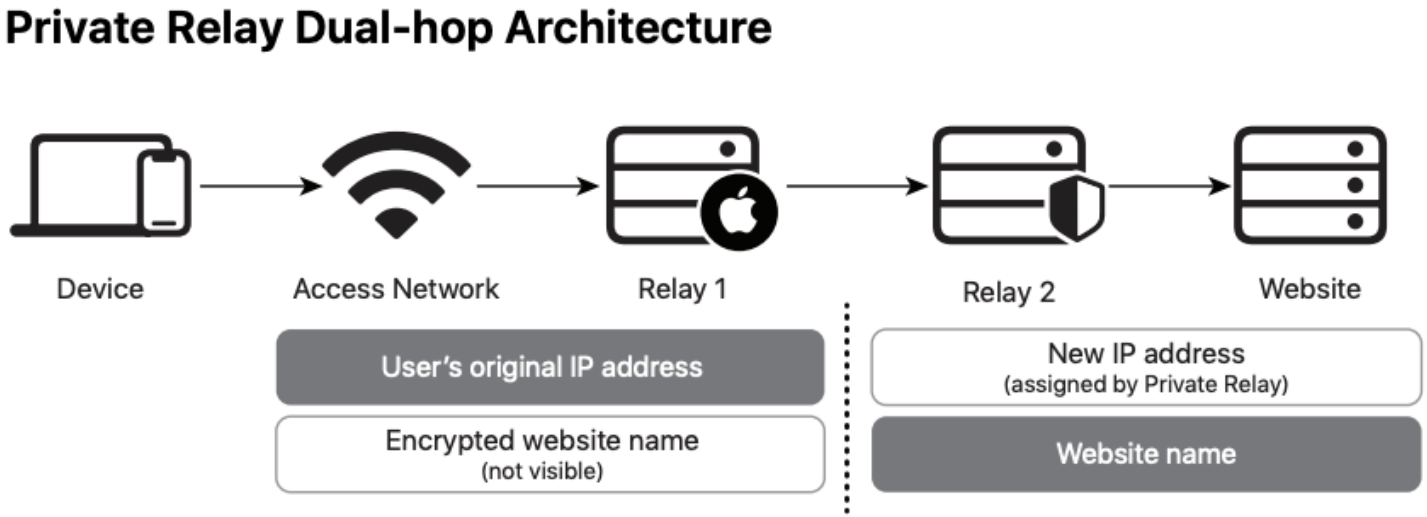

Decoupling users from content

SP Domain has less (or “no”) insights on traffic

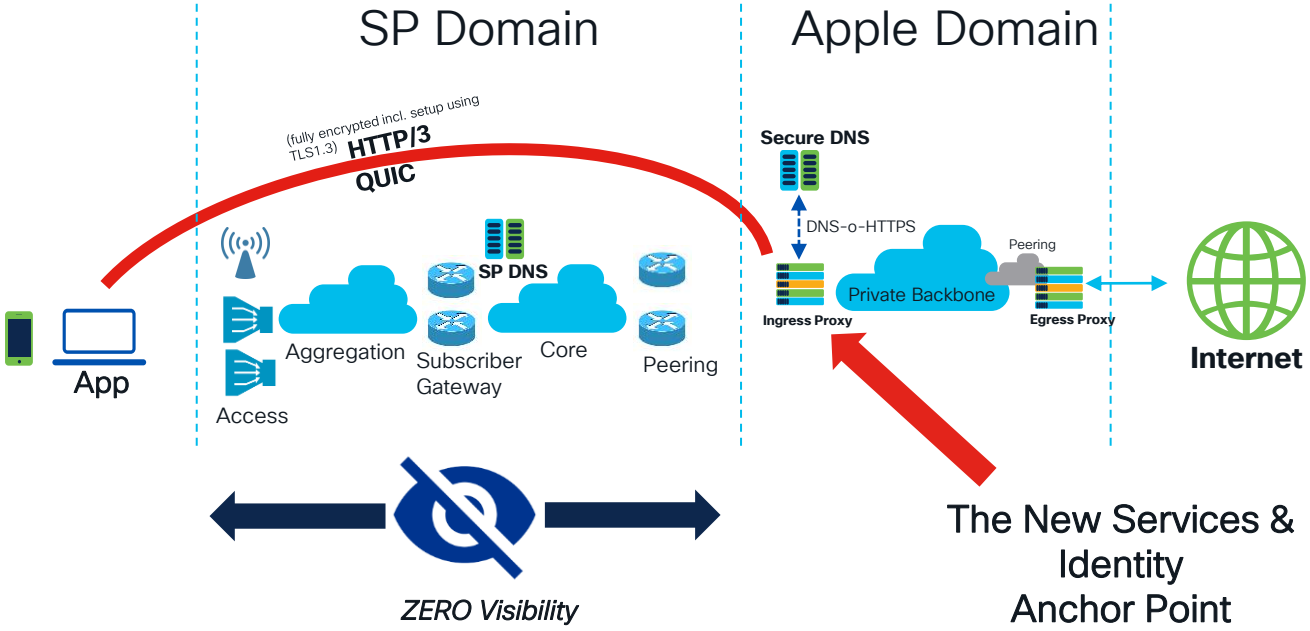

QUIC at MSFT*

- 70% of worldwide front-end servers deployed latest Windows Server with HTTP/3 support
- Chart below shows **all** EXO H2/H3 usage; including browser, mobile and desktop clients

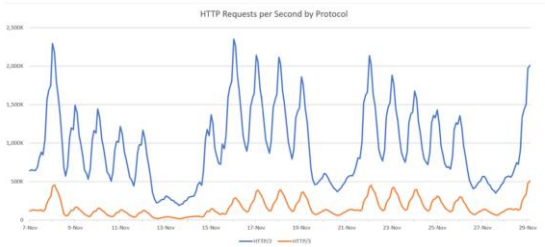

Easy to adopt

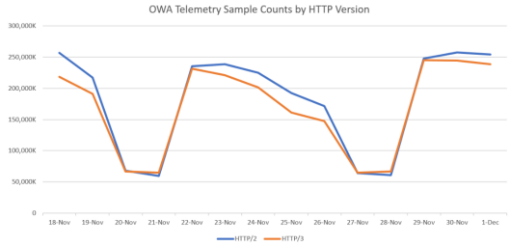

Outlook runs on Quic/H3

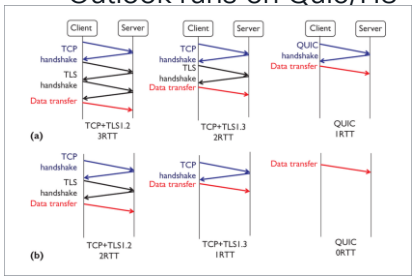

SMBoQUIC – No VPN

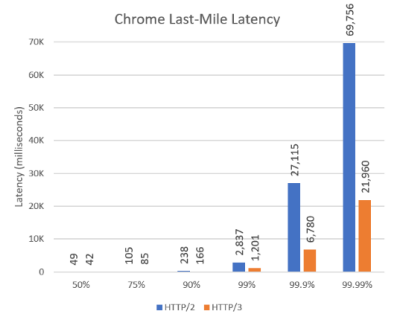

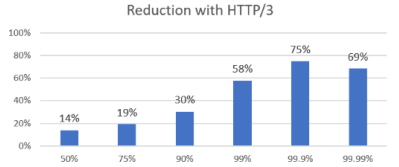

Outlook web access *actually* runs better using H/3

Pervasive across Products

* Source: EPIQ 20212, Nick banks, MSFT

QUIC/H3/DoH stack is in business

Content Delivery Security Privacy Loadbalancing App Infrastructure App Experience

Net Neutrality has effectively been subverted

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)

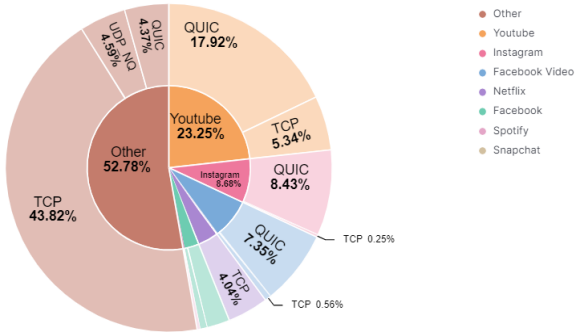

- Net Neutrality implicit assumption is that during network congestion the network will impartially impact all flows – **and that all flows will respond in the same way (TCP assumption)**
- App owned flow control breaks this assumption conclusively
- Therefore ~50% of the traffic in the internet is no longer conformant to neutrality principles

Standards are there

Continuous standard adoption

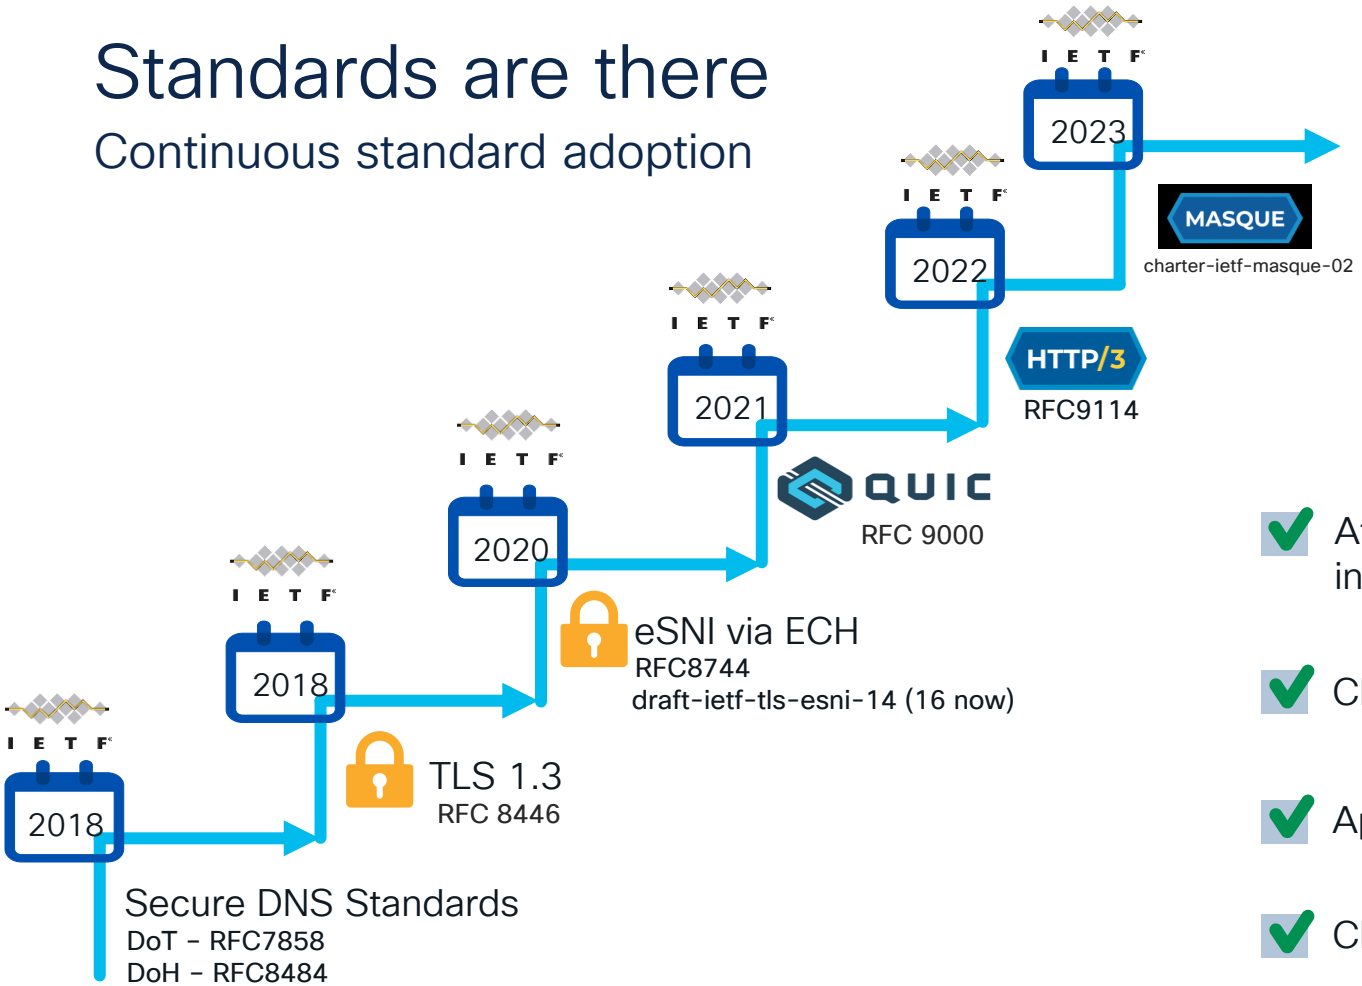

- At scale, in production
- Client (chromium, android, Microsoft Edge, Mozilla)
- Application
- Cloud

The consumers are observing benefits

QoE Drives QUIC Adoption

1.8B Daily Active Users – 3B Monthly
QUIC and H/3 are protocols of choice*

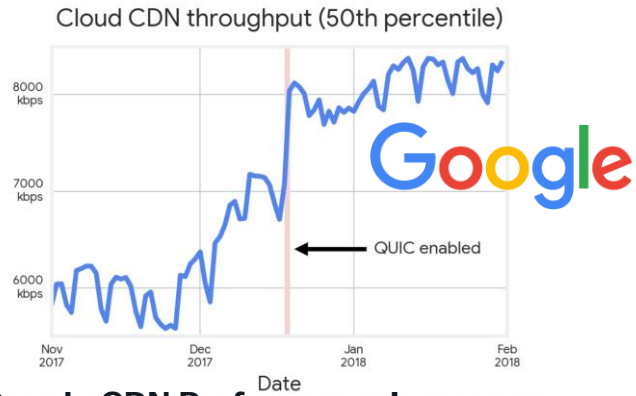

Google CDN Performance increase

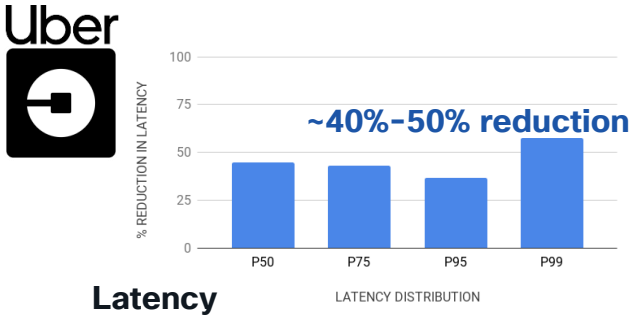

Latency reduced significantly**

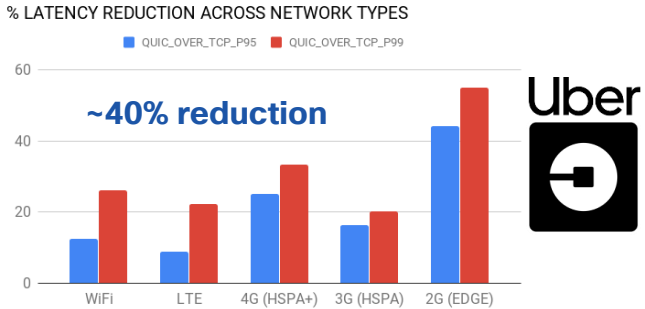

The more fragile the network, the more QUIC excels**

*source Facebook engineering

** source Uber engineering

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*

Business Services

- *VPN*
- *Security*

So what is left?

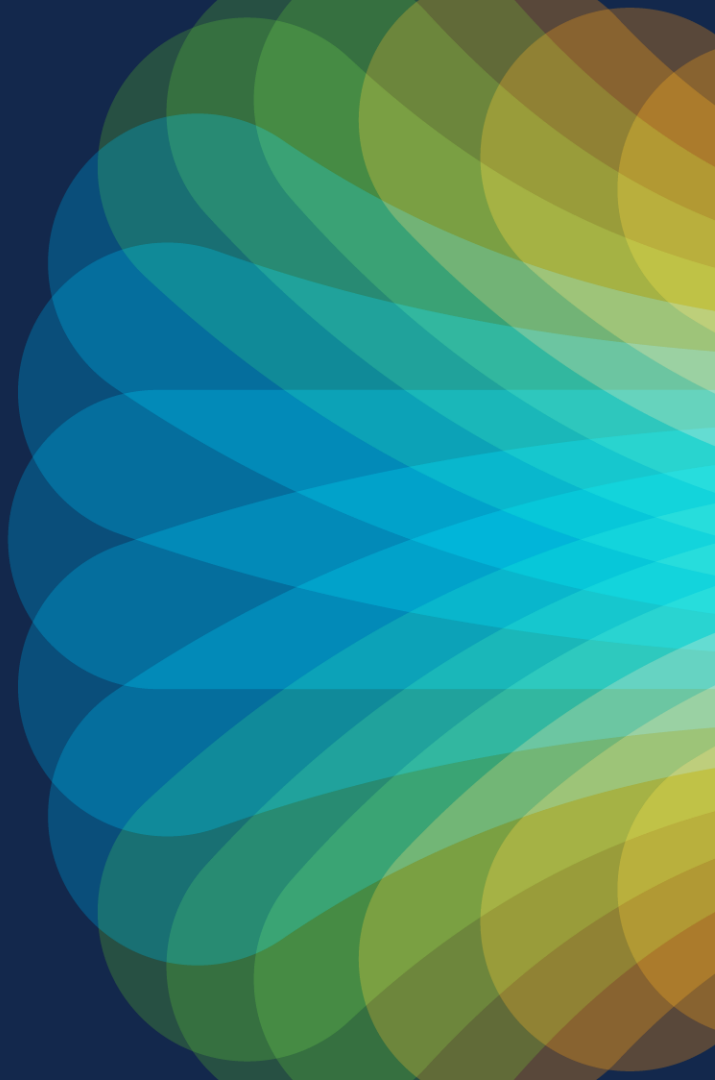

Customers are looking for solutions

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

There is some information that will not go away

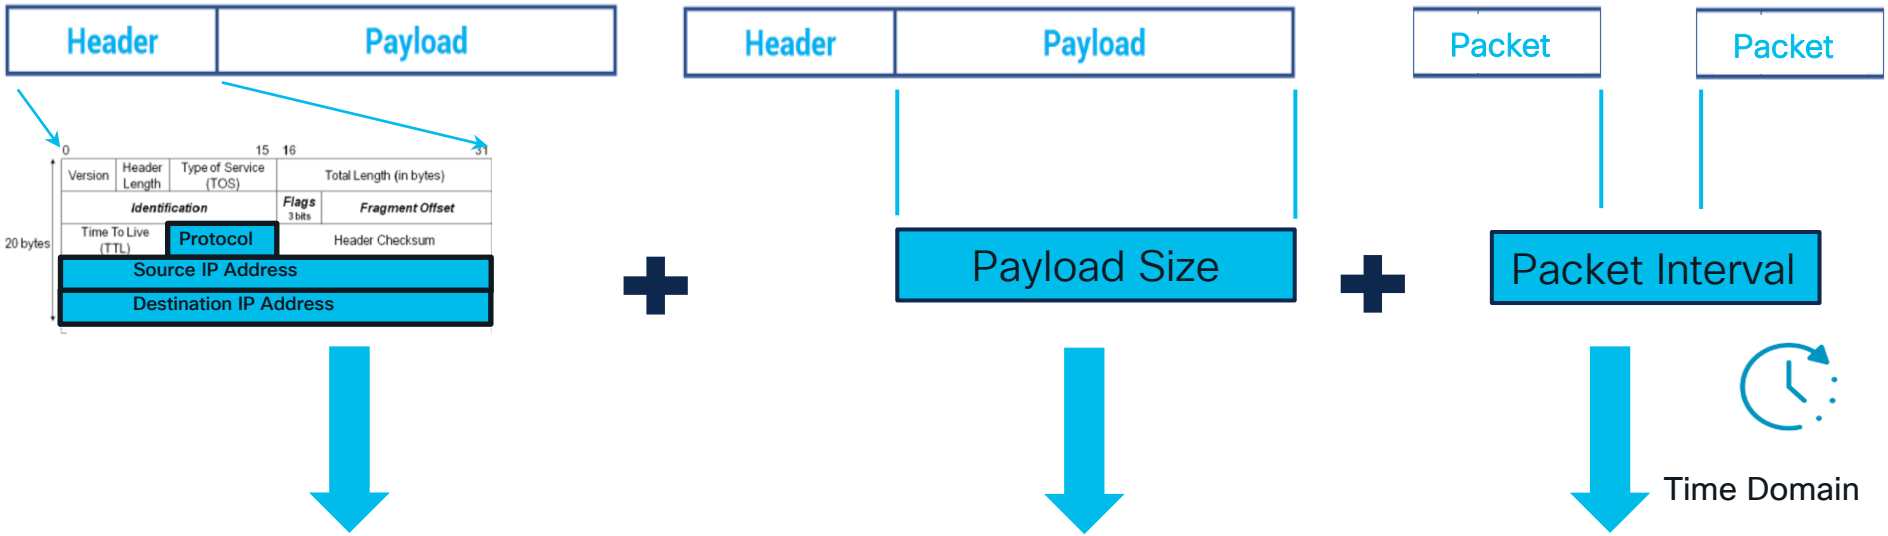

CDN Information

Traffic Volume in Time Information

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

TCP Video Stream Detection

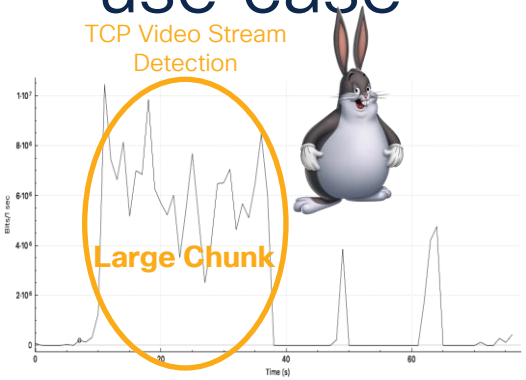

TCP based ABR video players prefer larger, sustained downloads due to high cost of establishing the TCP session and reducing time spent in TCP slow start. Often use HTTP/2 connection. (DASH/HLS) to fix HOL.

QUIC Video Stream Detection

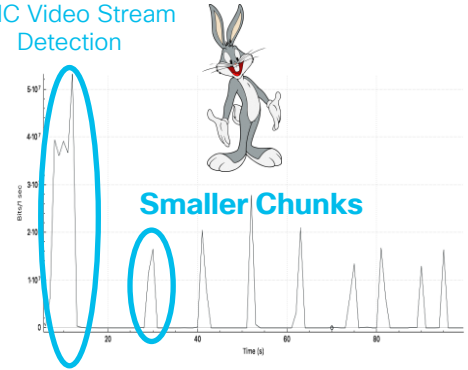

QUIC based ABR video players prefer requesting video in smaller chunks.

Multiple QUIC Streams in many cases to (different) ers

UDP Video Live Stream Detection

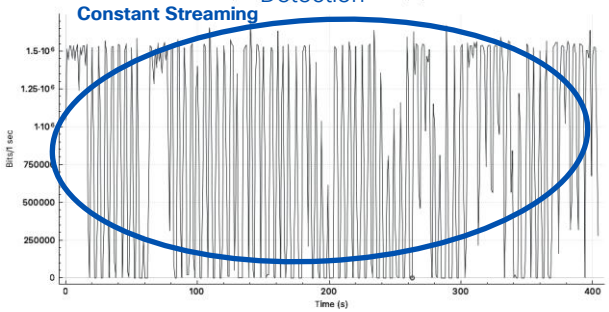

UDP based video players are extremely reliant on consistent network performance. Small buffer, sustained T'put
Applications: YouTube Live, WebEx, Microsoft Teams, Zoom

Download Stream Detection

Premium

NETFLIX

Time Domain Flow recognition

- Observe all flows
- Profile per flow (Time domain matched)
- In real time
- 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

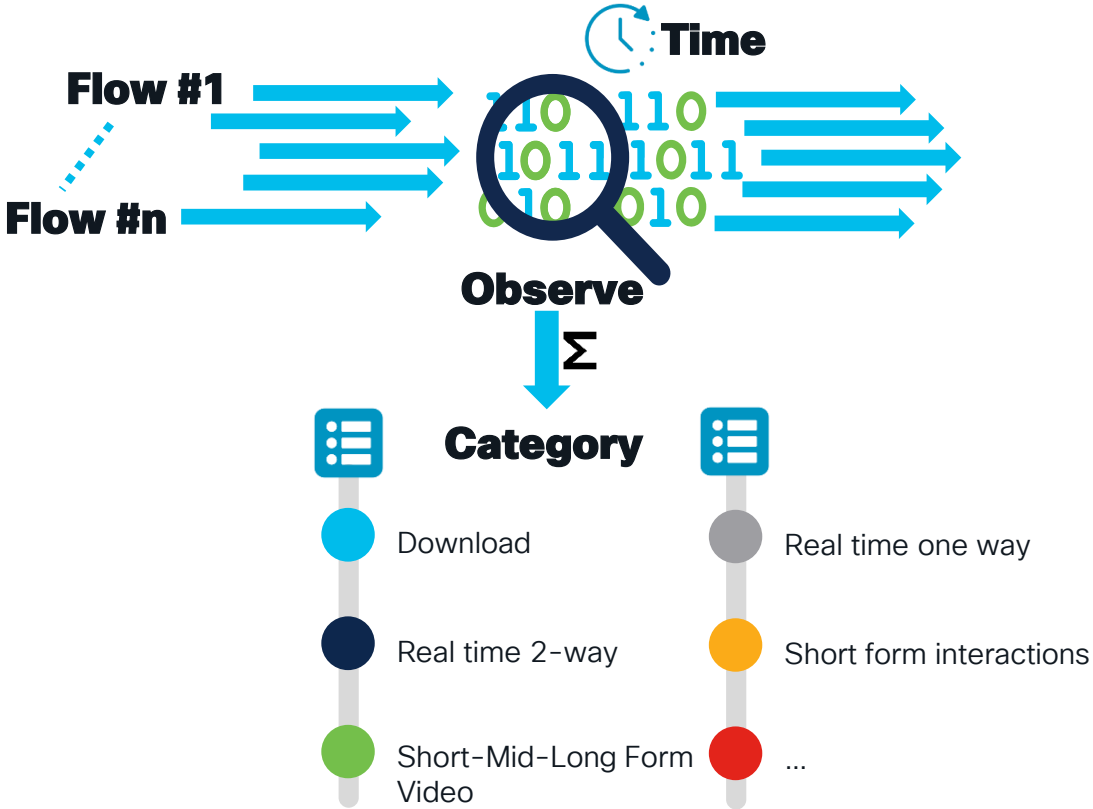

Inferring Traffic Behaviour (e.g. congestion)

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

Reno vs CUBIC vs BBR behaviour*

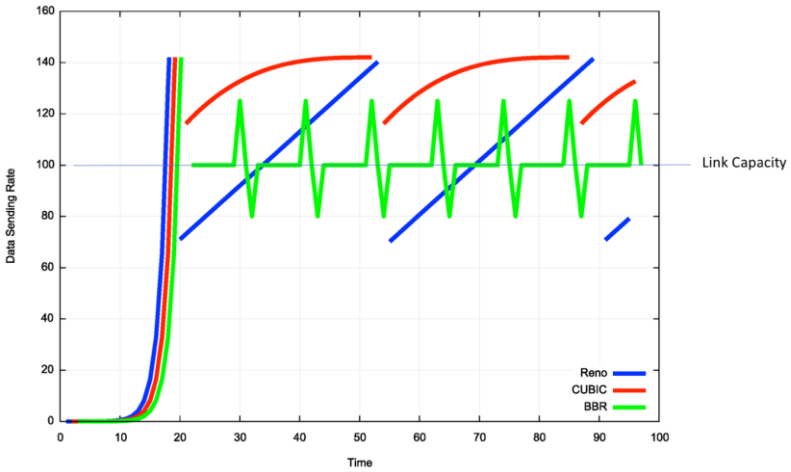

- Assessment of various flows in parallel
- Understand Protocol behaviour: congested or not
- This serves as input for Policy Application

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

Correlation with other network data

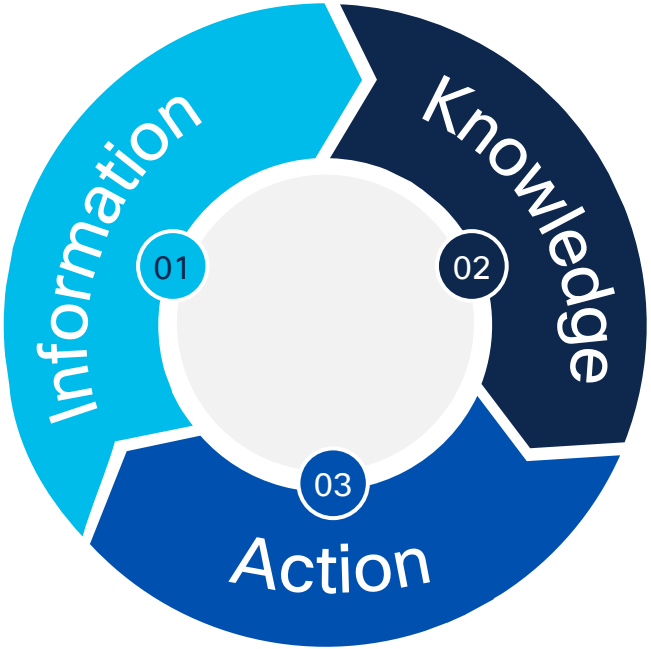

Information

01

- Power consumption
- IP traffic flows
- User data (Profile, KPI, SLA)
- Access data

Knowledge

02

- Correlate IP traffic to users
- Correlate IP traffic to power consumers
- Understand capacity of consumers
- Geo distribution of consumption

Action

03

- Manage usage of power consumers
- Manage services profile (KPI/SLA)
- Plan and upgrade

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

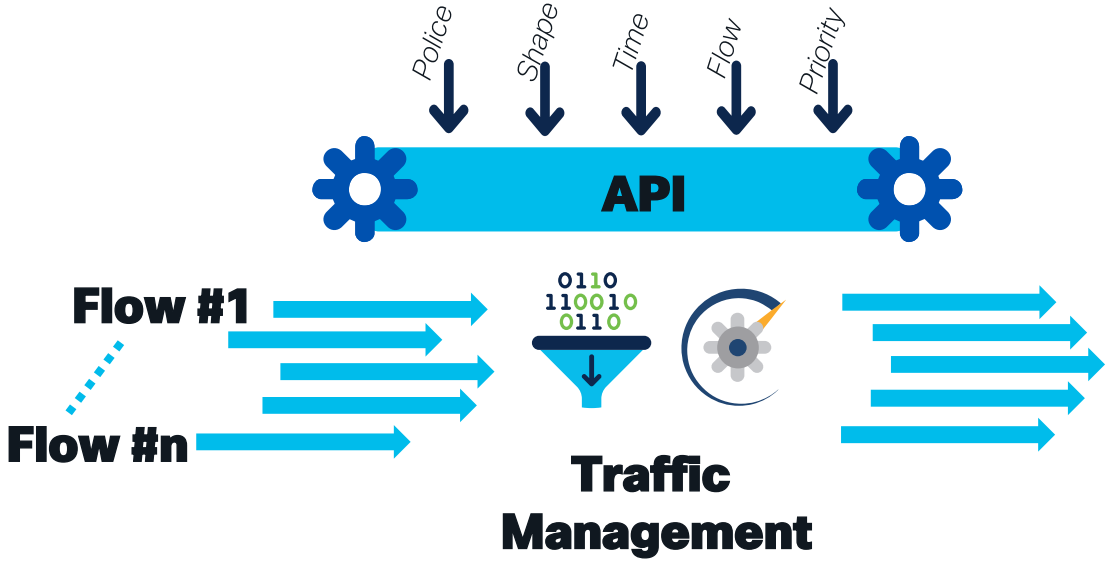

Real Data from Real traffic

Basis for building use cases

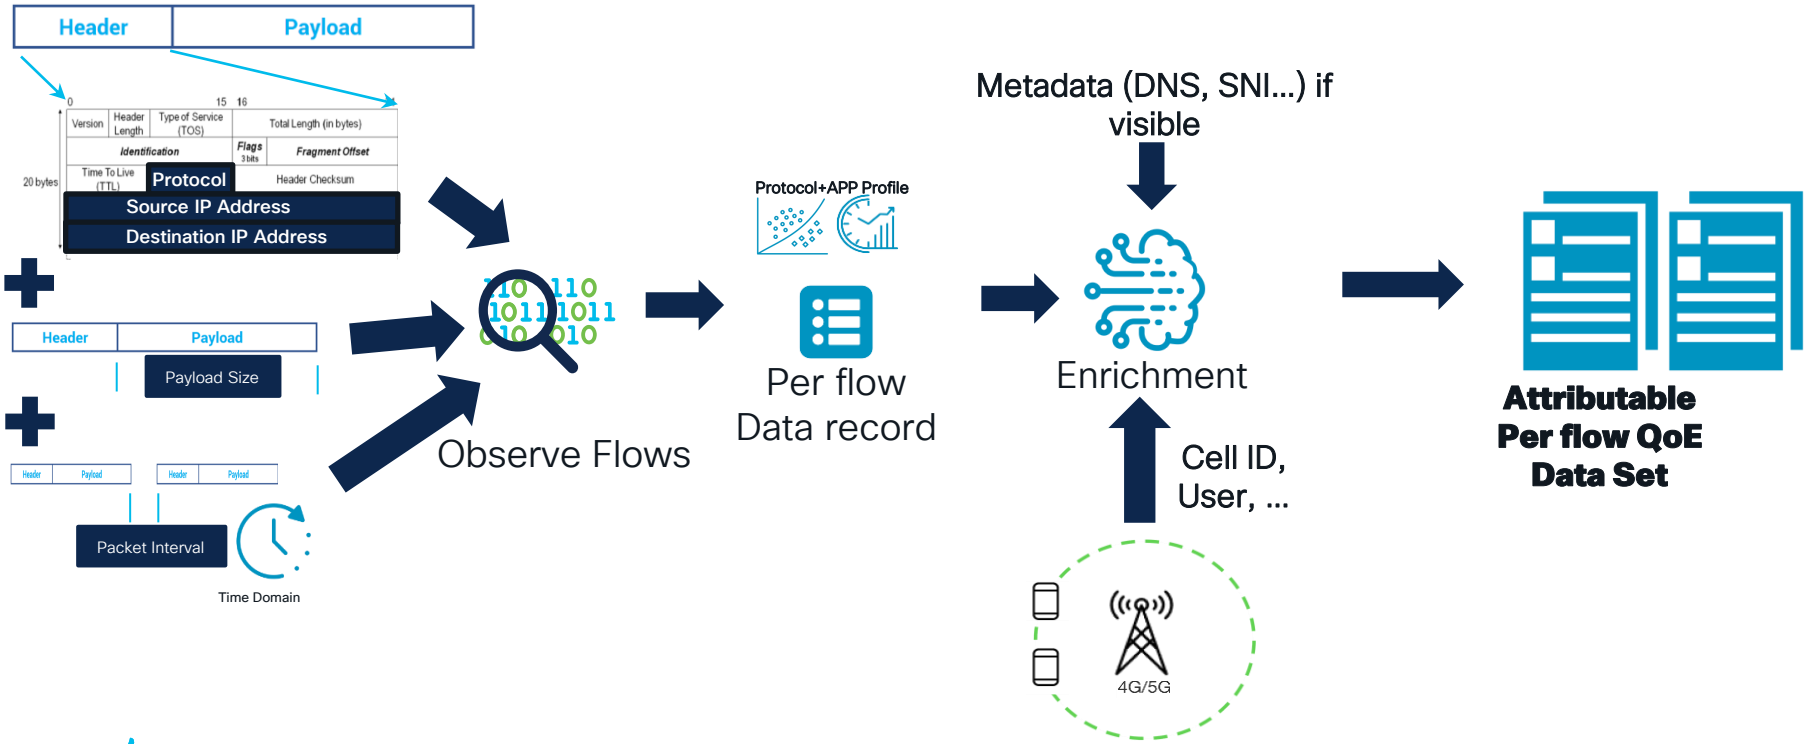

Use Case : Monitoring and analytics inc. QoE

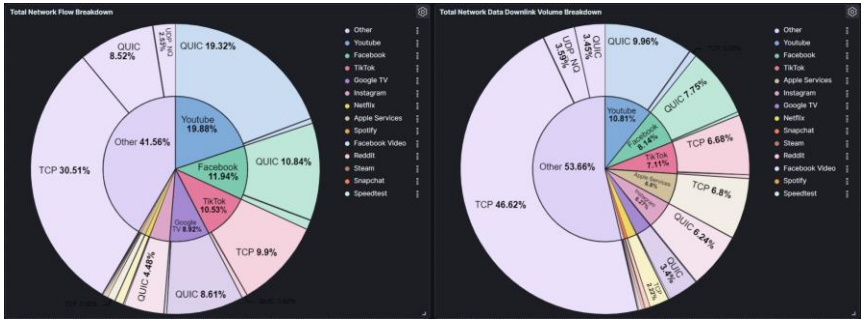

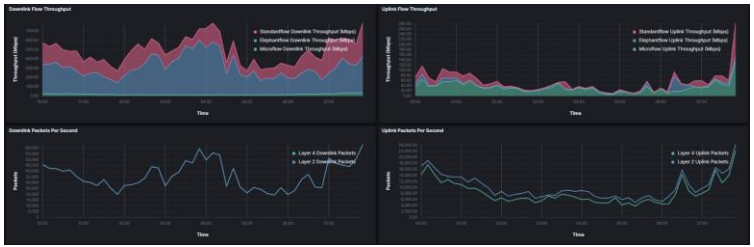

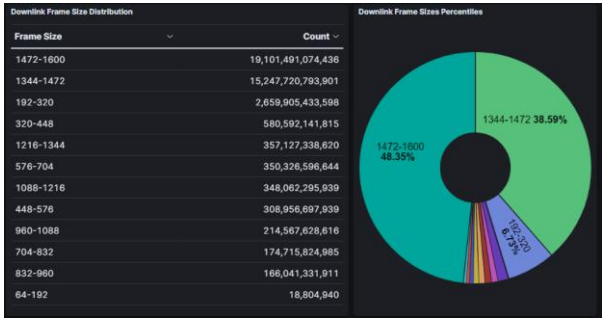

Monitor all flows

- Infer information for Source (DNS, SNI/eSNI), CDN (ECH), Flow Type (Time domain behaviour)
- Enrich with Control Plane Info (S1-MME, S-11…)
- ELK (elastic Search, Logstash, Kibana) analytics engine
- Extensible to enriched CDR production

Compute Metrics

- Application Metrics for Top apps (YouTube, Facebook, Instagram, Netflix, (operator defined app)) : Median throughput, Count of data sessions,Time to first 500KB, Latency, Data volume
- Per Cell performance (aggregate accross all apps / data sessions) : Average Simultaneous Active Subscribers, Median Burst Rate Throughput (measured only during transmit time), Total Data Volume, Time to first 500KB, Latency

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)

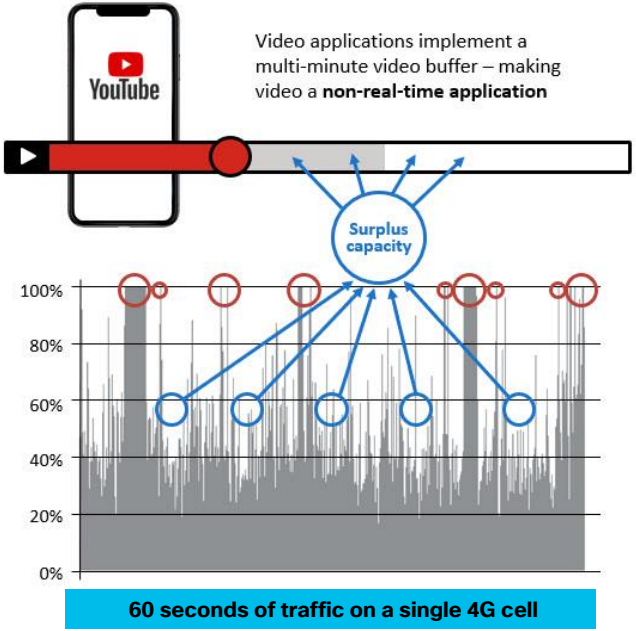

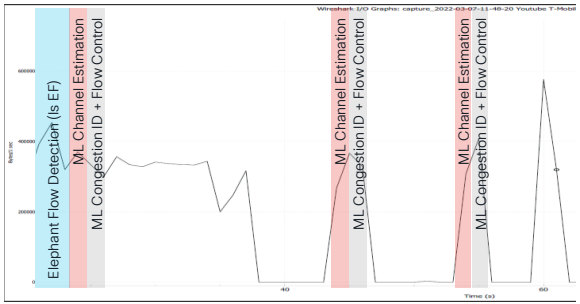

Real World outcome: Tier-1 EU Operator

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

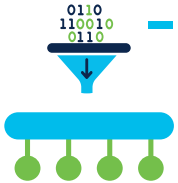

-50%
Network Congestion

90% **Speed Test Improvement**

OOKLA SPEEDTEST

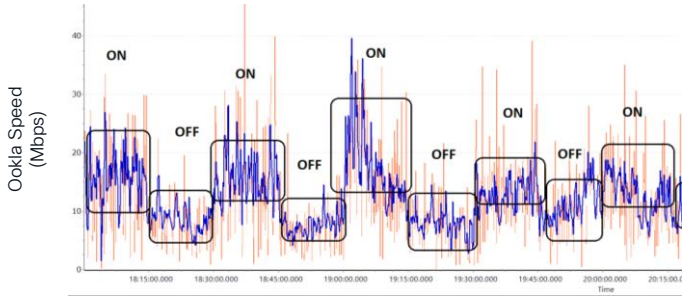

CUTO* Wireline

Enhanced User Experience within SLA Boundaries

Situation

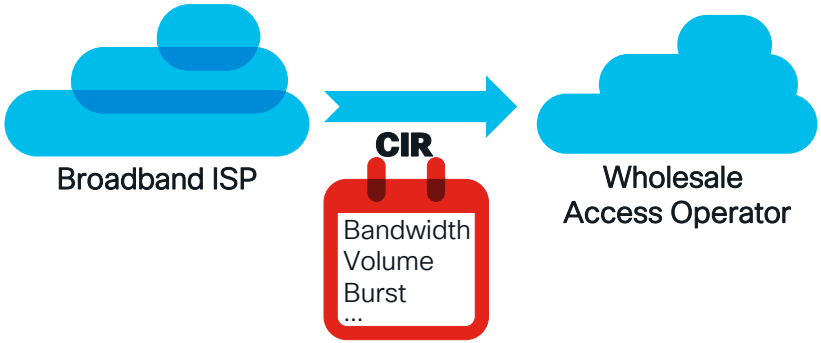

Conform to SLA results in predictable cost ✓

Violate SLA results in additional cost ✗

Indiscriminate Policing leads to bad user experience

Solution

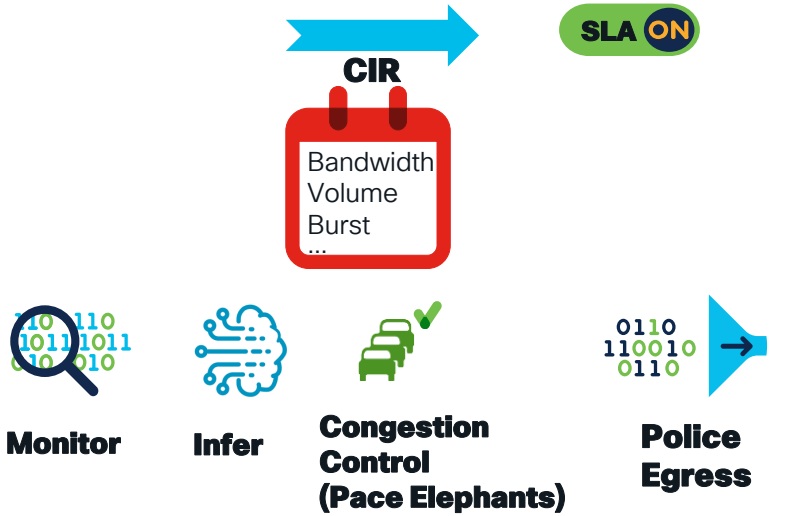

- ***Conform to SLA***
- ***Ensure QoE for every user***
- ***Fair use capability***

*Cisco Ultra Traffic Optimization

Custom Policy Enforcement

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

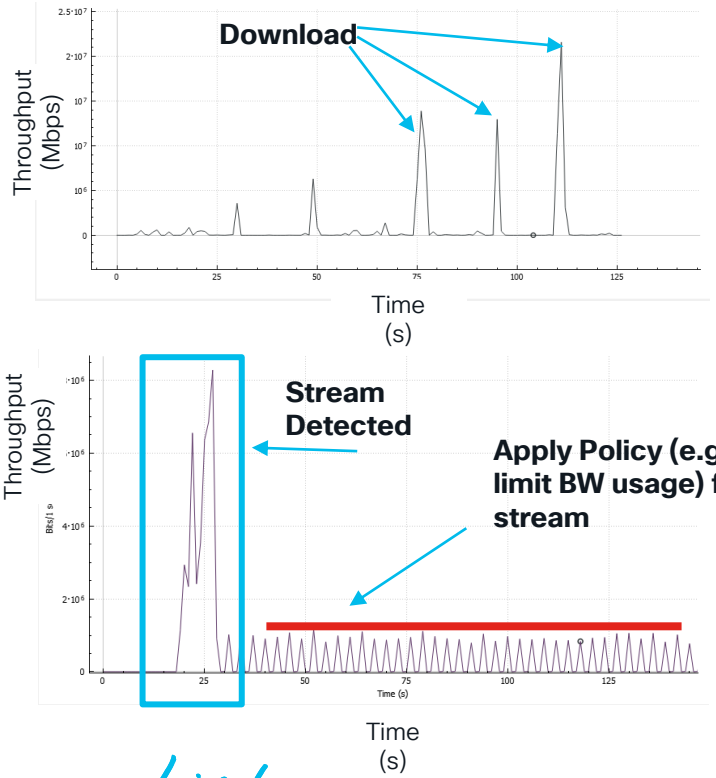

- Same Source/Destination Address
- Differentiate between download versus streaming *on the same SA/DA*
- *Apply Policy per flow type, e.g.*
  - *Download Policy: no action*
  - *Streaming Policy: Limit to set BW profile (police/buffer/...)*

RAN Energy Savings based on IP utilization metrics (BRKSPG-2583)

Honor User Experience

IP Flow Measurements

(Correlate data)

RAN KPI

Decide where to Optimize RAN

Save Energy (RAN Optimization)

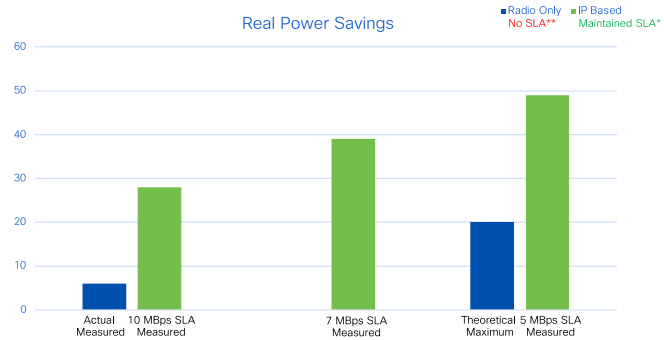

Systematic Approach to Smart Energy Savings

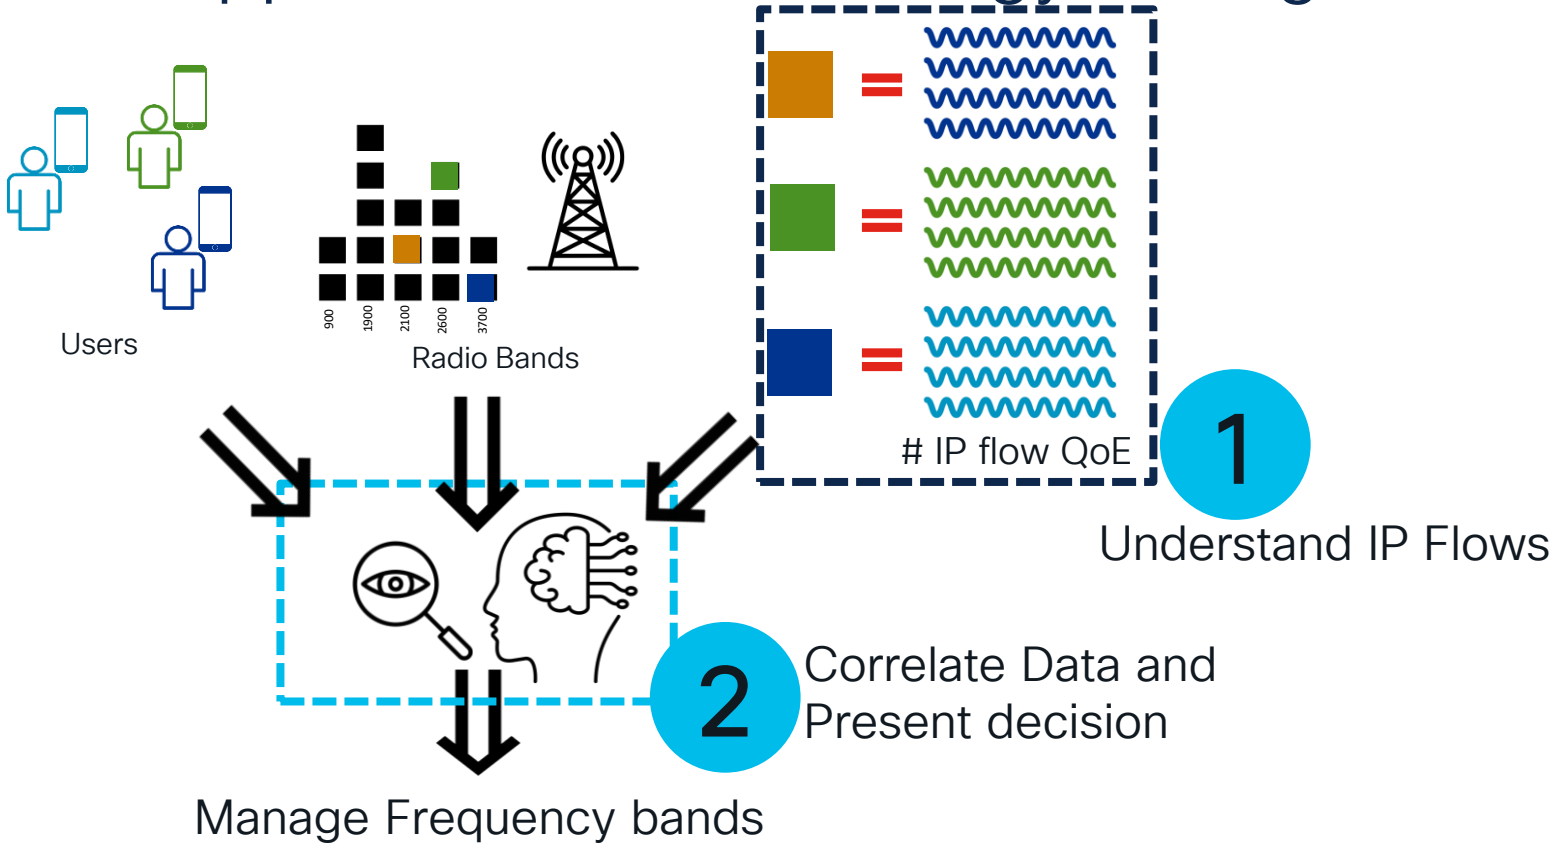

Use Case : Protecting Real-time Traffic

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

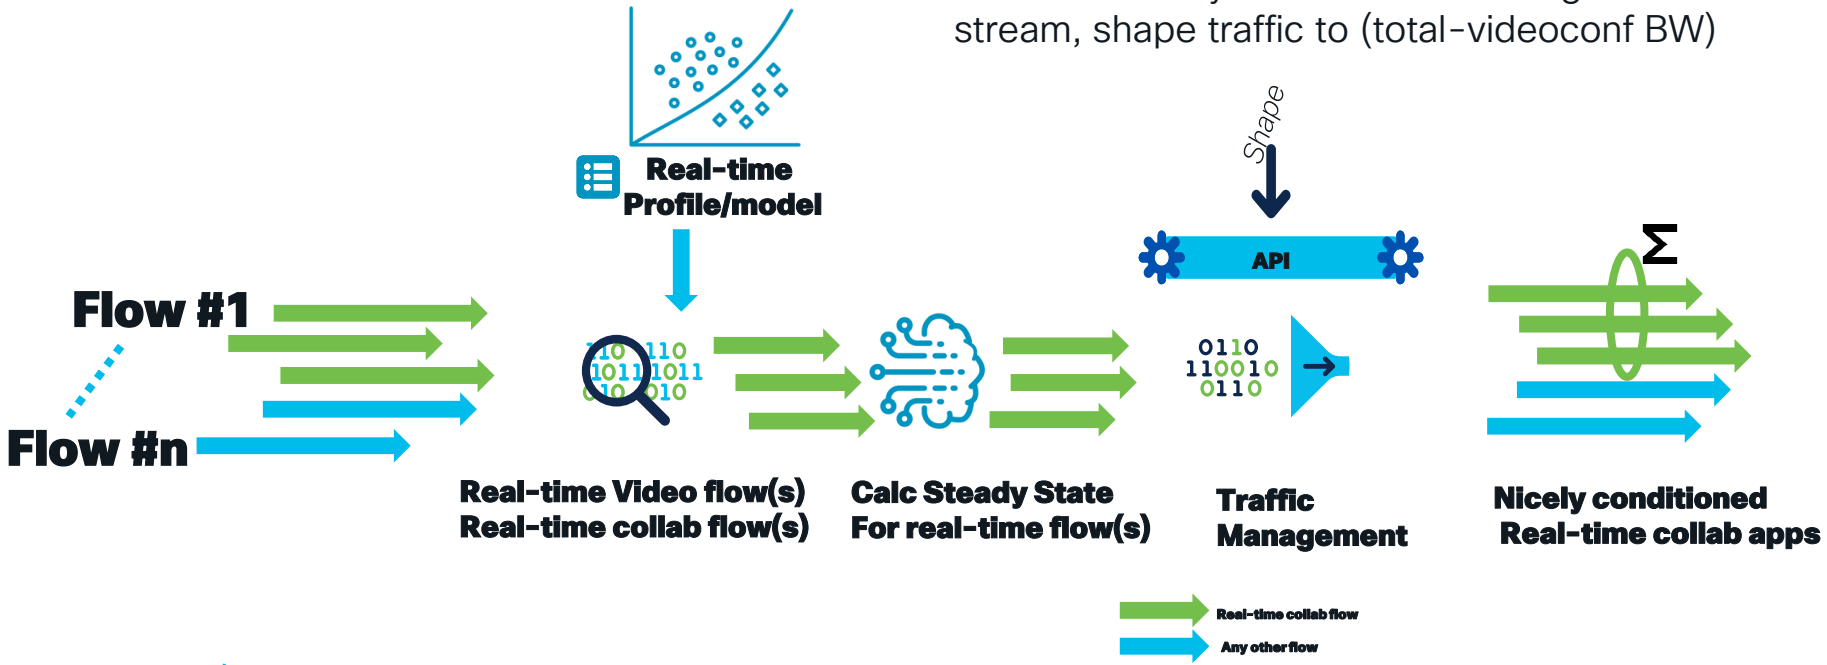

Use Cases Summary

Non-exhaustive list

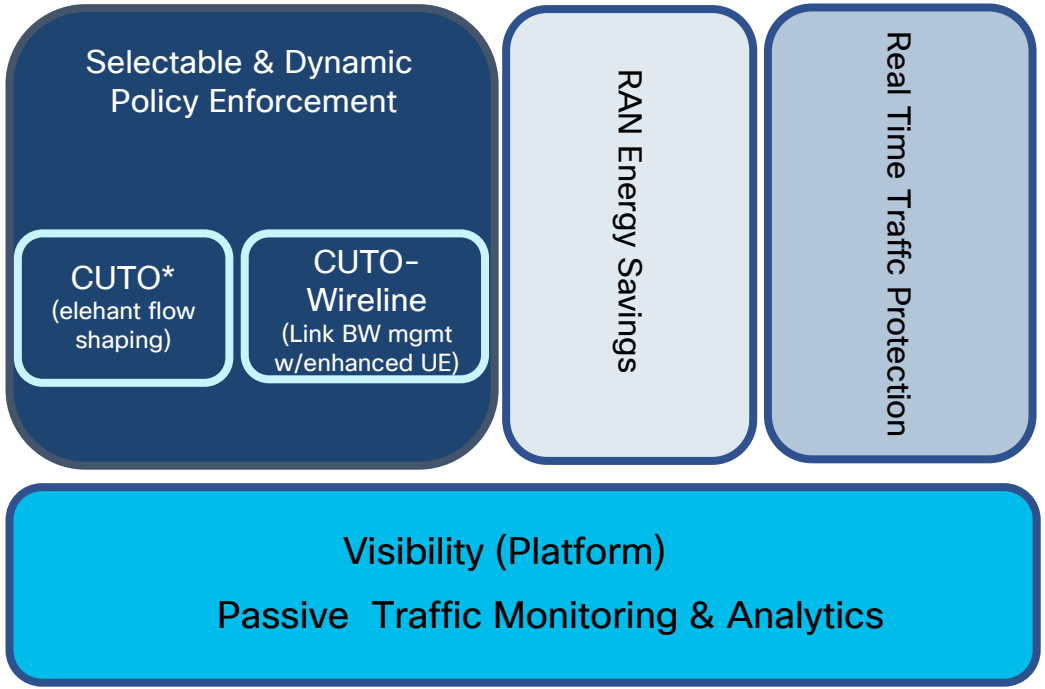

*Cisco Ultra Traffic Optimization

Visibility (Platform)

(Passive) Traffic Monitoring & Analytics
QoE derivation and monitoring

Policy Enforcement Engine

Dynamic Policy Enforcement per (APN|MSISDN|Link|Base Station|...)

CUTO (Dynamic Congestion Alleviation by Elephant Flow Shaping)
CUTO-Wireline (Interconnect link bandwidth management while maintaing an enhanced User Experience)

Protection for Real-Time Traffic

Manage overall link congestion dynamically to protect RTP traffic (videoconf, collaboration, etc)

RAN Energy Savings / Sustainability

Dynamically switching bands on/off at a cell site to match IP based real-time traffic demand & QoE from customers.

Why does this scale

Simple

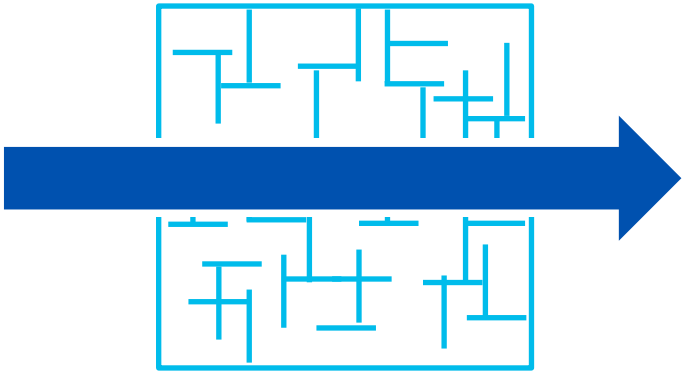

- I only use state on the important/interesting stuff
  - 20% of the flows generate 80% of the volume

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

The bridge to possible

Thank you

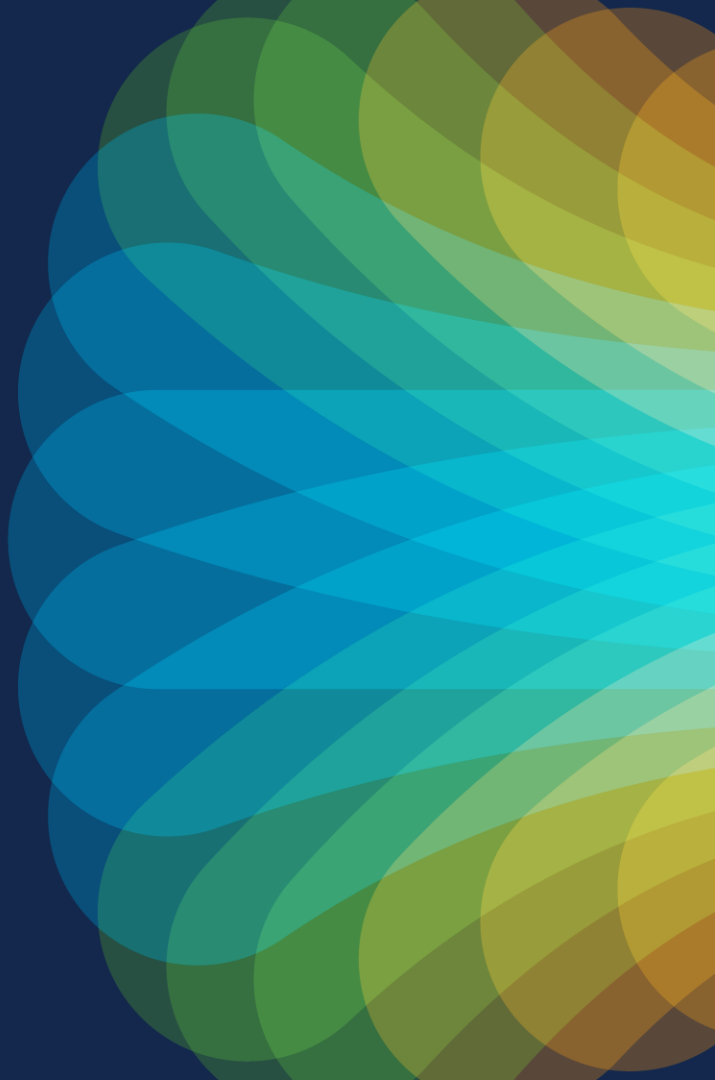

cisco Live!


cisco Live!
Let's go