

IAB Workshop on Managing Radio Networks in an Encrypted World (MaRNEW)
Report
draft-nrooney-marnew-report-00

Abstract

The MaRNEW workshop aimed to discuss solutions for bandwidth optimisation on mobile networks for encrypted content, as current solutions rely on unencrypted content which is not indicative of the security needs of today's internet users. The workshop gathered IETF attendees, IAB members and various organisations involved in the telecommunications industry including original equipment manufacturers and mobile network operators.

The group discussed the current internet encryption trends and deployment issues identified within the IETF, and the privacy needs of users which should be adhered. Solutions designed around sharing data from the network to the endpoints and vice versa were then discussed as well as analysing whether the current issues experienced on the transport layer are also playing a role here. Content providers and CDNs gave an honest view of their experiences delivery content with mobile network operators. Finally, technical responses to regulation was discussed to help the regulated industries relay the issues of impossible to implement or bad-for-privacy technologies back to regulators.

A group of suggested solutions were devised which will be discussed in various IETF groups moving forward.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 11, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

Mobile networks have a set of requirements and properties which places a large emphasis on sophisticated bandwidth optimization. Encryption is increasing on the internet which is positive for consumer and business privacy and security. Many existing mobile bandwidth optimization solutions primarily operate on non-encrypted communications; this can lead to performance issues being amplified on mobile networks. Encryption on networks will continue to increase; and with this understanding the workshop aimed to understand how we can solve the issues of bandwidth optimization and performance on radio networks in this encrypted world.

1.1. Understanding "Bandwidth Optimization"

For the purposes of this workshop, bandwidth optimization encompasses a variety of technical topics related to traffic engineering, prioritisation, optimisation, efficiency enhancements, as well as user-related topics such as specific subscription or billing models. These can include:

- o Caching
- o Prioritisation of interactive traffic over background traffic
- o Per-user bandwidth limit
- o Business-related topics such as content delivery arrangements with specific content providers.

Many of these functions can continue as they're performed today, even with more encryption. Others use methods which require them to inspect parts of the communication that are encrypted, and these will have to be done differently in an encrypted Internet.

Finally, while not strictly speaking traffic management, some networks employ policy-based filtering (e.g., requested parental controls) and all networks support some form of legal interception functionality per applicable laws.

1.2. Topics

The workshop aimed to answer questions including:

- o Understanding the bandwidth optimization use cases particular to radio networks
- o Understanding existing approaches and how these do not work with encrypted traffic
- o Understanding reasons why the Internet has not standardised support for LI and why mobile networks have
- o Determining how to match traffic types with bandwidth optimization methods
- o Discussing minimal information to be shared to manage networks but ensure user security and privacy
- o Developing new bandwidth optimization techniques and protocols within these new constraints
- o Discussing the appropriate network layer(s) for each management function
- o Cooperative methods of bandwidth optimization and issues associated with these

The further aim was to gather architectural and engineering guidance on future work in the bandwidth optimisation area based on the discussions around the proposed approaches. The workshop also explored possible areas for standardization, e.g. new protocols that can aid bandwidth optimization whilst ensuring user security inline with new work in the transport layer.

1.3. Organization of this report

This workshop report summarizes the contributions to and discussions at the workshop, organized by topic. The workshop began with scene setting topics which covered the issues around deploying encryption, the increased need for privacy on the internet and setting a clear understanding that ciphertext should remain unbroken. Later sessions focused on key solution areas; these included evolution on the transport layer and sending data up or down the path. A session on application layers and CDNs aimed to highlight both issues and solutions experienced on the application layer. The workshop ended with a session dedicated to technical response to regulation with regards to encryption. The contributing documents were split between identifying the issues experienced with encryption on radio networks and suggested solutions. Of the solutions suggested some focused on transport evolution, some on trusted middleboxes and others on collaborative data exchange. Solutions were discussed within the sessions. All accepted position papers and detailed transcripts of discussion are available at [MARNEW].

The outcomes of the workshop are discussed in Section [X], and discuss progress after the workshop toward each of the identified work items as of the time of publication of this report.

Although policy related topics were out of scope for this workshop they were infrequently referred to. Report readers should be reminded that this workshop did not and did not aim to discuss policy or policy recommendations.

1.4. Use of Note Well and Charter House Rule

The workshop was conducted under the IETF [NOTE WELL] with the exception of the "Technical Analysis and Response to Potential Regulatory Reaction" session which was conducted under [CHATHAM HOUSE RULE].

1.5. IETF and GSMA

The IETF and GSMA have divergent working practices, standards and processes. IETF is an open organisation with community driven standards with the key aim of functionality and security for the internet's users, the GSMA is membership based and serves the needs of its membership base most of whom are mobile network operators.

Unlike IETF, GSMA makes few standards. Within the telecommunications industry standards are set in various divergent groups depending on their purpose. Perhaps of most relevance to the bandwidth optimisation topic here is the work of the [3GPP] which work on radio

network and core network standards with their members which include mobile operators and original equipment manufacturers.

One of the [3GPP] standards relevant to this workshop is PCC-QoS. Traditionally mobile networks have managed different applications and services based on the resources available and priorities given; for instance, emergency services have a top priority, data has a lower priority and voice services are somewhere inbetween. [3GPP] defined the PCC-QoS mechanism to support this functionality, some of which cannot occur for encrypted communications.

2. Scene Setting Sessions

Scene setting sessions aimed to bring all attendees up to a basic understanding of the problem and the scope of the workshop. There were three scene setting sessions: Scene Setting (defining scope), Encryption Deployment Considerations and Trust Models and User Choice (Privacy).

2.1. Scene Setting

The telecommunications industry and internet standards are extremely different in terms of ethos and business practices. Both industries drive technical standards in their domain and build technical solutions with some policy-driven use cases. These technologies, use cases and technical implementations are different; not only this but motivators between the two industries are also diverse.

To ensure all attendees were aligned with contributing to discussions and driving solutions this Scene Setting session worked on generating a clear scope with all attendees involved. In short: it was agreed that ciphertext should not be broken by any solution, that the radio access network (RAN) is different and does experience issues with increased encrypted traffic, that we need to understand what those problems are precisely and that our goal is to improve user experience on the Internet. Technical solutions for regulation was not in scope. The full scope is given below.

2.1.1. Scope

The attendees identified and agreed the following scope:

- o In discussion we should assume: No broken crypto, Ciphertext increasingly common, congestion does need to be controlled as do other transport issues and Network management including efficient use of resources, in RAN and elsewhere, has to work

- o How/why is RAN different for transport; help us understand the complexities of the RAN and how hard it is to manage and why those matter
- o What are the precise problems caused by more ciphertext
- o Identify players, incl. Users, and resulting tensions and how ciphertext changes those
- o Some solutions will be radically changed by ciphertext, it's ok to talk about that
- o As good as possible Quality of experience for end user is a goal
- o Our aim for the next two days is to analyse the situation and identify specific achievable tasks that could be tackled in the IETF or GSMA (or elsewhere?) and that improve the Internet given the assumptions above
- o We should not delve into:
- o Ways of doing interception (legal or not), see RFC2804 for why
- o Unpredictable political actions.

2.1.1.2. Encryption Statistics and Radio Access Network Differences

Attendees were shown that encrypted content is reaching around 50% according to recent statistics [STATE BROWSER] and [STATE SERVER]. The IAB are encouraging all IETF groups to consider encryption by default on their new protocol work and the IETF are also working on encryption on lower layers, for example TCP encryption within the [TCPINC] Working Group. The aims of these items of work are greater security and privacy for users and their data.

Within telecommunications middleboxes exist on operator networks which have previously considered themselves trusted; but qualifying trust is difficult and should not be assumed. Some interesting use cases exist with these middleboxes; such as anti-spam and malware, but these need to be balanced against their ability to open up cracks in the network for attacks such as pervasive monitoring. Some needs to improve the radio access network quality of service could come from increasing radio access network cells ("Base Stations"), but this adds to radio pollution; this shows the balancing act when devising radio access network architecture.

2.2. Encryption Deployment Considerations

Encryption across the internet is on the rise. However, some organisations and individuals come across a common set of operational issues when deploying encryption, mainly driven by commercial perspectives. The [UBIQUITOUS] draft explains these network management function impacts, detailing areas around incident monitoring, access control management, and regulation on mobile networks. The data was collected from various internet players, including system and network administrators across enterprise, governmental organisations and personal use. The aim of the document is to gain an understanding of what is needed for technical solutions to these issues, maintaining security and privacy for users. Attendees commented that worthwhile additions would be: different business environments (e.g. cloud environments) and service chaining. Incident monitoring in particular was noted as a difficult issue to solve given the use of URL in today's incident monitoring middleware.

Some of these impacts to mobile networks can be resolved using difference methods and the [NETWORK MANAGEMENT] draft details these methods. The draft focuses heavily on methods to manage network traffic without breaching user privacy and security.

By reviewing encryption deployment issues and the alternative methods of network management MarNEW attendees were made aware of the issues which affect radio networks, the deployment issues which are solvable and require no further action, and those which aren't currently solvable and which should be addressed within the workshop.

2.3. Trust Models and User Choice (Privacy)

Solutions of how to improve delivery of encrypted content could affect some of all of the privacy benefits that encryption brings. Understanding user needs and desires for privacy is therefore important when designing these solutions.

From a recent study [Pew2014] 64% of users said concerns over privacy have increased, 67% of mobile internet users would like to do more to protect their privacy. The W3C and IETF have both responded to user desires for better privacy by recommending encryption for new protocols and web technologies. Within the W3C new security standards are emerging and the design principles for HTML hold that users are the stakeholders with most priority, followed by implementors and other stakeholders, further enforcing the "user first" principle. Users also have certain security expectations from particular contexts, and sometimes use new technologies to further

protect their privacy even if those technologies weren't initially developed for that purpose.

Technologies which can impact user privacy sometimes do this ignorant of the privacy implications or incorrectly assume that the benefits users gain from the new technology outweigh the loss of private information. Any new technology which introduces bad security vectors will be used by attackers. If these technologies are necessary they should be opt-in.

Internet stakeholders should understand the priority of other stakeholders. Users should be considered the first priority, other stakeholders include implementors, developers, advertisers, operators and other ISPs. Some technologies have been abused by these parties, such as cookie use or JavaScript injection. This has caused some developers to encrypt content to circumnavigate these technologies which they find intrusive or bad for their users privacy.

Some suggested solutions for network management of encrypted traffic have suggested "trust models". If users and content providers are to opt-in to user network management services with negative privacy impacts they should see clear value from using these services, and understand the impacts on clear interfaces. Users should also have easy abilities to opt-out. Some users will always automatically click through consent requests, so any trust model is flawed for these users. Understanding the extent of "auto click through" may help make better decisions for consent requests in the future. One trust model (Cooperative Traffic Management) works as an agent of the user; by opting-in metadata can be shared. Issues with this involve trust only being applied on end.

3. Network or Transport Solution Sessions

Network or Transport Solution Sessions aimed to discuss suggested and new solutions for managing encrypted traffic on radio access networks. Most solutions focus on the sharing of metadata; either from the endpoint to the network, from the network to the endpoint, or cooperative sharing between both. Evolutions on the transport layer could be another approach to solve some of the issues radio access networks experience which cause them to require network management middleboxes. By removing problems on the transport layer the need to expensive middleboxes could decrease.

3.1. Sending Data Up / Down for Network Management Benefits

Middleboxes in the network have a number of uses, some which are more beneficial than they are controversial. Collaboration between these

network elements and the endpoints could bring about better content distribution. A number of suggestions were given, these included:

- o Mobile Throughput Guidance: exchanges data between the network elements and the endpoints via TCP Options. It also allows for gaining a better idea of how the transport protocol behaves and improving user experience further, although the work still needs to evolve.
- o SPUD: a UDP-based encapsulation protocol to allow explicit cooperation with middleboxes while using new, encrypted transport protocols.
- o Network Status API: An API for operators to share congestion status or the state of a cell before an application starts sending data could allow applications to change their behaviour.
- o Traffic classification: classifying traffic and adding this as metadata for analysis throughout the network. This idea has trust and privacy implications.
- o ConEx: a mechanism where senders inform the network about the congestion encountered by previous packets on the same flow, in-band at the IP layer.
- o Latency versus Bandwidth: allowing the content provider to indicate whether a better bandwidth or lower latency is of greater priority and allowing the network to react. Where this bit resides and how to authenticate it would need to be decided.
- o No network management tools: disabling all network management tools from the network and allow the protocols to manage congestion alone.
- o FlowQueue Codel: a hybrid packet scheduler/AQM algorithm, aiming to reduce bufferbloat and latency. FQ-CoDel mixes packets from multiple flows and reduces the impact of head of line blocking from bursty traffic [FQ CODEL].

Many of these suggestions could be labeled "Network-to-App", a better approach may be "Network-to-User", to achieve this these ideas would need to be expanded. Others aim to create "hop-to-hop" solutions, which could be more inline with how congestion is managed today, but with greater privacy IMPLICATIONS.

"App-to-Network" style solutions have either existed for a long time by implicit solutions, or explicitly defined but never implemented or properly deployed. Some workshop attendees agreed that applications

declaring was quality of service they require was not a good route given the lack of success in the past.

3.1.1. Trust and the Mobile Network Complexities

One of the larger issues in the sharing of data is the matter of trust; networks operators find difficulties in relinquishing data for reasons such as revealing competitive information and applications wish to protect their users and only reveal little information to the network. Authentication in that case could be a key design element of any new work, as well as explicitness rather than the transparent middleboxes used more recently. Some workshop attendees suggested any exchange of information should be bidirectional, in an effort to improve trust between the elements. A robust incentive framework could provide a solution to the trust issue, or at least help mitigate it.

The radio access network is complex and manages a number of realities. Base stations understand many of these realities, and information within these base stations can be of value other entities on the path. Solutions for managing congestion on radio networks should involve the base station if possible. For instance, understanding how the Radio Resource Controller and AQM interact (or don't interact) could provide valuable information for solving issues. Although many workshop attendees agreed that even though there is a need to understand the base station not all agreed that the base station should be part of a future solution.

Some suggested solutions were based on network categorisation and providing this information to the protocols or endpoints. Categorising radio networks could be impossible due to their complexity, but categorising essential network properties could be possible and valuable.

4. Transport Layer: Issues, Optimisation and Solutions

TCP has been the dominant transport protocol since TCP/IP replaced NCP on the Arpanet in March 1983. TCP was originally devised to work on a specific network model that did not anticipate the high error rates and highly variable available bandwidth scenarios experienced on modern radio access networks. Furthermore new network elements have been introduced (NATs and network devices with large buffers creating bufferbloat), and considerable peer-to-peer traffic is competing with traditional client-server traffic. Consequently the transport layer today has requirements beyond what TCP was designed to meet. TCP has other issues as well; too many services rely on TCP and only TCP, blocking deployment of new transport protocols like SCTP and DCCP. This means that true innovation on the transport

layer becomes difficult because deployment issues are more complicated than just building a new protocol.

The IETF is trying to solve these issues through the "Stack Evolution" programme, and the first step in this programme is to collect data. Network and content providers can provide data including: the cost of encryption, the advantages of network management tools, the deployment of protocols, and the effects when network management tools are disabled. Network operators do not tend to reveal network information mostly for competition reasons and so is unlikely to donate this information freely to IETF. The GSMA is in the position to collect this data and anonymise it before bringing it to IETF which should alleviate the network operator worries but still provide IETF with some usable data.

A considerable amount of work has already been done on TCP, especially innovation in bandwidth management and congestion control; although congestion is usually detected by detecting loss, and better methods based on detecting congestion would be beneficial.

Furthermore, although the deficiencies of TCP are often considered as key issues in the evolution of the stack, the main route to resolve these issues may not be a new TCP, but an evolved stack. SPUD and ICN are two suggestions which may help here. QUIC engineers stated that the problems solved by QUIC are general problems, rather than TCP issues. This view was not shared by all attendees of the workshop. Moreover, TCP has had some improvements in the last few years which may mean some of the network lower layers should be investigated to see whether improvements can be made here.

5. Application Layer Optimisation, Caching and CDNs

Many discussions on the effects of encrypted traffic on radio access networks happen between implementers and the network operators; this session aimed to gather the opinions of the content and caching providers including their experiences running over mobile networks, the experience their users expect, and what they would like to achieve by working with or using the mobile network.

Content providers explained how even though this workshop cited encrypted data over radio access networks as the main issue the real issue is network management generally, and all actors (applications providers, networks and devices) need to work together to overcome these general network management issues. Content providers explained how they assume the mobile networks are standard compliant. When the network is not standards compliant (e.g. using non standards compliant intermediaries) content providers can experience real costs

as users contact their support centres to report issues which are difficult to test for and build upon.

Content providers cited other common issues concerning data traffic over mobile networks. Data caps cause issues for users; users are confused about how data caps work or are unsure how expensive media is and how much data it consumes. DNS and DNS caching cause unpredictable results. Developers build products on networks not indicative of the networks their customers are using and not every organisation has the finances to build a caching infrastructure.

Strongly related to content providers, CDNs are understood to be a trusted deliver of content and have shown great success in fixed networks. Now traffic is moving more to mobile networks there is a need to place caches at the edge of the network (e.g. in the Gi LAN or the radio network) within the mobile network. Places caches at the edge of the mobile network is a solution, but requires standards developed by content providers and mobile network operators. The CNDi working group at IETF aims to allow global CDNs to interoperate with mobile CDNs; but this causes huge trust issues for the caching of encrypted data between these CDNs. Some CDNs are experimenting with "Keyless SSL" to enable safer storage of content without passing private keys to the CDN. Blind Caching is another proposal aimed at caching encrypted content closer to the user and managing the authentication at the original content provider servers.

At the end of the session the panelists were asked to identify one key collaborative work item, these were: evolving caching to cache encrypted content, using one-bit for latency / bandwidth trade-off (explained below), better collaboration between the network and application, better metrics to aid bug solving and innovation, and indications from the network to allow the application to adapt.

6. Technical Analysis and Response to Potential Regulatory Reaction

This session was conducted under Chatham House Rule. The session aimed to discuss regulatory and political issues; but not their worth or need, rather to understand the laws that exist and how technologists can properly respond to these.

Mobile networks are regulated, compliance is mandatory (and can result in service license revocation in some nations round the world) and can incur costs on the mobile network operator. Regulation does vary geographically. Some regulations are court orders, others are "block lists" of websites such as the Internet Watch Foundation list [IWF]. Operators are not expected to decrypt sites, so those identified sites which are encrypted will not be blocked.

Parental control-type filters also exist on the network and are easily bypassed today, vastly limiting their effectiveness. Better solutions would allow for users to easily set these restrictions themselves. Other regulations are also hard to meet - such as user data patterns, or will become harder to collect - such as IoT cases. Most attendees agreed that if the governments cannot get the information from network operators they will approach the content providers. Some governments are aware of the impact of encryption and are working with or trying to work with content providers. The IAB have concluded blocking and filtering can be done at the endpoint of the communication.

These regulations do not always apply to the internet, and the internet community is not always aware of their existence. Collectively the internet community can work with GSMA and 3GPP and act collectively to alleviate the risk imposed by encrypted traffic for lawful intercept. The suggestion from attendees was that if any new technical solutions built should have the ability to be easily switched off.

Some mobile network operators are producing transparency reports covering regulations including lawful intercept. Operators who have done this already are encouraging others to do the same.

7. Requirements and Suggestions for Future Solutions

Based on the talks and discussions throughout the workshop a set of requirements and suggested solutions has been collected. This is not an exhaustive list.

- o Encrypted Traffic: any solution should encourage and support encrypted traffic.
- o Flexibility: radio access network qualities vary vastly and different network needs in content can be identified, so any new solution should be flexible to either the network type or content type or both.
- o Privacy: new solutions should not introduce ways where information can be discovered flows and attribute them to users.
- o Minimum data only for collaborative work: user data, app data and network data all needs protecting, so new solutions should use the minimum information to make a working solution.

A collection of solutions suggested throughout the workshop is given below. These solutions haven't been matched to the requirements above, so this step will need to come later.

- o Evolving TCP or evolution on the transport layer: this could take a number of forms and some of this work is already existing within IETF. Other suggestions include:
- o Congestion Control: many attendees cited congestion control as a key issue, further analysis, investigation and work could be done here.
- o SPROUT: research at MIT which is a transport protocol for interactive applications that desire high throughput and low delay. [SPROUT]
- o PCC: Performance-oriented Congestion Control: is a new architecture that aims for consistent high performance even in challenging scenarios. PCC endpoints observe the connection between their actions and their known performance, which allows them to adapt their actions. [PCC]
- o CDNs and Caches: placing caches closer to the mobile user or making more intelligent CDNs would result in faster content delivery and less strain on the network. Related work includes:
- o Blind Caching: a proposal for caching of encrypted content.
- o CDN improvements: including keyless SSL and better CDN placement.
- o Mobile Throughput Guidance: a mechanism and protocol elements that allow the cellular network to provide near real-time information on capacity available to the TCP server. [MTG]
- o One bit for latency / bandwidth tradeoff: using one bit to identify whether a stream prefers low latency at the expense of throughput. This rids solutions of the trust issue as applications will need to select the best scenario for their traffic type.
- o Base Station: some suggestions involved "using the Base Station", but this was not defined in detail. The Base Station holds the Radio Resource Controller and scheduler which could provide either a place to host solutions or data from the Base Station could help in devising new solutions.
- o Identify traffic types via 5-tuple: information from the 5-tuple could provide understanding of the traffic type which network management could then be applied.
- o Heuristics: Networks can sometimes identify traffic types through specifics such as data flow rate and then apply network management

to these identified flows. This is not recommended as categorisations can be incorrect.

- o APIs: An API for operators to share congestion status or the state of a cell before an application starts sending data could allow applications to change their behaviour. Alternatively an API could provide the network with some information on the data type to provide network management to, although this method exposes privacy issues.
- o Standard approach for operator to offer services to Content Providers: mobile network operators could provide caching services or other services for content providers to use for faster and smoother content delivery.
- o AQM [AQM] and ECN [ECN] deployments: queueing and congestion management methods have existed for sometime in the form of AQM, ECN and others which can help the transport and internet layer adapt to congestion faster.
- o Trust Model or Trust Framework: some solutions in this area (e.g. SPUD) have a reliance on trust when content providers or the network are being asked to add classifiers to their traffic.
- o Keyless SSL: allows content providers to maintain their private keys on a "key server" and host the content elsewhere (e.g. on a CDN). This could become standardised in IETF. [LURK]
- o Meaningful capacity sharing: including the ConEx [CONEX] work which exposes information about congestion to the network nodes.
- o Hop-by-hop: some suggestions offer hop-by-hop methods allowing nodes to adapt flow given the qualities of the networks around them and the congestion they're experiencing.
- o Metrics and metric standards: in order to evolve current protocols to be best suited to today's networks data is needed on the current network situations, protocol deployments, packet traces and middlebox behaviour. Further than this proper testing and debugging on networks could provide great insight for stack evolution.
- o 5G: Mobile operator standards bodies are in the process of setting the requirements for 5G, requirements for network management could be added.

In the workshop attendees identified other areas where greater understand could help the standards process. These were identified as:

- o Greater understanding of the RAN at IETF
- o Reviews and comments on 3GPP perspective
- o How to do congestion controlling in RAN

7.1. Better Collaboration

Throughout the workshop attendees placed emphasis on the need for better collaboration between the IETF and telecommunications bodies and organisations. The workshop was one such way to achieve this, but the good work and relationships built in the workshop should continue so the two groups can work on solutions which are better for both technologies and users.

8. Next Steps

The next steps for MarNEW attendees are to begin work on a select list of the above recommended solutions and other suggestions within the IETF and within other organisations. At IETF95 the ACCORD BoF will be held which will bring the workshop discussion to the wider IETF attendance and select key areas to progress on; these are likely to be definitions of the metrics to be collected, more information on the stack evolution ideas and their impact to network management, Mobile Throughput Guidance evolution, evolution of the Blind Caching work and draft definitions of the "1 bit for latency / bandwidth tradeoff" idea. As identified in the "Better Collaboration" section together we need to ensure that both groups continue the positive relationship to move these ideas forward into being real and workable solutions and both groups need to understand that even though collaboration between the operator network and the internet is of great importance the item of most importance is the experience and security for the users using these services.

Author's Address

Natasha Rooney

Email: nrooney@gsma.com

URI: <https://gsma.com>

Tsvwg Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 14, 2016

J. You
Huawei
M. Welzl
University of Oslo
B. Trammell
M. Kuehlewind
ETH Zurich
K. Smith
Vodafone Group
March 13, 2016

Latency Loss Tradeoff PHB Group
draft-you-tsvwg-latency-loss-tradeoff-00

Abstract

This document defines a PHB (Per-Hop Behavior) group called Latency Loss Tradeoff (LLT). The LLT group is intended to provide delivery of IP packets in two classes of services: a low-loss service (Lo service) and a low-latency service (La service). The LLT group enables an application to request treatment for either low-loss or low-latency at a congested network link.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 14, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
2.1. Abbreviations and acronyms	3
2.2. Definitions	4
3. Problem Statement	5
3.1. Existing DSCP PHBs	5
3.1.1. Default PHB	5
3.1.2. Class Selector PHB	5
3.1.3. Assured Forwarding PHB Group	5
3.1.4. Expedited Forwarding PHB	6
3.1.5. Voice Admit PHB	6
3.1.6. Delay Bound PHB	6
3.2. Incentives	6
4. Definition of LLT PHB	7
4.1. Goal and Scope of LLT	7
4.2. Description of LLT behavior	8
4.2.1. Implementation Considerations	8
4.3. Microflow misordering	9
4.4. Recommended Codepoints	9
4.5. Mutability	9
4.6. Tunneling	9
4.7. Interaction with other PHBs	9
5. Security Considerations	10
6. IANA Considerations	10
7. References	10
7.1. Normative References	10
7.2. Informative References	11
Authors' Addresses	11

1. Introduction

Different applications have different communication requirements [QoS]. In interactive applications of real-time sound transmission, as well as in virtual reality, the overall one-way delay needs to be short in order to give the user an impression of a real-time response. Yet, these applications may be able to tolerate high loss rates. In conventional text and data networking, delay thresholds are the least stringent. The response time in these types of applications can increase from 2 to 5 seconds before becoming unacceptable. However, given that increased loss reduces the throughput of TCP, these applications desire minimal loss.

The network resources consist primarily of buffers and link bandwidth. Operators often favor high utilization of bottleneck links at the price of high queuing delay. This is beneficial for non-real time applications. However, this may be considered unacceptable for some real-time applications. The proposed LLT group enables an application to choose between low-latency and low-loss at a congested network link [ABE] [RD]. Typically, an interactive application with real-time deadlines, such as audio, will mark most of its packets as a low-latency service. In contrast, an application that transfers bulk data will mark most of its packets as a low-loss service. The LLT group can be thought of as allowing an application to trade loss for delay by marking packets low-latency service (La) or to trade delay for loss by marking packets low-loss service (Lo).

2. Terminology

This section contains definitions for terms used frequently throughout this document.

2.1. Abbreviations and acronyms

DS: Differentiated Service

PHB: Per-Hop Behavior

LLT: Latency Loss Tradeoff

TCA: Traffic Conditioning Agreement

TCP: Transmission Control Protocol

2.2. Definitions

DS-capable: capable of implementing differentiated services as described in this architecture; usually used in reference to a domain consisting of DS-compliant nodes.

DS codepoint: a specific value of the DSCP portion of the DS field, used to select a PHB.

DS-compliant: enabled to support differentiated services functions and behaviors as defined in [RFC2474], this document, and other differentiated services documents; usually used in reference to a node or device.

DS field: the IPv4 header TOS octet or the IPv6 Traffic Class octet when interpreted in conformance with the definition given in [RFC2474]. The bits of the DSCP field encode the DS codepoint, while the remaining bits are currently unused.

Low-latency service (La service): puts an emphasis on low queuing delay at a congested network link. It allows an application to trade loss for delay.

Low-loss service (Lo service): puts an emphasis on low packet loss rate at a congested network link. It allows an application to trade delay for loss.

Per-Hop-Behavior (PHB): the externally observable forwarding behavior applied at a DS-compliant node to a DS behavior aggregate.

PHB group: a set of one or more PHBs that can only be meaningfully specified and implemented simultaneously, due to a common constraint applying to all PHBs in the set such as a queue servicing or queue management policy. A PHB group provides a service building block that allows a set of related forwarding behaviors to be specified together (e.g., four dropping priorities). A single PHB is a special case of a PHB group.

Traffic Conditioning Agreement (TCA): an agreement specifying classifier rules and any corresponding traffic profiles and metering, marking, discarding and/or shaping rules which are to apply to the traffic streams selected by the classifier. A TCA encompasses all of the traffic conditioning rules explicitly specified within a SLA along with all of the rules implicit from the relevant service requirements and/or from a DS domain's service provisioning policy.

3. Problem Statement

3.1. Existing DSCP PHBs

3.1.1. Default PHB

A default Per-Hop Behavior (PHB) [RFC2474] MUST be available in a DiffServ (DS)-compliant node. This is the common, best-effort forwarding behavior available in existing routers as standardized in [RFC1812]. Codepoint '000000' from Pool 1 is used as the default PHB value. In this document, packets received with the Default PHB is treated as Lo service on the LLT-compliant router.

3.1.2. Class Selector PHB

The Class Selector (CS) PHB [RFC2474] is introduced for backwards compatibility with use of the IPv4 Precedence field. Any of the eight codepoints in the range 'xxx000' (where 'x' may equal '0' or '1') from Pool 1 is assigned as Class Selector codepoint. The CS PHB does not match the services that LLT PHB is trying to deliver.

3.1.3. Assured Forwarding PHB Group

The Assured Forwarding (AF) PHB group [RFC2597] allows the operator to provide assured forwarding of IP packets as long as the aggregate traffic does not exceed the subscribed rate. Traffic that exceeds the subscribed rate is not delivered with as high a probability as the traffic that is within the rate.

The AF PHB group provides delivery of IP packets in four independently forwarded AF classes. Within each AF class, an IP packet can be assigned one of three different levels of drop precedence. The combination of classes and drop precedence yields twelve separate DSCP encodings from AF11 through AF43 as follows:

	Class 1	Class 2	Class 3	Class 4
Low Drop Prec	001010	010010	011010	100010
Medium Drop Prec	001100	010100	011100	100100
High Drop Prec	001110	010110	011110	100110

The AF PHB does not match the services that LLT PHB is trying to deliver.

3.1.4. Expedited Forwarding PHB

Expedited Forwarding (EF) PHB [RFC3246] is intended to provide a building block for low delay, low jitter and low loss services by ensuring that the EF aggregate is served at a certain configured rate. EF traffic requires a strict admission control mechanism. Codepoint '101110' is recommended for the EF PHB. The EF PHB does not match the services that LLT PHB is trying to deliver.

3.1.5. Voice Admit PHB

The Voice Admit (VA) PHB [RFC5865] has identical characteristics to the Expedited Forwarding PHB. However Voice Admit traffic is also admitted by the network using a Call Admission Control (CAC) procedure. The recommended DSCP for Voice Admit is '101100', parallel with the existing EF codepoint '101110'. The VA PHB does not match the services that LLT PHB is trying to deliver.

3.1.6. Delay Bound PHB

The Delay Bound (DB) PHB [RFC3248] requires a bound on the delay of packets due to other traffic in the network. Two parameters - capped arrival rate (R) and a 'score' (S), are defined and related to the target delay variation bound. An experimental codepoint '101111' is suggested for DB behavior. In this document, there's no specific bound on the delay, the LLT PHB only indicates the tradeoff.

3.2. Incentives

The primary goal of differentiated services is to allow different levels of service to be provided for traffic streams on a common network infrastructure. Hence, an adversary may be able to obtain better service by modifying the DS field to codepoints indicating behaviors used for enhanced services or by injecting packets with the DS field set to such codepoints. Such theft-of-service ([RFC2474], [RFC2475]) becomes a denial-of-service attack when the modified or injected traffic depletes the resources available to forward it and other traffic streams.

DS ingress nodes must condition all traffic entering a DS domain to ensure that it has acceptable DS codepoints. This means that the codepoints must conform to the applicable TCA(s) (Traffic Conditioning Agreement) [RFC2475] and the domain's service provisioning policy. Packets received with an unacceptable codepoints must either be discarded or must have their DS codepoints modified to acceptable values before being forwarded. For example, an ingress node receiving traffic from a domain with which no enhanced service agreement exists may reset the DS codepoint to the

Default PHB codepoint. However, the Default PHB (i.e. best-effort forwarding) cannot meet the diverse needs of different Internet applications.

The objective of the LLT PHB group is to retain the best-effort service while providing low delay to real-time applications at the expense of increased loss or providing low loss to non real-time applications at the expense of increased delay. This requires Internet applications to mark their traffic with appropriate codepoint values. Since the low-loss service is neither better nor worse than the low-latency service but is merely different, there is no incentive for Internet applications to abuse such codepoints, and no need for admission control.

4. Definition of LLT PHB

The LLT group provides forwarding of IP packets in two classes of service: a low-loss service (Lo) and a low-latency service (La). The LLT group enables an application to choose between low latency and low loss at a congested network link. The packets marked as low-latency service receive little queuing delay. The packets marked as low-loss service receive at least as much throughput as they would in a legacy best effort network. La-marked packets are more likely to be dropped during periods of congestion than the Lo-marked packets. Note that among the two services, neither of the two has priority over the other.

4.1. Goal and Scope of LLT

The LLT group may be used by a network operator in two distinct ways: either as a separate service, or as a replacement of the flat (existing) best-effort IP service.

A DS (Differentiated Services) node SHOULD implement the LLT group. It MAY allocate a configurable, minimum amount of forwarding resources (buffer space and bandwidth) to LLT group.

The LLT group MAY also be configurable to receive more forwarding resources than the minimum when excess resources are available from other PHB groups. This is beyond the scope of this document.

The LLT PHB definition does NOT mandate or recommend any particular method for achieving LLT behavior.

4.2. Description of LLT behavior

To support the LLT group on an output link, the router can maintain two FIFO (First-In First-Out) queues referred to as a Lo (Loss-sensitive) queue and La (Latency-sensitive) queue for packets destined to the link. Depending on whether an incoming packet is marked for the low-loss or low-latency service, the router appends the packet to the Lo or La queue respectively. The packets within each queue are served in the FIFO order. The scheduling is work-conserving.

A router can support the desired delay differentiation between the Lo and La services through buffer sizing for the Lo and La queues, and by ensuring that the La queue does not grow larger than the Lo queue. As common in current Internet routers, the size of the Lo buffer is chosen large enough so that the oscillating transmission of TCP (Transmission Control Protocol) and other legacy end-to-end congestion control protocols utilizes the available link rate fully. The La buffer is configured to a much smaller dynamic size to ensure that queuing delay for each forwarded packet of the La class is low. The assurance of low maximum queuing delay is attractive for delay-sensitive applications and easily verifiable by outside parties.

4.2.1. Implementation Considerations

This document does not specify any particular implementation method for achieving LLT behavior. Some LLT-like implementations may refer to [I-D.hurley-alternative-best-effort], [RD] and [I-D.briscoe-aqm-dualq-coupled].

[I-D.hurley-alternative-best-effort] marks every best effort packet as either green or blue. Green packets receive a low, bounded delay at every hop, the value of the per-hop delay bound configured by the operator. However, when transmitting more aggressively, the green users can enjoy both a higher rate and lower queuing delay than those of the blue users, which weakens the incentives for incremental deployment. [RD] proposes Rate-Delay (RD) service enabling a user to choose either a higher transmission rate or low queuing delay. The R (Rate) service is like Lo service while D (Delay) service is like La service.

Note that both classes defined in this document do not provide any absolute guarantees on the loss rate or delay a packet will experience. Using these classes only provides a relative treatment compared to the other class. Depending on the amount of traffic arriving per class, it is possible for traffic in the La class to experience more delay than traffic in the Lo class. However, this may be circumvented by using scheduling mechanisms, for example, by

adjusting the scheduling function that assigns traffic to the Lo and La queues, or by adjusting the scheduling weight based on the average load in each class. Moreover, the delay experienced by La traffic is always bounded by the length of the La queue. The particular implementation is beyond the scope of this document.

When a DS-compliant node claims to implement the LLT PHB, the implementation **MUST** conform to the specification given in this document.

4.3. Microflow misordering

The packets within each queue are served in the FIFO order. Packets belonging to a single microflow within the LLT aggregate **SHOULD NOT** experience re-ordering in normal operation of the device when passing through.

4.4. Recommended Codepoints

Recommended codepoints for the LLT PHB group are given below.

Low-loss service: 000001
Low-latency service: 000101

4.5. Mutability

Packets marked for LLT PHB **MAY** be remarked at a DS domain boundary only to other codepoints that satisfy the LLT PHB. Packets marked for LLT PHBs **SHOULD NOT** be demoted or promoted to another PHB by a DS domain.

4.6. Tunneling

When LLT packets are tunneled, the tunneling packets must be marked as LLT.

4.7. Interaction with other PHBs

Other PHBs and PHB groups may be deployed in the same DS node or domain with the LLT PHB.

Packets received with the Default PHB **SHOULD** be treated as Lo service as default by the LLT PHB aware device. [RD] has proved that La service does not hurt Lo service.

Packets received with the LLT PHB **SHOULD** be treated as Default PHB as default by the LLT PHB unaware device.

5. Security Considerations

Internet applications cannot benefit from wrongly indicating low-loss or low-latency as they have to pay the expense of high delay or high loss as a tradeoff. Hence there is no incentive for Internet applications to set the wrong codepoints.

6. IANA Considerations

This document suggests two experimental codepoints, 000001/000101, in Pool 3 of the code space defined by [RFC2474].

7. References

7.1. Normative References

- [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <<http://www.rfc-editor.org/info/rfc1812>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.
- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <<http://www.rfc-editor.org/info/rfc2475>>.
- [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, DOI 10.17487/RFC2597, June 1999, <<http://www.rfc-editor.org/info/rfc2597>>.
- [RFC3246] Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, DOI 10.17487/RFC3246, March 2002, <<http://www.rfc-editor.org/info/rfc3246>>.

- [RFC3248] Armitage, G., Carpenter, B., Casati, A., Crowcroft, J., Halpern, J., Kumar, B., and J. Schnizlein, "A Delay Bound alternative revision of RFC 2598", RFC 3248, DOI 10.17487/RFC3248, March 2002, <<http://www.rfc-editor.org/info/rfc3248>>.
- [RFC5865] Baker, F., Polk, J., and M. Dolly, "A Differentiated Services Code Point (DSCP) for Capacity-Admitted Traffic", RFC 5865, DOI 10.17487/RFC5865, May 2010, <<http://www.rfc-editor.org/info/rfc5865>>.

7.2. Informative References

- [ABE] Hurley, P., Le Boudec, J., Thiran, P., and M. Kara, "ABE: Providing a Low-Delay Service within Best Effort", IEEE Network Magazine 15(3): 60-69, May 2001.
- [I-D.briscoe-aqm-dualq-coupled] Schepper, K., Briscoe, B., Bondarenko, O., and I. Tsang, "DualQ Coupled AQM for Low Latency, Low Loss and Scalable Throughput", draft-briscoe-aqm-dualq-coupled-00 (work in progress), August 2015.
- [I-D.hurley-alternative-best-effort] Hurley, P., Iannaccone, G., Kara, M., Le Boudec, J., Thiran, P., and C. Diot, "The ABE Service", November 2000.
- [QoS] Chen, C., Farley, T., and N. Ye, "QoS Requirements of Network Applications on the Internet", Information Knowledge Systems Management 2004, 4(1): 55-76, 2004.
- [RD] Podlesny, M. and S. Gorinsky, "RD network services: differentiation through performance incentives", ACM SIGCOMM Computer Communication Review, 38(4): 255-266, 2008.

Authors' Addresses

Jianjie You
Huawei
101 Software Avenue, Yuhua District
Nanjing 210012
China

Email: youjianjie@huawei.com

Michael Welzl
University of Oslo
PO Box 1080 Blindern
Oslo N-0316
Norway

Email: michawe@ifi.uio.no

Brian Trammell
ETH Zurich
Zurich
Switzerland

Email: ietf@trammell.ch

Mirja Kuehlewind
ETH Zurich
Zurich
Switzerland

Email: mirja.kuehlewind@tik.ee.ethz.ch

Kevin Smith
Vodafone Group
One Kingdom Street,
London
UK

Email: kevin.smith@vodafone.com