Smarter Traffic Management
Version 1.0
14 March 2017

This is a White Paper of the GSMA

Security Classification: Non-confidential

Access to and distribution of this document is restricted to the persons permitted by the security classification. This document is confidential to the Association and is subject to copyright protection. This document is to be used only for the purposes for which it has been supplied and information contained in it must not be disclosed or in any other way made available, in whole or in part, to persons other than those permitted under the security classification without the prior written approval of the Association.

Copyright Notice

Disclaimer

The GSM Association ("Association") makes no representation, warranty or undertaking (express or implied) with respect to and does not accept any responsibility for, and hereby disclaims liability for the accuracy or completeness or timeliness of the information contained in this document. The information contained in this document may be subject to change without prior notice.

Antitrust Notice

The information contain herein is in full compliance with the GSM Association's antitrust compliance policy.

Table of Contents

1 Executive Summary 4
1.1 Abbreviations 4
1.2 References 5
2 Introduction 7
3 Network management today 8
3.1 Rationale for network traffic management 8
3.2 Network implementation considerations 9
3.2.1 Bearers and QoS 9
3.3 Traffic trends, Q4 2016 10
3.3.1 Use of encryption 10
3.3.2 Adaptive Bitrate Streaming 10
3.3.3 QUIC 11
3.4 TCP congestion controls in a gigabit Internet 12
3.4.1 The problem: TCP reaction to perceived congestion 12
3.4.2 Another problem: TCP on mobile networks 12
4 Trends in Internet stack evolution, and the impact on networks 14
4.1 Encryption by design 14
4.2 Impact of network middleboxes on transport protocol design 14
4.2.1 Network middleboxes: overview and impact 14
4.2.2 Getting new protocols across network middleboxes 15
4.2.3 Substrates and Signalling 17
5 Smarter Traffic Management 17
5.1 Network management of encrypted traffic 17
5.2 Security considerations 18
5.2.1 Vulnerabilities of network Man in the Middle certificate interception 18
5.2.2 Let’s Encrypt and malware mitigation 18
5.3 Collecting the SMART solutions 19
5.3.1 Guidelines 19
5.3.2 Recommendations from Internet community collaboration 19
5.3.3 Measurement of protocols and network traffic management functions 20
5.3.4 Protocol Optimisation Project 20
5.4 Solution analysis 20
5.4.1 Zero-bit optimisations 21
5.4.1.1 Active queue management and L4S [28] 21
5.4.1.2 Improved congestion controls 22
5.4.2 One-bit optimisations 22
5.4.3 Multi-bit optimisations 23
5.4.3.1 Path Layer UDP Substrate (PLUS) 23
5.4.3.2 Mobile Throughput Guidance 24
5.4.4 User-based prioritisation 24
6 2020 outlook for network management 25
6.1 Regulation 25

6.1.1 Privacy 25
6.1.2 Net Neutrality 25
6.2 5G 25
6.2.1 NFV and SDN 25
6.2.2 Network slicing 26
6.3 Internet of Things (IoT) 26
6.4 Mobile Edge Computing 26
6.5 Secure caching 27
7 Beyond 2020 27
7.1 ETSI NGP 27
8 Conclusions 28
Annex A Document Management 29
A.1 Document History 29
A.2 Other Information 29

1 Executive Summary

This document captures the research of the GSMA SMART sub-group of the Internet Group (formerly Web Working Group). It describes approaches to 'Smarter Traffic Management' that allows operators to manage network data that is increasingly encrypted, and being delivered via new protocols and congestion-control algorithms. It also includes summaries of interworking between GSMA and IAB/IETF to improve cellular delivery of Internet content, detailing the background to issues that stimulated this engagement; and also, covers non-technical considerations including regulation and security. The document concludes with a set of recommendations for the coming 5G era where the range of traffic types, and general load on the network, is expected to increase radically. The group's associated document on 'Secure Content Distribution', 'Mobile Throughput Guidance' and 'Captive Portals' provide further detail on those topics.

1.1 Abbreviations

Term	Description
3G	3rd Generation Mobile Network
3GPP	3rd Generation Partnership Project
5G	5th Generation Mobile Network
ABR	Adaptive Bitrate Streaming
APN	Access Point Name
AQM	Active Queue Management
AR/VR	Augmented Reality / Virtual Reality
ARQ	Automatic Repeat Request (L1)
BBR	Bottleneck, Bandwidth and Roundtrip propagation time
COAP	Constrained Application Protocol
DCTCP	Data Center TCP
DiffServ	Differentiated Services
DPI	Deep Packet Inspection
DSCP	Differentiated Services Code Point
DVD	Digital Versatile Disc
ECN	Explicit Congestion Notification
GPU	Graphical Processing Unit
HARQ	Hybrid Automatic Repeat Request (L2)
HDMI	High Definition Multimedia Interface
HSPA	High Speed Packet Access
HTTP 2.0	Hypertext Transfer Protocol version 2
HTTPS	Hypertext Transfer Protocol Secure
IAB	Internet Architecture Board
ICCRG	Internet Congestion Control Research Group
IETF	Internet Engineering Task Force
IMS	IP Multimedia Subsystem

Term	Description
L4S	Low Latency Low Loss Scalable throughput
LTE	Long Term Evolution
MAP	Measurement and Analysis of Protocols
MTG	Mobile Throughout Guidance. A network-calculated information element which recommends a sustainable bandwidth to flow endpoints.
NAT	Network Address Translation
PLUS	Path Layer UDP Substrate
QoS	Quality of Service
QUIC	Quick Internet UDP Connections, Googles protocols for faster Internet
SC	Service Characteristics
SEMI	Stack Evolution in a Middlebox Internet
SNI	Signal to Noise Interference
SPDY	Pronounced as Speedy, Google's networking protocol to decrease latency
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol
UE	User Equipment
URL	Uniform Resource Locator
W3C	World Wide Web Consortium

1.2 References

Ref	Doc Number	Title
[1]	IP 5 tuples	http://www.globalspec.com/reference/67153/203279/ipv4-five-tuple-classification
[2]	QoS SC	QoS Service Characteristics, 3GPP Technical Specification, TS 23.107 http://www.etsi.org/deliver/etsi_ts/123100_123199/123107/10.01.00_60/ts_123107v100100p.pdf
[3]	Let's Encrypt	Let's Encrypt, Free SSL/TLS certificates. https://letsencrypt.org
[4]	ABR	Adaptive Bitrate streaming for mobile networks https://en.wikipedia.org/wiki/Adaptive_bitrate_streaming
[5]	Google's BBR	Google Bottleneck, Bandwidth and Roundtrip; TCP Congestion Control http://cacm.acm.org/magazines/2017/2/212428-bbr-congestion-based-congestion-control/fulltext
[6]	IAB Internet Confidentiality	Statement on Internet Confidentiality, Internet Architecture Board, November 2014. https://www.iab.org/2014/11/14/iab-statement-on-internet-confidentiality/
[7]	ABI SEMI workshop	Tunnelling through inner space, https://www.iab.org/wp-content/IAB-uploads/2014/12/semi2015_briscoe.pdf, SEMI (Stack Evolution in a Middlebox Internet) IAB workshop, January 2015
[8]	GSMA-IAB	Managing Radio Networks in an Encrypted World', joint IAB/GSMA

Ref	Doc Number	Title
	MaRNEW workshop	workshop, https://www.iab.org/activities/workshops/marnew/
[9]	IETF Measurement and Analysis for Protocols (MAP)	Measurement and analysis for Protocols (MAP), active IETF research group, https://datatracker.ietf.org/rg/maprg/charter/
[10]	HOPS Research Group	How ossified is the protocol stack? (HOPS), proposed IETF research group superseded by MAP. https://datatracker.ietf.org/doc/charter-irtf-hopsrg/
[11]	Pervasive monitoring	IETF-RFC7258; Pervasive monitoring is an attack', S. Farrell, https://tools.ietf.org/html/rfc7258
[12]	IETF AQM	Active Queue Management and Packet Scheduling
[13]	IETF TCP RACK	A time-based fast loss detection for TCP, Y. Cheng , IETF 94 https://www.ietf.org/proceedings/94/slides/slides-94-tcpm-6.pdf
[14]	TCP PRAGUE	
[15]	LEDBAT	IETF-RFC6817; Low Extra Delay Background Transport (LEDBAT), Hazel et al., https://tools.ietf.org/html/rfc6817
[16]	DSCP	IETF-RFC2474; Definition of the Differentiated Services Field in IPv4 and IPv6 headers , https://tools.ietf.org/html/rfc2474
[17]	SIGCOMM	Network Cookies
[18]	ETSI Next Generation Protocols	[NGP], Scenarios for Next Generation Protocols, ETSI ISG NGP, 2016 http://www.etsi.org/news-events/news/1135-2016-10-news-etsi-next-generation-protocols-group-releases-first-specification
[19]	RINA	Recursive Internetwork Architecture, IRATI, http://irati.eu/the-recursive-internetwork-architecture/
[20]	SEMI Workshop	IN-network processing, User-level Stacks and the Future of Internet Evolution; for IAB SEMI workshop; Huici, Honda, Raiciu; January 2015
[21]	PLUS minutes	https://datatracker.ietf.org/doc/minutes-96-plus/, IETF 96, June 2016
[22]	1-bit LO/LA Loss/Latency Tradeoff	Loss/Latency tradeoff, ACCORD BoF IETF 95, https://www.ietf.org/proceedings/95/slides/slides-95-accord-1.pdf https://tools.ietf.org/html/draft-you-tsvwg-latency-loss-tradeoff-00
[23]	WWG-04 GSMA whitepaper	Network Management of Encrypted Traffic', GSMA, February 2015, http://www.gsma.com/newsroom/wp-content/uploads/WWG-04-v1-0.pdf
[24]	W3C Technical Architecture Group	End-to-End encryption and the Web, http://www.w3.org/2001/tag/doc/encryption-finding/

Ref	Doc Number	Title
[25]	ConEX	Congestion exposure use cases, B. Briscoe, https://tools.ietf.org/html/rfc6789
[26]	MTG	Mobile Throughput Guidance, K.Smith for GSMA SMART, December 2016
[27]	ICCRG	Congestion Control for 4G/5G networks, I. Johannson, IETF 96. https://www.ietf.org/proceedings/96/slides/slides-96-iccrg-1.pdf
[28]	L4S	Low Loss, Low Latency Scalable Throughput BoF, IETF 96, https://datatracker.ietf.org/wg/l4s/meetings/
[29]	Vision 2020	[VISION2020], 'Vision 2020', GSMA Intelligence, 2014, https://www.gsmaintelligence.com/research/2014/02/vision-2020-white-paper/421/
[30]	Effect of Ubiquitous Encryption	'Effect of Ubiquitous encryption', K. Moriarty & A. Morton, IETF https://tools.ietf.org/html/draft-mm-wg-effect-encrypt-05
[31]	PIE and FQ_Codel	https://www.ietf.org/proceedings/88/slides/slides-88-iccrg-4.pdf
[32]	STAR	https://datatracker.ietf.org/doc/draft-sheffer-acme-star-lurk/
[33]	QUIC charter	https://datatracker.ietf.org/wg/quic/charter

2 Introduction

In July 2014, the GSMA ENCRY sub-group was formed with the initial goal of 'mitigating the impact of HTTP 2.0/SPDY, and the rise in encryption' – in particular the rise in HTTPS (HTTP over TLS) from roughly 5% by volume in 2011 to less than 50% today. The perception at the time was that operators' network traffic management systems, which had for a long time had cleartext access to URLs and another traffic metadata, could become 'blinded' by encryption. This in turn would hinder traffic categorisation, flow control per traffic type, content filtering and other traffic management functions.

18 months of research followed, including observation of the impact of increased encryption on live networks. ENCRY was formed primarily of mobile network operators and mobile equipment vendors. Additionally, engagement with the broader Internet content community included the Internet Architecture Board (IAB) and Internet Engineering Task Force (IETF). The key conclusions of the ENCRY sub-group were:

- **There is no evidence that HTTPS encryption harms customers**
  Operators were unable to provide or did not share evidence that HTTPS was causing a problem to either the network throughput, or to perceived customer experience. The improvement in end-device computing power has allowed client encryption/decryption to support encrypted sessions with no significant impact to the end user, albeit with a battery drain on lower-end devices. 3GPP networks do support a model of multiple access bearers which can provide different Quality of Service (QoS) characteristics

across traffic types; however, it was discovered that operators do not in practice utilise multiple bearers for this purpose – rather they schedule all Internet traffic as 'best effort'. Hence the QoS benefits of traffic categorisation are generally non-applicable today.

- **There is no evidence that HTTPS encryption harms networks**
  The 'network protection' toolkit available to network operators does not rely on access to application layer metadata. IP 5-tuples [1], transport-layer rate limiting, protocol detection and traffic heuristics are not affected by session encryption at higher layers. Network congestion is not increased through the use of encryption, and where congestion commonly appears (at the radio access network) there is little practical advantage to knowing the traffic metadata.

- **Content filtering is impacted**
  Content filtering is hindered due to the full URL being unavailable to the network in an HTTPS session: so for a domain with mixed family/adult content, the network cannot distinguish which type it is.

The requirements for any filtering are out of scope of a technical forum, but rather a topic for regulatory fora.

Following these conclusions, ENCRY evolved into a platform to discuss technologies that impact the following 'mission statement':

Improving customers' Internet experience without breaching their privacy.

The encryption plays two roles: firstly to help secure the customer session outside of the operator network (and hence improve their Internet experience), and secondly by making operators and vendors review, evaluate and improve upon the various network techniques that have built up over the last 15 years or so as they were largely built to handle cleartext.

The phrase coined to capture these revised techniques was: **Smarter traffic management**, and hence the sub-group was renamed to **SMART.**

3 Network management today

3.1 Rationale for network traffic management

This section applies primarily to mobile networks as specified by 3GPP.

A mobile network will ensure secure, stable and perform an operation compliant to 3GPP network standards; and will be under obligation to meet regulatory requirements. Traffic entering a mobile network is typically managed by a range of network functions to meet these goals.

These functions may include:

- Access control - namely the endpoints allowed to connect to the network and the protocols they can utilise.
- Policy enforcement - including rate-limiting of high throughput flows.

- Network Address Translation
- Application detection - involving header inspection and/or heuristics to determine the Internet service being utilised. This may be used to determine a particular network treatment for a certain flow.
- Content filtering
- TCP optimisation - This function overrides end-to-end TCP congestion dialogue, so that radio conditions are taken into account when deciding whether to increase or decrease segment throughput.
- Video optimisation - This involves either overriding the ABR (Adaptive Bitrate) requests sent from client to server, again to account for radio conditions.
- Image resizing

3.2 Network implementation considerations

As detailed in previous document [WWG 04], operators have in general had to rethink traffic management in the light of the rise in HTTPS traffic. This session layer encryption hides application layer data and metadata from the network, thus preventing intrusive network monitoring such as Deep Packet Inspection (DPI). SMART continues from previous documents published in 2016 in capturing the emerging technologies and techniques which can persist the goals of network traffic management but without breaching customer privacy.

3.2.1 Bearers and QoS

3GPP architecture for Quality of Service (QoS) characteristics [2] describes how a mobile operator may create and configure bearer types to tailor delivery of different traffic types. A bearer is an encapsulation (or 'tunnel') created per user for that user's various traffic flows. 3GPP allows for a default bearer (with no guaranteed bitrate, and 'best effort' QoS, and dedicated bearers with a guarantee bitrate and a higher QCI (Quality of Service Characteristic).

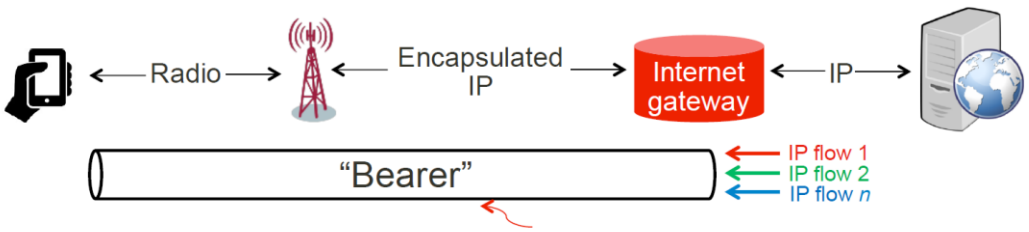

Figure 1. High-level diagram of 3GPP bearer architecture

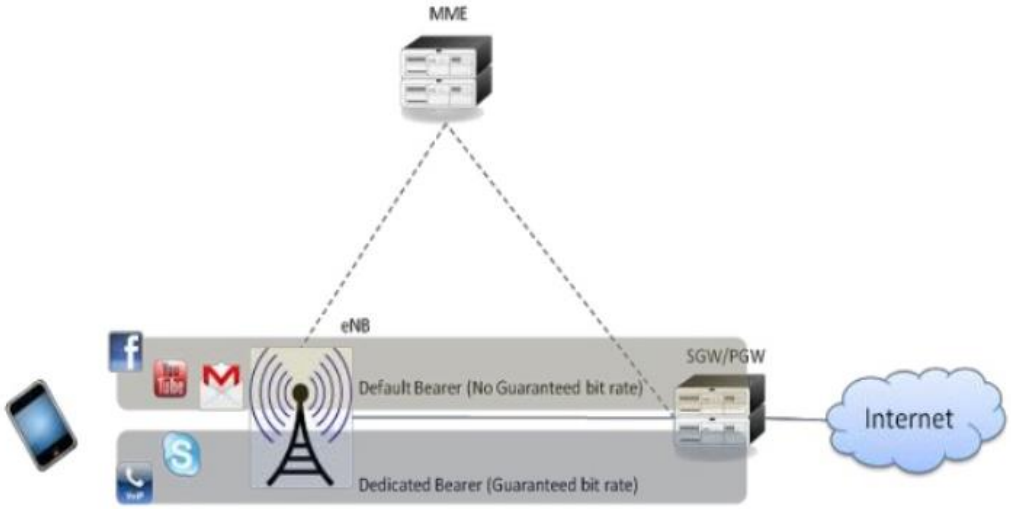

Figure 2. Theoretical example of default and dedicated Internet bearers

Bearer negotiation is performed between the user equipment (UE), such as a smartphone and the APN (Access Point Name) in the core network which maps to an external network – such as the Internet, or an IP Multimedia Subsystem (IMS) core, However, it is important to note that the availability and provisioning of bearers is strictly controlled by operator configuration, and is likely to have an associated operational cost.

A review by the GSMA SMART group in 2015 found that the member operators, across Europe and USA, operated a single default bearer type for all Internet traffic. This means that all Internet traffic, regardless of source or type, is treated equally at the radio scheduler as it decides which information to transmit. Therefore, the traffic detection functions in the operator core network will not result in prioritisation of any traffic flows at the radio access layer. Topics related to Privacy and security are covered in 'Security considerations' below.

3.3 Traffic trends, Q4 2016

3.3.1 Use of encryption

WWG [3], a document published by the GSMA, describes the drivers behind the rise in HTTPS encryption since 2011. The GSMA SMART group has continued to monitor the development of HTTPS (i.e. TLS over TCP) and QUIC; which together will account for less than 50% of traffic transited by a typical mobile network. The success of "Let's Encrypt" [3], which simplifies the certificate acquisition process and is free of charge, and 'HTTPS as a service' offered by cloud providers and CDNs (Content Delivery Networks) will see this trend continue for the foreseeable future. Therefore, network management practices going forward should assume and accept that session-layer encryption is present.

3.3.2 Adaptive Bitrate Streaming

Adaptive Bitrate Streaming (ABR) is a content delivery technique used by major Internet video providers. As the name suggests, the bitrate of the stream is capable of being changed throughout the lifetime of the stream, based on the perceived bandwidth available to the connection. Typically this involves a client estimating the available bandwidth based on the observation of throughput over the previous 5-10 seconds. The client will use this throughput measurement to signal to the server to either raise or lower the bitrate for the

next five second ‘chunk’ of content. The server will attempt to fulfil this request by increasing or lowering its sending rate accordingly, namely by selecting a higher or lower video resolution, in the range of e.g. 240p to 1080p.

Since content providers want to deliver the best customer experience, there is a tendency to:

1. Fill client buffers with high resolution video under perceived good network conditions, on the assumption that users will watch the whole video (or at least a certain amount of video ahead of the current point). This can mean a seamless video experience, but also risks radio resource contention and user data spend based on an assumption of use.
2. Base the estimates of currently available bandwidth on previously observed data. On mobile networks with rapidly-varying available bandwidth (see Another problem: TCP on mobile networks), any estimate of future network conditions based on recent observations is less likely to be accurate. The requested chunk risks being inappropriate to the bandwidth available at the point of delivery.
3. Send the highest available video resolution, regardless of screen size. Higher video resolution takes up significantly more data – 1080p being approx. 2.5x the data size of 480p. However, higher resolutions were designed for television screens, with 480p being the standard resolution introduced for DVD. The differences between resolutions are harder to perceive on far smaller 5-inch smartphone screens, leading to a rapidly-diminishing return on investment for the customer’s data spend. This, along with the greatly increased radio resource contention of 1080p video, has stimulated tariffs such as T-Mobile USA’s ‘Binge On’, placing a 480p resolution limit in return for ‘all you can eat’ video consumption. This can sustain multiple user video streaming sessions with a reduced network load, but without a significant/perceptible decrease in quality of experience. Note that users on larger tablets, or tethering a handset to a large screen over HDMI, will however benefit from HD streaming, hence the user having the final say on preferred resolution is important.

Please see the companion paper, ‘Adaptive Bitrate Streaming on Mobile Networks’ [4] for further details.

3.3.3 QUIC

Quick UDP Internet Connections (QUIC) is an IETF standards-track protocol which brings security, congestion-control and reliability disciplines to UDP. As such, it is an alternative to using TLS/TCP. Claimed improvements are a faster ‘time to data’, since QUIC aims to connect, secure and send content in a single round-trip; session redundancy, as each QUIC connection is uniquely identified and can be re-established following a break or change in connection; and a delay-based pacing algorithm, aiming to reduce jitter and especially reduce streaming video stalling incidents. Early observations indicate a performance improvement over TLS/TCP, especially in poor radio network conditions.

As of January 2017, QUIC is available in the Chrome browser (including Chrome for Android) and google.com servers, and also in the YouTube Android app – which itself can account for 15% increase of traffic by volume for mobile operators.

A key difference for operators is that QUIC, unlike even TLS/TCP, encrypts transport headers (with the exception of the connection identifier and some other public flags). This

includes acknowledgement signals (ACKs), retransmission flags, and segment identifiers. These have historically been used by network operators to monitor health of network nodes and pinpoint any 'problematic' flows – for example a flow which keeps retransmitting. Concerns on the lack of availability of this information are captured in [QUIC-PATH] and are expected to be discussed in the QUIC Working Group at IETF.

3.4 TCP congestion controls in a gigabit Internet

3.4.1 The problem: TCP reaction to perceived congestion

The "multiplicative decrease and additive increase" exhibited by TCP when it perceives congestion has an effect on throughput. Namely, TCP will rapidly back off (reduce its sending rate) and then gently increase the sending rate until it perceives congestion again. This results in a 'sawtooth' pattern of packets sent over time, with the size of the teeth (and hence size of the problem) proportional to the sending rate:

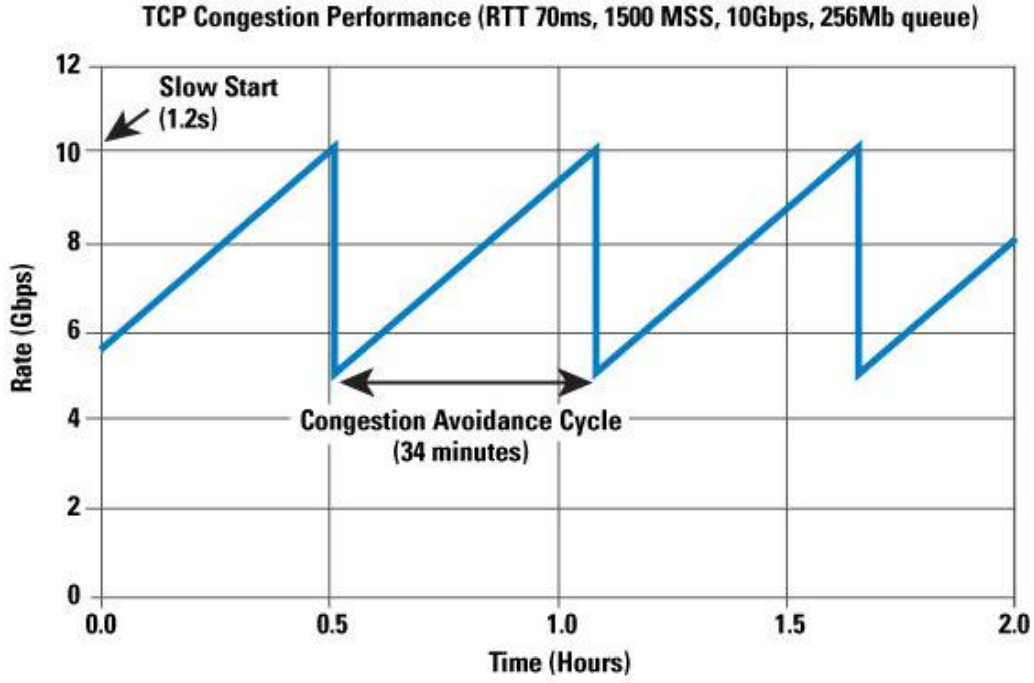

Figure 3. TCP at high speed, "Gigabit TCP', the Internet Protocol Journal Vol 9 no. 2, G. Houston, 2006

This 'cautious probing' also means that TCP incurs several round trips to 'get up to speed', meaning that under ideal network conditions there is a lag as TCP ramps up its sending rate to the available bandwidth. This issue is recognised and being explored at IETF/IRTF.

3.4.2 Another problem: TCP on mobile networks

Recent presentations to (ICCRG) [26] demonstrate issues with TCP over 3GPP mobile networks. The root of the problem is the volatility of the bandwidth available at the radio access layer, due to

- Signal-to-Noise Interference(SNI) or signal fading
- rapidly changing cell load
- handover of buffers during mobility

- device battery state
- device radio state transitions
- etc.

As well as bandwidth variance, this requires a robust polling and retransmission process between the radio network and the device baseband. This is implemented in two complementary ways in LTE:

- fast retransmit of parts of segments believed to have been lost
- full retransmit of full segments.

However, the radio layer is not synchronised with the transport layer: meaning that TCP is not aware of any retransmissions of the segments it has previously sent.

If these retransmissions exceed TCP's retransmission timer, then TCP can therefore wrongly infer packet loss and retransmit segments itself, wastefully adding to the network queues with data that is already queued at the radio layer.

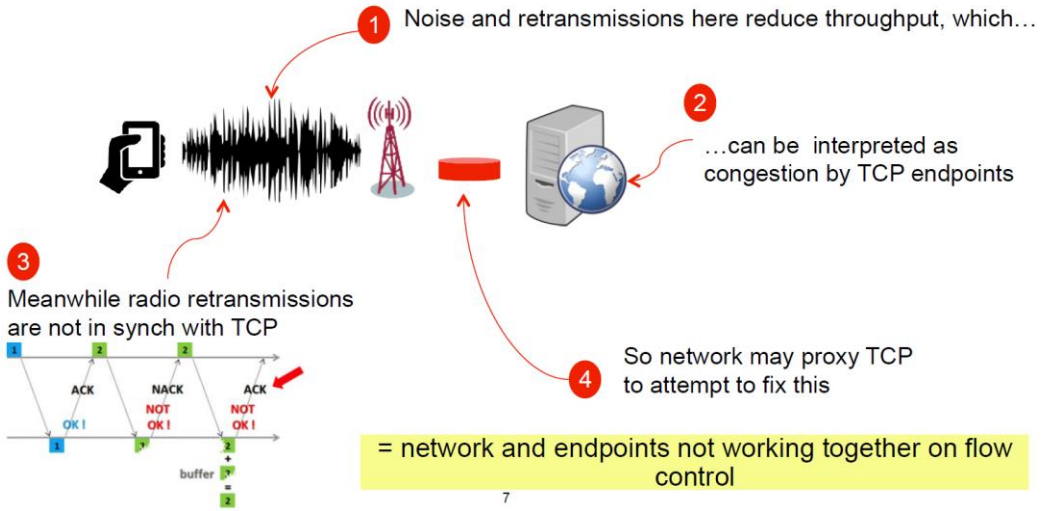

Figure 4. Radio volatility and TCP retransmission

To mitigate this, (CCRG) [27] and (ACCORD) [22] makes recommendations including:

- Packet pacing over traffic bursts
- Choose either high peak throughput or low latency
- Avoiding 'loss'-based congestion controls. Alternatives include TCP-RACK (time-based) and Google's BBR (congestion-based) [5]
- Closer interworking between cellular and Internet standard bodies to solve the problem.

Meanwhile network operators are known to utilise TCP proxies (aka Split TCP or TCP optimisers). These decouple the client-server connection into two TCP connections: client-proxy and server-proxy. TCP proxies typically introduce a bespoke congestion control algorithm which may be tailored either to common cell and core network conditions for that network, or in some cases integrate with the radio layer to react rapidly to changing radio conditions (as per Mobile Throughput Guidance).

4 Trends in Internet stack evolution, and the impact on networks

This section describes some significant design trends in Internet standards and associated technology, including how they relate to network traffic management.

4.1 Encryption by design

The Internet Architecture Board (IAB) published a 'Statement on Internet Confidentiality' [6] in November 2014, with the following key recommendations:

- Newly designed protocols should prefer encryption to cleartext operation
- Encryption be deployed throughout the protocol stack

This includes all Internet protocols: from consumer Internet traffic served over HTTPS to Internet of Things standards such as Constrained Application Protocol (CoAP).

This decision was driven by both privacy (in light of pervasive monitoring revelations) and as a means to allow new protocols to pass through network middleboxes which may otherwise not recognise (and hence reject) them, or drop key packet headers.

4.2 Impact of network middleboxes on transport protocol design

4.2.1 Network middleboxes: overview and impact

A "network middlebox" is any interim node between client and server that falls within the network operator domain. Typically, this is not used to mean Internet routing functions (such as Autonomous System routers) nor physical or link layer nodes: rather it implies a network function that inspects and potentially modifies information communicated between client and server. These can include:

- access controls for Internet flow (firewalls),
- attached client endpoint mapping (Network and Port Address Translation),
- traffic and application detection functions (inspection of 5-tuple for categorisation),
- traffic shaping (rate limiters and load balancers)
- content filters (inspection of 5-tuple, and request URL where available),
- malware filters (inspection of 5-tuple, request URL where available, deep packet inspection to parse content payload where available),
- TCP optimisers (split TCP that attempts to account for volatile radio conditions, by manipulation of ACKs and retransmission timers)

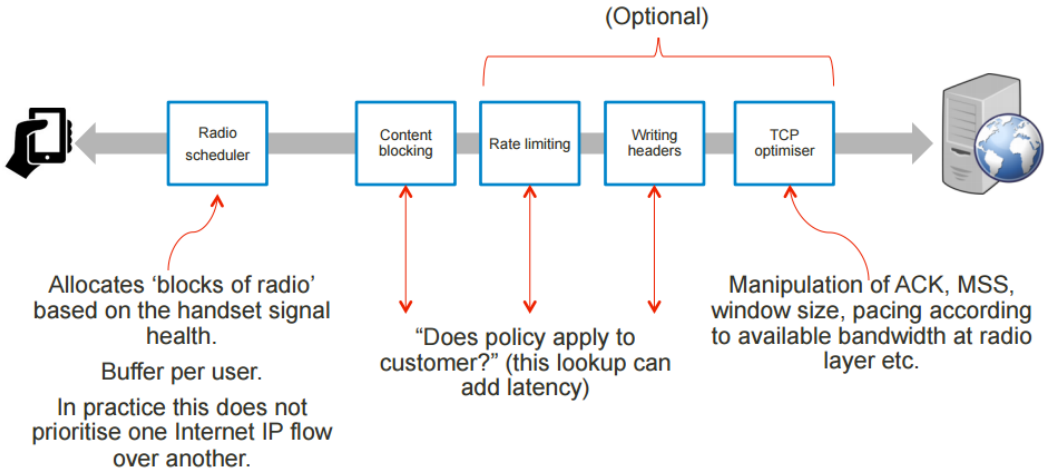

Figure 5. Illustrative example of network middleboxes (from ACCORD BoF, IETF 95)

The impact of these functions has been discussed at one of the IAB workshop [7] and the GSMA – IAB workshop [8], with a general conclusion that middleboxes can slow down adoptions or maybe in some cases hinder the deployment of new or evolved transport protocols – typically this would be due to either the new protocol being unrecognised at access control, or new headers being dropped at routers. The resulting 'ossification' (making rigid that which should be flexible) has led to a difference between what the endpoints expect, and what the networks can deliver – exacerbated by the fact that networks present a 'black box' to endpoints:

"It might be tempting to single out middleboxes as the root of the Internet's ossification problem, but their existence is not a fundamental show stopper. Rather, the issue lies in the fact that their presence in a path is often unknown to the end points, that some of them may modify or alter protocols in ways unexpected to end points, and that there is no mechanism to explicitly address them in order to negotiate and resolve the tussle between what the operator needs in order to ensure the correct functioning of its network and what end users would expect in terms of the service they would like to experience." – quote from SEMi workshop [20]

The measurement of the impact of network middleboxes on transport protocols is underway at IETF research projects [9], with input from the former GSMA Protocol Optimisation Project. The project will also identify areas where protocols may themselves allow information transfer between endpoints and middleboxes – a challenge being to ensure that any such information, whether in isolation or in a broader context, does not breach users' privacy or security.

4.2.2 Getting new protocols across network middleboxes

Content providers hindered by a perceived inability to deploy new, improved protocols over access network middleboxes, and have hence have utilised "network-approved" protocols as a carrier layer. This has been referred to as 'the Narrow Waist Model of the Internet', [10] wherein the Internet Protocol is difficult to extend, leaving higher layers as the area of innovation.

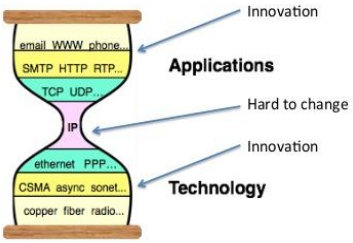

Figure 6. The Narrow Waist Model of the Internet

Well-deployed recent examples have been:

- SPDY utilising session-layer encryption to tunnel the new application protocol through the network, using the recognised port 443. This was widely deployed by Google, Facebook, Twitter and Yahoo prior to standardisation as HTTP/2.
- QUIC over UDP, using UDP to traverse both Network Address Translations (NAT) and firewalls. At the time of writing this is used by Google on google.com and youtube.com servers towards Chrome browsers and Android YouTube apps, and the QUIC working group and its charter [33] has been approved by the IETF.

The ability to traverse middleboxes is helped by host-to host encryption: deployment of improved transport/application protocols has been one of the drivers for significantly increased Internet encryption in the last five years. By utilising Transport Layer Security (TLS) with well-known ports familiar to firewalls, the new protocol is encrypted within the TLS layer and not typically subject to additional access control checks (i.e. the new protocol identifier is not typically parsed to determine access). QUIC applies encryption by default and utilises UDP as a protocol known to NATs and firewalls, which has led to a rapid and successful penetration across live networks:

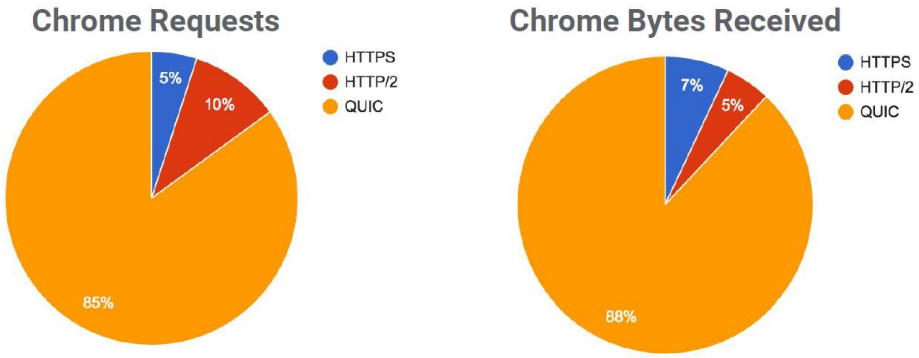

Figure 7. QUIC penetration across live networks, IETF 96 June 2016

Finally, there have been proposals to utilise bytes within the protocol payload as a means to extend the protocol [Inner-space]. In theory, this will protect the extensions from being dropped or modified as they are in an area that remains untouched by middleboxes that only read headers. Inner-space extensions would only be hindered by a middlebox that performed Deep Packet Inspection and was specifically looking for such extensions, for which there would be little benefit and increased network cost. If Inner-space extensions are delivered via an encrypted channel (such as TLS) then the likelihood of middleboxes

dropping or modifying any extensions diminishes further. The only drawback is the reduction in payload space caused by extensions, hence making efficient extension modelling advisable.

4.2.3 Substrates and Signalling

The ability to provide signals toward the network, or from the network towards endpoints, can allow flows to be treated according to network state. This state can relate directly to the flow, such as preferred delivery characteristics (for example preferring drop over queue applicable to the flow, as per (1-bit) [22], or it can relate to the network state as a whole (for example, congestion exposure (ConEx) [25] or Mobile Throughput Guidance (MTG) [26]. In-band signalling will require signals to be passed unmodified to the intended node where they will be processed. This is non-trivial, as there is no 'control plane' for signalling flow metadata in the Internet stack: ICMP (Internet Control Message Protocol) operates at the Internet layer and is concerned with relaying IP connection errors between routers; and other header space (TCP Options, IPv6 header extensions) may be dropped by routers in order to reduce latency. Further details can be found at PLUS (Plus Layer UDP Substrate) [21].

5 Smarter Traffic Management

Smarter Traffic Management is the range of techniques that allow a network to optimise data flows, without access to those flows' private data or metadata. 'Optimise' can be read in the context of a single flow, the combined flows for a single client, or the total flows across all customers' on the network path.

5.1 Network management of encrypted traffic

Networks utilise various management techniques to ensure efficient throughput, congestion management, anti-SPAM and security measures. Historically these functions have utilised visibility of the Internet application layer, including Shallow (headers) and Deep Packet (payload) inspection.

This visibility is rapidly diminishing - encrypted Internet traffic is expected to continue its upward trend, driven by increased privacy awareness, uptake by popular services, and advocacy from the IAB [11] and Technical Architecture Group at the World Wide Web Consortium (W3C) [24].

The IAB produced RFC7258 technical specification and IETF Security Group [29] recognise that network management functions may be impacted by encryption, and that solutions to persist these management functions must not threaten user security or privacy. Such solutions can ensure the benefits of encryption do not degrade network efficiency. IAB observes "many network operations activities today, from traffic management and intrusion detection to spam prevention and policy enforcement, assume access to cleartext payload" and promises to "work with those affected to foster development of new approaches for these activities which allow us to move to an Internet where traffic is confidential by default." The joint GSMA/IAB/IETF activities have helped acknowledgement of the need to support both user privacy and non-intrusive network management.

The catalogue of techniques that can assist with network management of encrypted traffic, whilst protecting user privacy and security, is maintained at [smith-encrypt], the IETF Internet

draft reflecting input from [23]. This document evolves as various technologies mature and new proposals introduced.

5.2 Security considerations

5.2.1 Vulnerabilities of network Man in the Middle certificate interception

The GSMA ENCRY sub-group (the precursor to the SMART sub-group) recommended that operators do not apply 'Man-in-the-Middle' (MITM) interception to encrypted Internet traffic. The reasons for this are to protect the customer, the network and the content provider from security risks and poor quality of experience.

Man-in-the-Middle interception of TLS involves the initial handshake being spoofed by an intermediary: the result is that instead of a client-server secure connection, there are two 'back-to-back' connections between client-intermediary and intermediary-server. The intermediary therefore has cleartext access to the headers and content payload, before re-encrypting the contents to forward to the receiver. The interceptor needs to spoof the client (and less commonly, the server) into believing that the certificate has come from the origin server.

Browsers and OS certificate validation work to authenticate that the certificate is genuine, including:

Denial of access:

- Certificate Pinning (restricting which certificates are authenticated),
- Certificate Transparency (a cloud repository of rogue certificates),

Reduced quality of experience

- Visual warnings to the end user (which may also inure the user to warnings of other threats)

In addition, the process of introducing a MITM certificate also introduces security vulnerabilities:

- The user must install the certificate themselves. A rogue actor may use this as an attack vector, for example a phishing mail asking the user to 'update their certificate' when in fact it is a new, rogue, certificate allowing access to the user's private data.
- Any operator interception proxy itself becomes a high-risk attack vector, as breaching the proxy or the certificate can result in all encrypted user data being compromised.
- Downgrade of trust, where the interception proxy encrypts with lower security than the origin server and client had intended to use.
- Inappropriate upgrade of trust, where the interception proxy 'upgrades' a connection which would otherwise have used an expired origin server certificate and raised an alarm at the client.

5.2.2 Let's Encrypt and malware mitigation

Let's Encrypt makes HTTPS both free and simple, and as such has been very successful in enabling encryption for domains which may otherwise struggle to afford or configure it. This has, however, meant that malware actors – who may have previously been put off using

HTTPS due to both the expense and registration/payment that may identify them – can rapidly generate free certificates. If that certificate is recognised as that of a malware actor, then the actor can simply generate another one, also for free.

Let's Encrypt are aware of this issue and utilise Google's Safe Browsing API to help filter requests from known malware providers, however cases have been observed where malware actors have posed as genuine domains to circumvent this check. Operators should be aware that while HTTPS can validate that the connection is to the stated origin server, that does not in itself mean the origin server should be trusted as a good actor by the user. Any compromise of ad-servers resulting in malware being served to otherwise trusted sites is a concern that should be tracked by operators, as the only network mitigation would be to block the (otherwise benign) ad-server domain as a whole.

5.3 Collecting the SMART solutions

This section describes the basis for collecting suitable Smarter Traffic Management solutions for consideration.

5.3.1 Guidelines

1. Solutions must not breach customer privacy. Privacy is difficult to quantify, since information available for routing (IP source and destination) and other metadata (timings, packet sizes etc.) can *in totem* indicate a certain kind of Internet communication from a known source towards a known person. The solutions we consider are not intended to introduce any further metadata towards such a goal: rather they aim to use routing, link and transport metadata below the private user session.
2. Solutions will not favour one customer's traffic over another or one content provider's traffic over another. The ability for a customer to prioritise between their own traffic flows is valid, as long as this does not affect other customers' traffic flows.

5.3.2 Recommendations from Internet community collaboration

Smarter traffic management is not solely an operator responsibility but also an opportunity for content providers. The IAB MaRNEW workshop proved hugely important in allowing the requirements and tensions of various actors to be discussed: including content providers, OS and equipment vendors, network operators, security experts and privacy advocates.

Key recommendations towards the mobile network operators were to:

1. Measure the impact of encryption at the radio scheduler – since all operators in the group reported that they only use the default bearer ('best effort QoS') and hence no impact could be measured, i.e. transport layer encryption has no impact on scheduling.
2. Look at 'zero-bit' optimisations that require no transfer of metadata between the network and endpoints. These include Active Queue Management solutions, designed to efficiently drain network buffers to increase overall throughput.
3. Then consider 'one bit' optimisations, e.g. flagging packets as 'drop/queue' in case of congestion.

The GSMA Web Working Group (including SMART members) continues its engagement at IETF meetings, including the MARNEW follow-up ACCORD at IETF 96. Continued

monitoring and contribution to emerging standards (QUIC, TLS 1.3, CAPPORT etc.) helps ensure early impact awareness and solution design for operators

5.3.3 Measurement of protocols and network traffic management functions

Smarter Traffic Management can include the removal of existing network functions, since the rationale for their original introduction may have been overcome by improvements in Internet protocols or other technologies. Therefore, capturing the range of today's network functions, and their effect, provides an important basis to understanding the impact of network transport on Internet protocols: specifically, the impact on the protocol behaviour between client and server. This effort is managed within the GSMA by the Protocol Optimisation Project, in close liaison with the IETF MAP research group.

5.3.4 Protocol Optimisation Project

The GSMA Protocol Optimisation Project aims to analyse how new and existing Internet protocols (including TCP and QUIC/UDP) may be best delivered over operator networks. The initial work includes the testing of two hypotheses, around how QUIC flows co-exist with TCP flows on the same core and radio networks; and whether the flagging of flows with intended network treatment result in an improved customer quality of experience. Further information can be obtained through the GSMA Internet group.

5.4 Solution analysis

Per the IAB MaRNEW feedback, the table breaks up solutions into 'zero', 'one' or 'multi-bit' optimisations:

Category	Current examples
Zero-bit optimisations:	1. Low latency, low loss scalable throughput: DualQ AQM/DTTCP 2. Availability of H2C: HTTP/2 without encryption as Apache module
One-bit optimisations:	1. Delivery context sharing: Mobile Throughput Guidance, PLUS (formerly SPUD) 2. Assisting transport layer congestion controls: ConEx 3. Explicit flagging towards network: 1-bit loss/latency trade-off [22]
Multi-bit optimisations	1. Mobile throughout guidance 2. PLUS

The following analysis continues, and may result in technical solutions that fit into the table above:

4. Impact assessment of ABR streaming and how to reduce wasted bytes for the customer
5. Assessment of continued value of network functions: analysing output of MAP (formerly HOPS [10]) and the Protocol Optimisation Project.

5.4.1 Zero-bit optimisations

5.4.1.1 Active queue management and L4S [28]

The IETF Active Queue Management and Packet Scheduling WG [12] works on algorithms to manage network queues, with the aim of reducing packet delay and taming aggressive/misbehaving flows. This includes allowing flow sources to control their sending rates to avoid unnecessary losses.

Whilst cutting-edge “AQMs such as PIE [31] and fq_CoDel [31] give a significant reduction in queuing delay relative to no AQM at all”, and should hence be considered for evaluation by network operators, ‘“ without addressing the large sawtoothing rate variations of Classic congestion controls, they cannot reduce queuing delay too far without significantly reducing link utilization”.

A solution proposed to IETF, Low Latency Low Loss Scalable throughput, (L4S), is of particular interest to Smarter Traffic Management. This is because it operates at the IP-layer, using standard Explicit Congestion Notification(ECN) flags, hence not requiring any shallow or deep packet inspection (unlike other AQM systems).

L4S aims to make ‘finer’ sawteeth: that is, reduce the huge variance in sending rate of TCP caused by TCP’s implicit perception of network packet drops as congestion. Benefits may be increased on mobile due to TCP’s perception of radio access volatility as loss. These ‘finer sawteeth’ can mean better network utilisation. Where ECN marking is not present, L4S will create two distinct queues: one ECN capable, one ‘classic TCP’. An algorithm then calculates which of these to process further (keep the ECN marking) and which to drop. The claimed result is low latency and low loss throughput. Lab results presented at IETF 93 and IETF 96 ACCORD meetings show benefits; highlighting the potential benefits across the different flow types from 5G use cases:

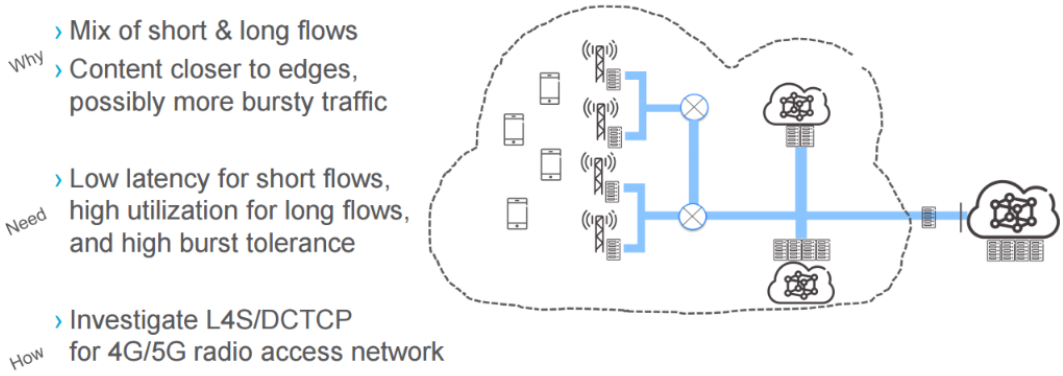

Figure 8. Motivation for L4S, I. Johansson, IETF 96 ACCORD meeting

Latency/bandwidth demands for 5G flows will span from sporadic, long-lived, low-bandwidth connections for IoT; to 4K and VR/AR flows requiring very high capacity and mobility. L4S claims to result in better buffer utilisation: because buffers are drained quickly, they are

better able to accommodate any bursty traffic. Further testing by mobile operators is recommended.

5.4.1.2 Improved congestion controls

TCP congestion controls in a gigabit Internet outlines how multiplicative decrease/additive increase can result in an undesirable sawtooth pattern of bitrate, and how the volatility of cellular radio access can result in an incorrect perception of packet loss.

A range of TCP congestion controls aims to solve these problems, and hence reduce bufferbloat, spurious retransmissions and sawtooth jitter.

These include:

- TCP RACK [13], which is a time-based loss detection. The algorithm is based on the premise that ‘Packet A is lost if some packet B is sufficiently later s/ACKed’. In other words, if the sender receives an ACK for packet B, but is still waiting for an ACK for packet A, then it may assume packet A was lost – allowing for any out-of-order ACK signalling by the receiver, by configuring what is meant by ‘sufficiently later’. The packet transmit times are monitored to make the RACK system aware of network performance. This means RACK is accounting for the average transmission time and the ACK pacing when inferring loss. However, on cellular networks, there still remains the decoupling of layer 1 and 2 retransmissions and buffer handover during mobility, which is not explicitly accounted for in RACK.
- TCP PRAGUE [14], based on the Data Center TCP (DCTCP) approach utilised in L4S. DCTCP is tuned for private networks with known traffic patterns, trusted senders and receivers, and a stable (or scalable) capacity. This results in a resource-intensive protocol that is considered too aggressive for the open Internet. However, parts of the congestion controls are under consideration as a means to flatten the typical TCP ‘sawtooth’ sending rate, and also to allow TCP to get up to speed quickly (to utilise available bandwidth and reduce starting latency) without several probing round trips.
- Scavenger protocols such as LEDBAT (RFC6817) [15] aim to allow utilisation of network resources without impacting network queues. LEDBAT is experimental, and advises ‘LEDBAT's responsiveness and throughput should be evaluated in the wide area and under conditions where abrupt changes in base delay might occur, such as with route changes and with cellular handovers.’ – and also the cellular network air interface in general exhibits rapid changes in base delay, as noted previously.
- Google’s [BBR], ‘Bufferbloat Reduction’, is available in Linux Kernels as of Q4 2016. It takes the approach of tracking bandwidth and delay independently, and assessing the variance in the pacing rate over time. A basic indication of how this can reduce the TCP sawtooth is available here.

This document’s recommendation is that operators allow the testing of new, non-inferred-loss based congestion control algorithms in mobile scenarios (including handover), especially as comparison with QUIC, which also moves away from the loss-based approach. Benchmarks against any existing TCP optimisers will be particularly valuable.

5.4.2 One-bit optimisations

DSCP (DiffServ CodePoints) [16] were intended as a means for content providers to flag IP headers as requiring a certain class of treatment by the operator network: for example, to

expedite or otherwise prioritise traffic. The DSCP value was added by the content server, but in general this value was ignored, or 'bleached', by the operator: the markings were too coarse, and there was no reason for a content provider to mark their traffic as anything other than 'priority', for fear of their traffic being treated as a lower class than their competitors.

One approach to reduce the motivation of flagging everything as 'priority', and without reducing customer privacy, is to allow the content server to rather state what should happen to a packet or flow at the point of resource contention. In other words, if the network cannot immediately forward the packet, a packet flag will indicate 'drop me' or 'queue me'. The choice will depend on the nature of the traffic, namely the preferred balance between low-latency and resilience. A real-time video call, with the human parties' ability to fill in any gaps caused by loss, will have a different requirement than a long-lived download, for example. This scenario is explored at IETF 95 meeting.

3GPP networks also support traffic classes, which are determined by analysing traffic data via deep packet inspection (which is rapidly diminishing due to HTTPS and privacy regulation), shallow packet inspection (such as IP 5-tuple) or heuristics (traffic rate and pattern estimation).

Note, that in both IETF and 3GPP cases, the ability of the radio scheduler to fulfil any queuing instructions in the radio access network is described in Bearers and QoS

5.4.3 Multi-bit optimisations

5.4.3.1 Path Layer UDP Substrate (PLUS)

A method to allow signalling without router upgrade has been proposed at IETF: PLUS, Path Layer UDP Substrate (formerly known as SPUD) encapsulates the flow within UDP. The outer UDP allows authenticated, encrypted signals to be relayed between endpoints, and allows the network to populate a 'scratchpad' area of the headers with network information useful to the endpoint. The intention of the proponents is to move from today's implicit network functions to network functions that are explicitly signalled to the endpoints, hence allowing them to account for middleboxes (utilisation or possibly bypass).

At time of writing (pre-IETF 97) the proposal received a mixed reception [PLUS-minutes]. There was recognition that mobile networks do not work well with TCP flow controls, however there were concerns that:

1. PLUS could be impeded performance, due to additional header overhead (early deigns have an estimated 12 additional bytes) for transport and parsing.
2. Scope cannot feasibly be locked to purely transport semantics. There were concerns that the network could abuse the scratchpad: for example, by setting conditions for access.
3. The proposal is also governed by the endpoints: a network cannot create a scratchpad to immediately signal a change in throughput guidance or congestion, which would reduce the value of using PLUS for those cases.

The SMART group will continue to monitor and contribute to the PLUS work.

5.4.3.2 Mobile Throughput Guidance

Mobile Throughput Guidance (MTG) allows the network to provide a hint to content servers: “here is the suggested bandwidth to target for this flow”. This is driven by two factors that the content server typically guesses at: the fluctuating radio signal quality for the client device, and the congestion state of the radio and core networks. Together these result in volatile bandwidth. Another issue is that the ‘probing’ nature of TCP means that a flow will take several round-trips to utilise available bandwidth: it simply does not ‘get up to speed’ quickly enough.

MTG can therefore help achieve a sustainable bandwidth target for a flow, and ensure more rapid use of available bandwidth.

5.4.4 User-based prioritisation

This approach aims to have the user state which of their flows should be prioritised. A user’s interpretation of ‘prioritised’ is likely to be less granular than a network, and can be abstracted as ‘the priority service should achieve best user experience. Non-priority services can degrade if necessary’. A core use case is that if the customer has a jitter-sensitive service then it can be prioritised over non-jitter-sensitive service: such as prioritising a live TV stream over app updates or application downloads.

A non-standard proposal was presented at SIGCOMM, August 2016 [17]. This allows a customer to prioritise among *their* traffic, but not at the expense of others. For example, if the customer is streaming a video and downloading a file concurrently, they can ask that the streaming video be given preference, and the download may as a result take longer. The download would basically occur in the gaps between video stream requests.

This is similar to the DiffServ approach, but is made user -centric. DiffServ is service-provider centric, and in practice has risked all services simply flagging their traffic as “priority” in an attempt to achieve best quality of service.

As the user is prioritising *their* traffic among itself, it does not result in the user getting a ‘fatter pipe’ at any point from the network. The network will allocate bandwidth and scheduling at best effort among users without discrimination. So, it won’t allow a preference for 4K video to harm neighbouring users, it just means the 4K gets priority over downloads (for example).

However, the document utilises elements which have run into deployment difficulties. Adding flags to IPv6 headers (which routers often drop, and placement is still debated) or TLS extensions (which may require browser and OS network stack support) are seriously non-trivial. Authentication and security will need to be added, as well as user state mappings and data model upgrades for the OSS/BSS. And the 3G/LTE radio schedulers do not today distinguish among Internet traffic, they just distinguish between IMS/VoLTE and Internet.

The general area of ‘user defined priorities’ remains valid, and may fit into some of the signalling architectures (such as PLUS). The integrity of any such flag, to ensure it has been set by the user and not by a client service unknown to the user, is important to determine. Whilst it may seem a valid approach for user priorities to be stored in the network, for example by interacting with an operator-provided app, the problem of operators identifying the nature of flows so that they can be prioritised remains.

6 2020 outlook for network management

The Vision 2020 White Paper from the GSMA [29] Intelligence analyses the scenarios that will drive a change in network capability and operation c. 2020, including convergent networks, an explosion in connectivity and mobility through Internet of Things uptake, and the bringing together of physical and virtual worlds through augmented reality over low-latency edge computing.

In terms of traffic management, operators are particularly encouraged to track and contribute to developments in the following areas

6.1 Regulation

The following topics are discussed in the context of technical impact, and should not be read as regulatory guidance. Readers are directed to their regulatory contacts, and GSMA CROG for further discussion.

6.1.1 Privacy

By 2020, the majority of Internet traffic will be encrypted at the session layer, based on the current rate of adoption. Operators will need to ensure that their technical realisation of traffic identification, for example to implement network slicing, does not breach data regulations, such as the General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679), whilst also meeting any country specific obligations (such as the UK Investigatory Powers Bill).

6.1.2 Net Neutrality

Network traffic management should not in general discriminate between customers or content providers. Rather, Technical criteria, such as the characteristics of a traffic flow, or the radio state of a client device, should be considered as the basis of management policies. An exception may be certain emergency use cases where defined by regulators, allowing priority over non-emergency traffic.

6.2 5G

Whilst ‘5G’ refers to the evolved radio access layer beyond LTE-Advanced, two network operation themes are likely to receive deployment in the same period: network slicing and NFV/SDN.

6.2.1 NFV and SDN

The provisioning or scaling of a network function should not require the purchase of dedicated telecoms hardware, or an integration of that hardware *in situ* to the operator’s network. Instead, new functions, capacity, storage and processing power can be rapidly provisioned in a Cloud, running a standard virtual machine infrastructure.

The benefit to industry and consumers is that the network can more rapidly handle shifts in data usage – including the ability to deploy short-term infrastructure to support a business or service that only requires a short-term enhancement. The elastic nature of the Cloud allows the infrastructure to shrink again when no longer required.

This model is named Network Function Virtualisation (NFV).

Networks implement a range of networking protocols - rules that govern the structure and behaviour of packet data systems. By abstracting the control of this range of protocols to a simpler management layer, it is possible to manage the network as a whole. This is known as Software Defined Networking (SDN), since control of the network is defined at a software level rather than through hardware routers, switches etc.

SDN and NFV together make a more flexible, reactive network: this can be tailored towards various industry and consumer sectors through ‘Network Slicing’.

6.2.2 Network slicing

Network slicing allows the allocation of network resources appropriate to the service being consumed and its particular delivery characteristics. These characteristics include:

1. Latency - how quickly packets are delivered
2. Pacing - the consistency of delivery rate, ensuring packets arrive in order
3. Durability - the ability for a connection to be long-lived, or dormant and then woken over a multi-year period
4. Resilience - guarantees that packets are delivered within a certain timescale
5. Capacity - the ability to deliver high volumes of packet data to a service
6. Mobility - services persist seamlessly as devices transit access points
7. Coverage - services are supported across a wide geographic area

By using NFV and SDN, this allocation does not require the adding of new hardware each time a service is brought online, but can rather be rapidly set-up and later discarded as required. Thus, a ‘slice’ provides the layers described above, as separate logical networks.

Network slicing can therefore enable *no discrimination between content providers or customers* - a flow-centric, service-provider-agnostic approach.

6.3 Internet of Things (IoT)

The network traffic management demands of massive IoT deployments will essentially depend upon the nature of the traffic flows. Sensors that are awoken every six months to deliver a short, but critical, status report will have different demands to drone-carried cameras streaming 4K video over short range/high capacity small cells. Network Slicing as described above should allow contention for network resources to be managed across the wide range of IoT use cases.

6.4 Mobile Edge Computing

Edge computing moves computation and storage functions close to the network edge, to:

1. Reduce ‘distance to content’ and hence latency
2. Remove any Internet hops to retrieve content, and hence reduce backhaul costs.

Use cases include augmented reality, as the ‘motion-to-photon’ demands of human optical overlays require very low latency for a satisfactory quality of experience; live video stream analysis (analysing video on the uplink); rapid location lookup; and radio network information to enhance service delivery (along the lines of Mobile Throughput Guidance).

In terms of network management, the implication is that the Radio Access Network will need to host certain management functions (caches, video optimisers) etc. which have typically resided further back in the operator core – for example, on (S)Gi LAN. These functions will be specified by ETSI/3GPP and will operate solely on the MEC traffic.

6.5 Secure caching

Content cached in the operator network can bring two benefits: the customer receives the content quicker, and the operator saves on Internet backhaul. Encrypted content, however, cannot be cached. A technique, variously known as ‘out of band content encoding’ or ‘blind caching’ has been proposed for caching encrypted content. Additionally, there are other techniques such as Use of Short-Term Automatically-Renewed (STAR) [32] Certificates that allow delegations of certificates. The details of those techniques are captured in Secure Content Delegation white paper.

In brief, the out of band process is based on:

1. Encrypted content is cached at the operator network
2. The content provider can signal the presence of this ‘nearest server’ to a client connected to that operator’s network, along with a means to decrypt the content.
3. The client requests the content from the operator cache and decrypts it.

This can therefore balance the need for user privacy and security with the benefit caching brings to operators. See the associated WWG paper, ‘Secure Content Distribution’, for further details.

7 Beyond 2020

7.1 ETSI NGP

The 2020s will very likely see an explosion in network and internet connections (driven by IoT), latency and throughput demands (virtual reality, augmented reality, 4K and 8K video), mobility demands (automotive and high-speed train use cases), multi-homing and convergence (across cellular, Wi-Fi, fixed, Narrowband-IoT and non-3GPP networks).

ETSI’s Next Generation Protocol interest group [18] is concerned with ensuring that networking/internetworking protocols will able to meet these scenarios, and has published a white paper explaining that the existing IP stack will not be able to meet these demands, due to inherent constraints. In addition, network security and mobility will become increasingly costly to operate as overlays top the existing IP stack.

ETSI NGP is therefore performing 3 tasks:

1. Capturing the scenarios and gap analysis against the existing IP stack.
2. Collecting requirements that next generation protocols will need to meet to deliver the scenarios.
3. Researching non-IP architectures (meaning ‘styles of building’) to SDOs, including 3GPP and IETF, so that the next generation protocols can be developed and specified.

Traffic management considerations for the recommended architectures include flow controls (including a distinct control and data plane) and access controls. As an early example under consideration please see [19].

8 Conclusions

Network traffic management over the next 5-10 years will be dominated by several themes, both established and emerging. Encryption of traffic will be a given constraint for traffic management, and operators should account for traffic management solutions that improve customer experience without breaching customer privacy. These will include heuristic approaches that avoid deep packet inspection, as well as explicit signalling of desired packet treatment from content servers. Such signalling allows the network to play a defined and predictable role in optimised traffic delivery.

Traffic management will shift from policies based on source (for example any treatment of a particular provider) to management based on traffic characteristics. This will allow the vast range of connections and traffic types predicted for the 5G era to receive the appropriate network treatment based on their requirements, whilst allowing networks to adhere to regulations concerning non-discrimination. Network slicing will allow the configuration of these characteristic-based policies, scaling and shrinking according to demand. Any operational constraint of a single Internet default bearer, delivering ‘best effort’ with no guaranteed bitrate, will however leave the radio scheduler as a potential bottleneck unable to distinguish between traffic with different characteristics. A ‘slice aware’ scheduler will therefore be required to allow end-to-end traffic management.

Adoption of cutting-edge congestion controls – whether from endpoints, or advanced queue management developments – should assist both the efficient usage of network capacity and improved customer experience. These may cause operators to rethink which middlebox functions are still applicable, as Internet technologies evolve to solve the problems which led to middlebox deployment.

Finally, the predicted explosion in connections, throughput and mobility in the next ten years will strain existing IP-based infrastructure; and operators should engage in investigation of scalable, sustainable alternatives in order to meet the performance and cost challenges of the 5G-era.

Annex A Document Management

A.1 Document History

Version	Date	Brief Description of Change	Approval Authority	Editor / Company
0.1	06/02/2017	New Whitepaper	PSG/PSMC	Kevin Smith (Vodafone) Istvan Lajtos (GSMA)
1.0	15/02/2017	Editorial changes	PSG/PSMC	Sanjay Mishra (Verizon)

A.2 Other Information

Type	Description
Document Owner	Internet Group
Editor / Company	Kevin Smith (Vodafone)

It is our intention to provide a quality product for your use. If you find any errors or omissions, please contact us with your comments. You may notify us at prd@gsma.com