Internet Architecture Board (IAB)
Request for Comments: 8462
Category: Informational
ISSN: 2070-1721

N. Rooney
S. Dawkins, Ed.

October 2018

Report from the IAB Workshop on Managing Radio Networks in an Encrypted World (MaRNEW)

Abstract

The Internet Architecture Board (IAB) and GSM Association (GSMA) held a joint workshop on Managing Radio Networks in an Encrypted World (MaRNEW), on September 24-25, 2015. This workshop aimed to discuss solutions for bandwidth optimization 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 participants from various organizations involved in the telecommunications industry including original equipment manufacturers, content providers, and mobile network operators.

The group discussed Internet encryption trends and deployment issues identified within the IETF and the privacy needs of users that should be adhered to. Solutions designed around sharing data from the network to the endpoints and vice versa were then discussed; in addition, issues experienced when using current transport-layer protocols were also discussed. Content providers and Content Delivery Networks (CDNs) gave their own views of their experiences delivering their content with mobile network operators. Finally, technical responses to regulation were 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

Copyright Notice

Table of Contents

1. Introduction

The Internet Architecture Board (IAB) and GSM Association (GSMA) held a joint workshop on Managing Radio Networks in an Encrypted World (MaRNEW), on September 24-25, 2015. This workshop aimed to discuss solutions for bandwidth optimization 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.

Mobile networks have a set of properties that place a large emphasis on sophisticated bandwidth optimization. The use of encryption is increasing on the Internet, which is positive for consumer and business privacy and security. Many existing solutions for mobile bandwidth optimization primarily operate on non-encrypted communications; this can lead to performance issues being amplified on mobile networks. The use of encryption on networks will continue to increase; 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, prioritization, optimization, and efficiency enhancements. It also encompasses user-related topics such as specific subscription or billing models, and it may touch upon regulatory aspects or other issues relating to government-initiated regulatory concerns.

The first category of bandwidth optimization includes the following:

- Caching
- Prioritization of interactive traffic over background traffic
- Per-user bandwidth limits

The second category of bandwidth optimization may depend on one or more of the first category optimization strategies, but may, in particular, also encompass business-related topics such as content delivery arrangements with content providers.

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

Many of these functions can continue as they are performed today, even with increased use of encryption. Others are using methods that inspect parts of the communication that are not encrypted today, but will be encrypted, and these functions will have to be done differently in an increasingly encrypted Internet.

1.2. Topics

The workshop aimed to answer questions that focused on:

- understanding the bandwidth optimization use cases particular to radio networks;
- understanding existing approaches and how these do not work with encrypted traffic;
- understanding reasons why the Internet has not standardized support for lawful intercept and why mobile networks have;
- determining how to match traffic types with bandwidth optimization methods
- discussing minimal information to be shared to manage networks but ensure user security and privacy;
- developing new bandwidth optimization techniques and protocols within these new constraints;
- discussing the appropriate network layer(s) for each management function; and
- 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 optimization 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 that user security is in line with new work in transport-layer protocols.

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 that 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 discussing a technical response to regulation with regards to encryption. The contributing documents identified 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 Sections 7 and 8; they discuss the progress made since the workshop toward each of the identified work items through the time this document was approved for publication.

Report readers should be reminded that this workshop did not aim to discuss regulation or legislation, although policy topics were mentioned in discussions from time to time.

1.4. Use of Note Well and the Chatham 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 the [CHATHAM_HOUSE_RULE].

1.5. IETF and GSMA

The IETF and GSMA [GSMA] have different working practices, standards, and processes. IETF is an open organization with community-driven standards, with the key aim of functionality and security for the Internet’s users, while 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 optimization topic here is the work of the 3rd Generation Partnership Project (3GPP) [SDO_3GPP], which works on radio network and core network standards. 3GPP members include mobile operators and original equipment manufacturers.

One of the 3GPP standards relevant to this workshop is Policy and Charging Control QoS [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 in-between. 3GPP defined the PCC-QoS mechanism to support this functionality, and this depends on unencrypted communications [EffectEncrypt].

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:

- Section 2.1: Scene Setting
- Section 2.2: Encryption Deployment Considerations
- Section 2.3: Awareness of User Choice (Privacy)

2.1. Scene Setting

The telecommunications industry and Internet standards community are extremely different in terms of ethos and practices. Both groups 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, and the 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 1) ciphertext encrypted by one party and intended to be decrypted by a second party should not be decrypted by a third party in any solution, 2) the Radio Access Network (RAN) does experience issues with increased encrypted traffic, 3) the RAN issues need to be understood precisely, and 4) the goal is to improve user experience on the Internet. Proposing new technical solutions based on presumed future regulation was not in scope. The full scope is given below.

2.1.1. Scope

The attendees identified and agreed to the scope described here.

We should do the following:

- in discussion, assume that there is no broken crypto; ciphertext is 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;
- identify how/why RAN is different for transport, and attempt to understand the complexities of RAN (i.e., how hard it is to manage) and why those complexities matter;
- identify the precise problems caused by increased use of encryption;
- identify players (in addition to end users), the resulting tensions, and how ciphertext changes those tensions;
- discuss how some solutions will be radically changed by ciphertext (it’s ok to talk about that)
- assume that the best possible quality of experience for the end user is a goal; and lastly,
- for the next two days, aim to analyze 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.

We should not delve into the following:

- ways of doing interception, legal or not, for the reasons described in [RFC2804]; and,
- unpredictable political actions.

2.1.2. Encryption Statistics and Radio Access Network Differences

According to then-current statistics, attendees were shown that encrypted content reaches around 50% [STATE_BROWSER] [STATE_SERVER]. The IAB is encouraging all IETF working groups to consider the effect encryption being "on by default" will have on new protocol work. The IETF is also working on encryption at lower layers. One recent

example of this work is opportunistic TCP encryption within the TCP Increased Security [TCPINC] Working Group. The aims of these work items are greater security and privacy for end users and their data.

Telecommunications networks often contain middleboxes that operators have previously considered to be 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 detection, but these need to be balanced against their ability to open up cracks in the network for attacks such as pervasive monitoring.

When operators increase the number of radio access network cells (base stations), this can improve the radio access network quality of service; however, it also adds to radio pollution. This is one example of the balancing act required when devising radio access network architecture.

2.2. Encryption Deployment Considerations

Encryption across the Internet is on the rise. However, some organizations and individuals that are mainly driven by commercial perspectives come across a common set of operational issues when deploying encryption. [RFC8404] 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 organizations, and personal use. The aim of the document is to gain an understanding of what is needed for technical solutions to these issues while 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 URLs in today’s incident monitoring middleware.

Some of these impacts to mobile networks can be resolved using different methods, and the [NETWORK_MANAGEMENT] document details these methods. The document 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 that affect radio networks, the deployment issues that are solvable and require no further action, and those issues that have not yet been solved but should be addressed within the workshop.

2.3. Awareness of User Choice (Privacy)

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

From a then-current study [Pew2014], 64% of users said concerns over privacy have increased, and 67% of mobile Internet users would like to do more to protect their privacy. The World Wide Web Consortium (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 maintain that users are the stakeholders with the highest priority, followed by implementors and other stakeholders, which further enforces 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.

Operators may deploy technologies that can either impact user privacy without being aware of those privacy implications or incorrectly assume that the benefits users gain from the new technology outweigh the loss of privacy. 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, such as cookie use and JavaScript injection, have been abused by these parties. This has caused some developers to encrypt content to circumvent these technologies that are seen as intrusive or bad for user privacy.

If users and content providers are to opt in to network management services with negative privacy impacts, they should see clear value from using these services and understand the impacts of using these services. Users should also have easy abilities to opt out. Some users will always automatically click through consent requests, so any model relying on explicit consent is flawed for these users. Understanding the extent of "auto click-through" may improve decisions about the use of consent requests in the future. One 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 at endpoints.

3. Network or Transport Solution Sessions

Network or Transport Solution Sessions discussed proposed solutions for managing encrypted traffic on radio access networks. Most solutions focus on metadata sharing, whether this sharing takes place from the endpoint to the network, from the network to the endpoint, or cooperatively in both directions. Transport-layer protocol evolution could be another approach to solve some of the issues radio access networks experience, which cause them to rely on network management middleboxes. By removing problems at the transport layer, reliance on expensive and complex middleboxes could decrease.

3.1. Sending Data Up/Down for Network Management Benefits

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

- Mobile Throughput Guidance [MTG]: exchanges metadata between network elements and endpoints via TCP options. It also allows for better understanding of how the transport protocol behaves and further improves the user experience, although additional work on MTG is still required.

- Session Protocol for User Datagrams [SPUD]: a UDP-based encapsulation protocol to allow explicit cooperation with middleboxes while using, new encrypted transport protocols.

- Network Status API: an API for operators to share congestion status or the state of a cell before an application starts sending data that could allow applications to change their behavior.

- Traffic Classification: classifying traffic and adding these classifications as metadata for analysis throughout the network. This idea has trust and privacy implications.

- Congestion Exposure [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.

- Latency versus Bandwidth: a bit that allows the content provider to indicate whether higher bandwidth or lower latency is of greater priority and allows the network to react based on that indication. Where this bit resides in the protocol stack and how it is authenticated would need to be decided.

- No Network Management Tools: disabling all network management tools from the network and relying only on end-to-end protocols to manage congestion.

- Flow Queue Controlled Delay (FQ-CoDel) [FLOWQUEUE]: a hybrid packet scheduler / Active Queue Management (AQM) [RFC7567] algorithm aiming to reduce bufferbloat and latency. FQ-CoDel manages packets from multiple flows and reduces the impact of head-of-line blocking from bursty traffic.

Some of these suggestions rely on signaling from network elements to endpoints. Others aim to create "hop-by-hop" solutions, which could be more aligned with how congestion is managed today but with greater privacy implications.

Still others rely on signaling from endpoints to network elements. Some of these rely on implicit signaling and others on explicit signaling. Some workshop attendees agreed that relying on applications to explicitly declare the quality of service they require was not a good path forward given the lack of success with this model in the past.

3.1.1. Competition, Cooperation, and Mobile Network Complexities

One of the larger issues in sharing data about the problems encountered with encrypted traffic in wireless networks is the matter of competition; network operators are reluctant to relinquish data about their own networks because it contains information that is valuable to competitors, and application providers wish to protect their users and reveal as little information as possible to the network. Some people think that if middleboxes were authenticated and invoked explicitly, this would be an improvement over current transparent middleboxes that intercept traffic without endpoint consent. Some workshop attendees suggested any exchange of information should be bidirectional in an effort to improve cooperation between the elements. A robust incentive framework could provide a solution to these issues or at least help mitigate them.

The radio access network is complex because it must deal with a number of conflicting demands. Base stations reflect this environment, and information within these base stations can be of value to other entities on the path. Some workshop participants thought solutions for managing congestion on radio networks should involve the base station if possible. For instance, understanding how the radio resource controller and AQM [RFC7567] 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 categorization and on providing this information to the protocols or endpoints. Completely categorizing radio networks could be impossible due to their complexity, but categorizing essential network properties could be possible and valuable.

4. Transport Layer: Issues, Optimization, and Solutions

TCP has been the dominant transport protocol since TCP/IP replaced the Network Control Protocol (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 the Stream Control Transmission Protocol (SCTP) and Datagram Congestion Control Protocol (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 IAB’s IP Stack Evolution program, and the first step in this program 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. For mostly competitive reasons, network operators do not tend to reveal network information and so are unlikely to donate this information freely to the IETF. The GSMA is in a position to try to collect this data and anonymize it before bringing it to IETF, which should alleviate the network operator worries but still provide IETF with some usable data.

Although congestion is only detected when packet loss is encountered and better methods based on detecting congestion would be beneficial, a considerable amount of work has already been done on TCP, especially innovation in bandwidth management and congestion control.

Furthermore, although the deficiencies of TCP are often considered key issues in the evolution of the Internet protocol stack, the main route to resolve these issues may not be a new TCP, but an evolved stack. Some workshop participants suggested that SPUD [SPUD] and Information-Centric Networking (ICN) [RFC7476] may help here. Quick UDP Internet Connection [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.

5. Application-Layer Optimization, 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 regarding their experiences running over mobile networks, the quality of experience their users expect, and the content and caching that providers 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 standards 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 centers to report issues that are difficult to test for and resolve.

Content providers cited other common issues concerning data traffic over mobile networks. Data subscription limits (known as "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. Developers build products on networks not indicative of the networks their customers are using, and not every organization has the finances to build a caching infrastructure.

Strongly related to content providers, content owners consider CDNs to be trusted deliverers of content, and CDNs have shown great success in fixed networks. Now that more traffic is moving to mobile networks, there is a need to place caches near the user at the edge of the mobile network. Placing caches at the edge of the mobile network is a solution, but it requires standards developed by content providers and mobile network operators. The IETF’s CDN

Interconnection [CDNI] Working Group aims to allow global CDNs to interoperate with mobile CDNs, but this causes huge issues for the caching of encrypted data between these CDNs. Some CDNs are experimenting with approaches like "Keyless SSL" [KeylessSSL] to enable safer storage of content without passing private keys to the CDN. Blind Caching [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, each panelist was asked to identify one key collaborative work item. Work items named were: evolving 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 troubleshooting 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 the Chatham House Rule. The session aimed to discuss regulatory and political issues, but not their worth or need, and to understand the laws that exist and how technologists can properly respond to them.

Mobile networks are regulated; compliance is mandatory and can incur costs on the mobile network operator, while non-compliance can result in service license revocation in some nations. Regulation does vary geographically. Some regulations are court orders and others are self-imposed regulations, for example, "block lists" of websites such as the Internet Watch Foundation [IWF] list. Operators are not expected to decrypt sites, so those encrypted sites will not be blocked because of content.

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 Internet of Things (IoT) cases. Most attendees agreed that if a government cannot get information it needs (and is legally entitled to have) from network operators, they will approach content providers. Some governments are aware of the impact of encryption and are working with, or trying to work with, content providers. The IAB has concluded that blocking and filtering can be done at the endpoints of the communication.

Not all of these regulations 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 together to alleviate the risk imposed by encrypted traffic. Some participants expressed concern that governments might require operators to provide information that they no longer have the ability to provide because previously unencrypted traffic is now being encrypted, and this might expose operators to new liability, but no specific examples were given during the workshop. A suggestion from some attendees was that if any new technical solutions are necessary, they should easily be "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. Suggested Principles and Solutions

Based on the talks and discussions throughout the workshop, a set of suggested principles and solutions has been collected. This is not an exhaustive list, and no attempt was made to come to consensus during the workshop, so there are likely at least some participants who would not agree with any particular principle listed below. The list is a union of participant thinking, not an intersection.

- Encrypted Traffic: Any solution should encourage and support encrypted traffic.

- Flexibility: Radio access network qualities vary vastly, and the network needs of content can differ significantly, so any new solution should be flexible across either the network type, content type, or both.

- Privacy: New solutions should not introduce new ways for information to be discovered and attributed to individual users.

- Minimum data only for collaborative work: User data, application data, and network data all need protection, so new solutions should use minimal information to make a working solution.

A collection of solutions suggested by various participants during the workshop is given below. Inclusion in this list does not imply that other workshop participants agreed. Again, the list is a union of proposed solutions, not an intersection.

- Evolving TCP or evolution on the transport layer: This could take a number of forms, and some of this work is already underway within the IETF.

- Congestion Control: Many attendees cited congestion control as a key issue. Further analysis, investigation, and work could be done in this space.

- Sprout [SPROUT]: Researched at MIT, Sprout is a transport protocol for applications that desire high throughput and low delay.

- PCC [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.

- CDNs and Caches: This suggests that placing caches closer to the edge of the radio network, as close as possible to the mobile user, or making more intelligent CDNs, would result in faster content delivery and less strain on the network.

- Blind Caching [BLIND_CACHING]: This is a proposal for caching of encrypted content.

- CDN Improvements: This includes Keyless SSL and better CDN placement.

- Mobile Throughput Guidance [MTG]: This is a mechanism and protocol elements that allow the cellular network to provide near real-time information on capacity available to the TCP server.

- One Bit for Latency / Bandwidth Trade-Off: This suggests determining whether using a single bit in an unencrypted transport header to distinguish between traffic that the sender prefers to be queued and traffic that the sender would prefer to drop rather than delay provides additional benefits beyond what can be achieved without this signaling.

- 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 a place to host solutions, or data from the base station could help in devising new solutions.

- Identify Traffic Types via 5-Tuple: Information from the 5-tuple could provide understanding of the traffic type, and network management appropriate for that traffic type could then be applied.

- Heuristics: Networks can sometimes identify traffic types by observing characteristics, such as data flow rate, and then apply network management to these identified flows. This is not recommended, as categorizations can be incorrect.

- 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 behavior. Alternatively, an API could provide the network with information on the data type, allowing appropriate network management for that data type; however, this method exposes privacy issues.

- Standard approach for the 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.

- AQM [RFC7567] and ECN [RFC3168] deployments: Queuing and congestion management methods have existed for some time in the form of AQM, ECN, and others, which can help the transport and Internet protocol layers adapt to congestion faster.

- 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.

- Keyless SSL [KeylessSSL]: This 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 standardized in the IETF. [LURK]

- Meaningful capacity sharing: This includes the ConEx [CONEX] work, which exposes information about congestion to the network nodes.

- Hop-by-hop: Some suggestions offer hop-by-hop methods that allow nodes to adapt flow given the qualities of the networks around them and the congestion they are experiencing.

- Metrics and metric standards: In order to evolve current protocols to be best suited to today’s networks, data is needed about current network conditions, protocol deployments, packet traces, and middlebox behavior. Beyond this, proper testing and debugging on networks could provide great insight for stack evolution.

- 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 understanding could help the standards process. These were identified as:

- greater understanding of the RAN within the IETF;
- reviews and comments on 3GPP perspective; and,
- how to do congestion control in the RAN.

7.1. Better Collaboration

Throughout the workshop, attendees placed emphasis on the need for better collaboration between the IETF and telecommunications bodies and organizations. 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 that are better for both technologies and users.

8. Since MaRNEW

Since MaRNEW, a number of activities have taken place in various IETF working groups and in groups external to IETF. The Alternatives to Content Classification for Operator Resource Deployment (ACCORD) BoF was held at IETF 95 in November 2015, which brought the workshop discussion to the wider IETF audiences by providing an account of the discussions that had taken place within the workshop and highlighting key areas to progress on. Key areas to progress on and an update on their current status are as follows:

- The collection of usable metrics and data were requested by a number of MaRNEW attendees, especially for use within the IRTF Measurement and Analysis for Protocols (MAP) Research Group; this data has been difficult to collect due to the closed nature of mobile network operators.
- Understanding impediments to protocol stack evolution has continued within the IAB’s IP Stack Evolution program and throughout transport-related IETF working groups such as the Transport Area Working Group (TSVWG).
- The Mobile Throughput Guidance document [MTG] has entered into a testing and data collection phase, although further advancements in transport technologies (QUIC, among others) may have stalled efforts in TCP-related proposals.

- Work on proposals for caching encrypted content continue, albeit with some security flaws that proponents are working on further proposals to fix. Most often, these are discussed within the IETF HTTPbis Working Group.

- The Path Layer UDP Substrate (PLUS) BOF at IETF 96 in July 2016 did not result in the formation of a working group, as attendees expressed concern on the privacy issues associated with the proposed data-sharing possibilities of the shim layer.

- The Limited Use of Remote Keys (LURK) BOF at IETF 96 in July 2016 did not result in the formation of a working group because the BOF identified more problems with the presumed approach than anticipated.

The most rewarding output of MaRNEW is perhaps the most intangible. MaRNEW gave two rather divergent industry groups the opportunity to connect and discuss common technologies and issues affecting users and operations. Mobile network providers and key Internet engineers and experts have developed a greater collaborative relationship to aid development of further standards that work across networks in a secure manner.

9. Security Considerations

This document is an IAB report from a workshop on interactions between network security, especially privacy, and network performance.

It does not affect the security of the Internet, taken on its own.

10. IANA Considerations

This document has no IANA actions.

11. Informative References

[BLIND_CACHING]
Thomson, M., Eriksson, G., and C. Holmberg, "Caching Secure HTTP Content using Blind Caches", Work in Progress, draft-thomson-http-bc-01, October 2016.

[CDNI] IETF, "Content Delivery Networks Interconnection (cdni)", <https://datatracker.ietf.org/wg/cdni/charter/>.

[CHATHAM_HOUSE_RULE]
Chatham House, "Chatham House Rule | Chatham House", <https://www.chathamhouse.org/about/chatham-house-rule>.

[CONEX] IETF, "Congestion Exposure (conex) - Documents", <https://datatracker.ietf.org/wg/conex/documents/>.

[EffectEncrypt] Xiong, C. and M. Patel, "The effect of encrypted traffic on the QoS mechanisms in cellular networks", August 2015, <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_25.pdf>.

[FLOWQUEUE] Hoeiland-Joergensen, T., McKenney, P., Taht, D., Gettys, J., and E. Dumazet, "FlowQueue-Codel", Work in Progress, draft-hoeiland-joergensen-aqm-fq-codel-01, November 2014.

[GSMA] GSMA, "GSMA Homepage", <http://gsma.com>.

[IWF] IWF, "Internet Watch Foundation Homepage", <https://www.iwf.org.uk/>.

[KeylessSSL] Sullivan, N., "Keyless SSL: The Nitty Gritty Technical Details", September 2014, <https://blog.cloudflare.com/keyless-ssl-the-nitty-gritty-technical-details/>.

[LURK] Migault, D., Ma, K., Salz, R., Mishra, S., and O. Dios, "LURK TLS/DTLS Use Cases", Work in Progress, draft-mglt-lurk-tls-use-cases-02, June 2016.

[MARNEW] IAB, "Managing Radio Networks in an Encrypted World (MaRNEW) Workshop 2015", <https://www.iab.org/activities/workshops/marnew/>.

[MTG] Jain, A., Terzis, A., Flinck, H., Sprecher, N., Arunachalam, S., Smith, K., Devarapalli, V., and R. Yanai, "Mobile Throughput Guidance Inband Signaling Protocol", Work in Progress, draft-flinck-mobile-throughput-guidance-04, March 2017.

[NETWORK_MANAGEMENT] Smith, K., "Network management of encrypted traffic", Work in Progress, draft-smith-encrypted-traffic-management-05, May 2016.

[NOTE_WELL] IETF, "IETF Note Well", <https://www.ietf.org/about/note-well.html>.

[PCC] Dong, M., Li, Q., Zarchy, D., Brighten Godfrey, P., and M. Schapira, "PCC: Re-architecting Congestion Control for Consistent High Performance", Proceedings of the 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI '15), USENIX Association, May 2015, <https://www.usenix.org/system/files/conference/nsdi15/nsdi15-paper-dong.pdf>.

[PCC-QOS] 3GPP, "Policy and charging control signalling flows and Quality of Service (QoS) parameter mapping", 3GPP TS 29.213, version 15.3.0, Release 15, June 2018, <http://www.3gpp.org/DynaReport/29213.htm>.

[Pew2014] Madden, M., "Public Perceptions of Privacy and Security in the Post-Snowden Era", November 2014, <http://www.pewinternet.org/2014/11/12/public-privacy-perceptions/>.

[QUIC] Hamilton, R., Iyengar, J., Swett, I., and A. Wilk, "QUIC: A UDP-Based Secure and Reliable Transport for HTTP/2", Work in Progress, draft-tsvwg-quic-protocol-02, January 2016.

[RFC2804] IAB and IESG, "IETF Policy on Wiretapping", RFC 2804, DOI 10.17487/RFC2804, May 2000, <https://www.rfc-editor.org/info/rfc2804>.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <https://www.rfc-editor.org/info/rfc3168>.

[RFC7476] Pentikousis, K., Ed., Ohlman, B., Corujo, D., Boggia, G., Tyson, G., Davies, E., Molinaro, A., and S. Eum, "Information-Centric Networking: Baseline Scenarios", RFC 7476, DOI 10.17487/RFC7476, March 2015, <https://www.rfc-editor.org/info/rfc7476>.

[RFC7567] Baker, F., Ed. and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <https://www.rfc-editor.org/info/rfc7567>.

[RFC8404] Moriarty, K., Ed. and A. Morton, Ed., "Effects of Pervasive Encryption on Operators", RFC 8404, DOI 10.17487/RFC8404, July 2018, <https://www.rfc-editor.org/info/rfc8404>.

[SDO_3GPP] 3GPP, "3GPP Homepage", <http://www.3gpp.org/>.

[SPROUT] Winstein, K., Sivaraman, A., and H. Balakrishnan, "Stochastic Forecasts Achieve High Throughput and Low Delay over Cellular Networks", 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI '13), USENIX Association, April 2013, <https://www.usenix.org/system/files/conference/nsdi13/nsdi13-final113.pdf>.

[SPUD] IETF, "Session Protocol for User Datagrams (spud)", <https://datatracker.ietf.org/wg/spud/about/>.

[STATE_BROWSER] Barnes, R., "Some observations of TLS in the web", July 2015, <https://www.ietf.org/proceedings/93/slides/slides-93-saag-3.pdf>.

[STATE_SERVER] Salz, R., "Some observations of TLS in the web", July 2015, <https://www.ietf.org/proceedings/93/slides/slides-93-saag-4.pdf>.

[TCPINC] "TCP Increased Security (tcpinc)", <https://datatracker.ietf.org/wg/tcpinc/charter/>.

Appendix A. Workshop Attendees

- Rich Salz, Akamai
- Aaron Falk, Akamai
- Vinay Kanitkar, Akamai
- Julien Maisonneuve, Alcatel Lucent
- Dan Druta, AT&T
- Humberto La Roche, Cisco
- Thomas Anderson, Cisco
- Paul Polakos, Cisco
- Marcus Ihlar, Ericsson
- Szilveszter Nadas, Ericsson
- John Mattsson, Ericsson
- Salvatore Loreto, Ericsson
- Blake Matheny, Facebook
- Andreas Terzis, Google
- Jana Iyengar, Google
- Natasha Rooney, GSMA
- Istvan Lajtos, GSMA
- Emma Wood, GSMA
- Jianjie You, Huawei
- Chunshan Xiong, Huawei
- Russ Housley, IAB
- Mary Barnes, IAB
- Joe Hildebrand, IAB / Cisco

- Ted Hardie, IAB / Google
- Robert Sparks, IAB / Oracle
- Spencer Dawkins, IETF AD
- Benoit Claise, IETF AD / Cisco
- Kathleen Moriarty, IETF AD / EMC
- Barry Leiba, IETF AD / Huawei
- Ben Campbell, IETF AD / Oracle
- Stephen Farrell, IETF AD / Trinity College Dublin
- Jari Arkko, IETF Chair / Ericsson
- Karen O'Donoghue, ISOC
- Phil Roberts, ISOC
- Olaf Kolkman, ISOC
- Christian Huitema, Microsoft
- Patrick McManus, Mozilla
- Dirk Kutscher, NEC Europe Network Laboratories
- Mark Watson, Netflix
- Martin Peylo, Nokia
- Mohammed Dadas, Orange
- Diego Lopez, Telefonica
- Matteo Varvello, Telefonica
- Zubair Shafiq, The University of Iowa
- Vijay Devarapalli, Vasona Networks
- Sanjay Mishra, Verizon
- Gianpaolo Scassellati, Vimplecom

- Kevin Smith, Vodafone

- Wendy Seltzer, W3C

Appendix B. Workshop Position Papers

- Mohammed Dadas, Emile Stephan, Mathilde Cayla, Iuniana Oprescu, "Cooperation Framework between Application layer and Lower Layers" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_33.pdf>

- Julien Maisonneuve, Vijay Gurbani, and Thomas Fossati, "The security pendulum" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_4.pdf>

- Martin Peylo, "Enabling Secure QoE Measures for Internet Applications over Radio Networks is a MUST" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_32.pdf>

- Vijay Devarapalli, "The Bandwidth Balancing Act: Managing QoE as encrypted services change the traffic optimization game" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_10.pdf>

- Humberto J. La Roche, "Use Cases for Communicating End-Points in Mobile Network Middleboxes" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_12.pdf>

- Patrick McManus and Richard Barnes, "User Consent and Security as a Public Good" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_13.pdf>

- Iuniana Oprescu, Jon Peterson, and Natasha Rooney, "A Framework for Consent and Permissions in Mediating TLS" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_31.pdf>

- Jari Arkko and Goran Eriksson, "Characteristics of Traffic Type Changes and Their Architectural Implications" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_15.pdf>

- Szilveszter Nadas and Attila Mihaly, "Concept for Cooperative Traffic Management" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_16.pdf>

- Gianpaolo Scassellati, "Vimpelcom Position paper for MaRNEW Workshop" at <https://www.iab.org/wp-content/IAB-uploads/2015/09/MaRNEW_1_paper_17.pdf>

- Mirja Kuhlewind, Dirk Kutscher, and Brian Trammell, "Enabling Traffic Management without DPI" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_18.pdf>

- Andreas Terzis and Chris Bentzel, "Sharing network state with application endpoints" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_19.pdf>

- Marcus Ihlar, Salvatore Loreto, and Robert Skog, "The needed existence of PEP in an encrypted world" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_20.pdf>

- John Mattsson, "Network Operation in an All-Encrypted World" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_21.pdf>

- Dirk Kutscher, Giovanna Carofiglio, Luca Muscariello, and Paul Polakos, "Maintaining Efficiency and Privacy in Mobile Networks through Information-Centric Networking" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_23.pdf>

- Chunshan Xiong and Milan Patel, "The effect of encrypted traffic on the QoS mechanisms in cellular networks" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_25.pdf>

- Thomas Anderson, Peter Bosch, and Alessandro Duminuco, "Bandwidth Control and Regulation in Mobile Networks via SDN/NFV-Based Platforms" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_26.pdf>

- Karen O’Donoghue and Phil Roberts, "Barriers to Deployment: Probing the Potential Differences in Developed and Developing Infrastructure" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_27.pdf>

- Wendy Seltzer, "Security, Privacy, and Performance Considerations for the Mobile Web" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_28.pdf>

- Jianjie You, Hanyu Wei, and Huaru Yang, "Use Case Analysis and Potential Bandwidth Optimization Methods for Encrypted Traffic" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_29.pdf>

- Mangesh Kasbekar and Vinay Kanitkar, "CDNs, Network Services and Encrypted Traffic" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_30.pdf>

- Yves Hupe, Claude Rocray, and Mark Santelli, "Providing Optimization of Encrypted HTTP Traffic" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_341.pdf>

- M. Zubair Shafiq, "Tracking Mobile Video QoE in the Encrypted Internet" at <https://www.iab.org/wp-content/IAB-uploads/2015/08/MaRNEW_1_paper_35.pdf>

- Kevin Smith, "Encryption and government regulation: what happens now?" at <https://www.iab.org/wp-content/IAB-uploads/2015/09/MaRNEW_1_paper_1.pdf>

Acknowledgements

Stephen Farrell reviewed this report in draft form and provided copious comments and suggestions.

Barry Leiba provided some clarifications on specific discussions about Lawful Intercept that took place during the workshop.

Bob Hinden and Warren Kumari provided comments and suggestions during the IAB Call for Comments.

Amelia Andersdotter and Shivan Kaul Sahib provided comments from the Human Rights Review Team during the IAB Call for Comments.

Authors' Addresses

Natasha Rooney
GSMA

Email: nrooney@gsma.com
URI: https://gsma.com

Spencer Dawkins (editor)
Wonder Hamster

Email: spencerdawkins.ietf@gmail.com