RFC 9490
Report from the IAB Workshop on Management Techniques in Encrypted Networks (M-TEN)

Abstract

The "Management Techniques in Encrypted Networks (M-TEN)" workshop was convened by the Internet Architecture Board (IAB) from 17 October 2022 to 19 October 2022 as a three-day online meeting. The workshop was organized in three parts to discuss ways to improve network management techniques in support of even broader adoption of encryption on the Internet. This report summarizes the workshop's discussion and identifies topics that warrant future work and consideration.

Note that this document is a report on the proceedings of the workshop. The views and positions documented in this report are those of the expressed during the workshop by participants and do not necessarily reflect IAB views and positions.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

The Internet Architecture Board (IAB) holds occasional workshops designed to consider long-term issues and strategies for the Internet, and to suggest future directions for the Internet architecture. This long-term planning function of the IAB is complementary to the ongoing engineering efforts performed by working groups of the Internet Engineering Task Force (IETF).

User privacy and security are constantly being improved by increasingly strong and more widely deployed encryption. This workshop aims to discuss ways to improve network management techniques in support of even broader adoption of encryption on the Internet.

Network management techniques need to evolve to work effectively and reliably in the presence of ubiquitous traffic encryption and to support user privacy. In an all-encrypted network, it is not viable to rely on unencrypted metadata for network monitoring and security functions, troubleshooting devices, and passive traffic measurements. New approaches are needed to track network behaviors, e.g., by directly cooperating with endpoints and applications, increasing use of in-band telemetry, increasing use of active measurement approaches, and privacy-preserving inference techniques.

The aim of this IAB online workshop from October 17-19, 2022, has been to provide a platform to explore the interaction between network management and traffic encryption and to initiate work on collaborative approaches that promote security and user privacy while supporting operational requirements. As such, the workshop addressed the following questions:

- What are actionable network management requirements?
- Who is willing to work on collaborative solutions?
- What are the starting points for collaborative solutions?

1.1. About This Workshop Report Content

This document is a report on the proceedings of the workshop. The views and positions documented in this report are those of the workshop participants and do not necessarily reflect IAB views and positions.

Furthermore, the content of the report comes from presentations given by workshop participants and notes taken during the discussions, without interpretation or validation. Thus, the content of this report follows the flow and dialog of the workshop but does not attempt to capture a consensus.

2. Workshop Scope and Discussion

The workshop was held across three days with all-group discussion slots, one per day. The following topic areas were identified, and the program committee organized paper submissions into three main themes for each of the three discussion slots. During each discussion, those papers were presented sequentially with open discussion held at the end of each day.

2.1. "Where We Are" - Requirements and Passive Observations

The first day of the workshop focused on the existing state of the relationship between network management and encrypted traffic from various angles. Presentations ranged from discussing classifiers using machine learning to recognize traffic, to advanced techniques for evading traffic analysis, to user privacy considerations.

After an introduction that covered the goals of the workshop and the starting questions (as described in Section 1), there were four presentations followed by open discussion.

2.1.1. Traffic Classification and Network Management

Many existing network management techniques are passive in nature: they don't rely on explicit signals from end hosts to negotiate with network middleboxes but instead rely on inspecting packets to recognize traffic and apply various policies. Traffic classification, as a passive technique, is being challenged by increasing encryption.

Traffic classification is commonly performed by networks to infer what applications and services are being used. This information is in turn used for capacity and resource planning, Quality-of-Service (QoS) monitoring, traffic prioritization, network access control, identity management, and malware detection. However, since classification commonly relies on recognizing unencrypted properties of packets in a flow, increasing encryption of traffic can decrease the effectiveness of classification.

The amount of classification that can be performed on traffic also provides useful insight into how "leaky" the protocols used by applications are and points to areas where information is visible to any observer, who may or may not be malicious.

Frequently, classification has been based on specific rules crafted by experts, but there is also a move toward using machine learning to recognize patterns. "Deep learning" machine-learning models generally rely on analyzing a large set of traffic over time and have trouble reacting quickly to changes in traffic patterns.

Models that are based on closed-world data sets also become less useful over time as traffic changes. [JIANG] describes experiments that show that a model that performed with high accuracy on an initial data set becomes severely degraded when running on a newer data set that contains traffic from the same applications. Even in as little time as one week, the traffic classification would become degraded. However, the set of features in packets and flows that were useful for models stayed mostly consistent, even if the models themselves needed to be

updated. Models where the feature space is reduced to fewer features showed better resiliency and could be retrained more quickly. Based on this, [JIANG] recommends more work and research to determine which set of features in IP packets are most useful for focused machine-learning analysis. [WU] also recommends further research investment in Artificial Intelligence (AI) analysis for network management.

2.1.2. Preventing Traffic Analysis

Just as traffic classification is continually adapting, techniques to prevent traffic analysis and to obfuscate application and user traffic are continually evolving. An invited talk from the authors of [DITTO] shared a novel approach with the workshop for how to build a very robust system to prevent unwanted traffic analysis.

Usually traffic obfuscation is performed by changing the timing of packets or by adding padding to data. The practices can be costly and negatively impact performance. [DITTO] demonstrated the feasibility of applying traffic obfuscation on aggregated traffic in the network with minimal overhead and inline speed.

While traffic obfuscation techniques are not widely deployed today, this study underlines the need for continuous effort to keep traffic models updated over time, the challenges of the classification of encrypted traffic, as well as the opportunities to further enhance user privacy.

2.1.3. Users and Privacy

The Privacy Enhancements and Assessments Research Group (PEARG) is working on a document to discuss guidelines for measuring traffic on the Internet in a safe and privacy-friendly way [LEARMONTH]. These guidelines and principles provide another view on the discussion of passive classification and analysis of traffic.

Consent for collection and measurement of metadata is an important consideration in deploying network measurement techniques. This consent can be given explicitly as informed consent, given by proxy, or may be only implied. For example, a user of a network might need to consent to certain measurement and traffic treatment when joining a network.

Various techniques for data collection can also improve user privacy, such as discarding data after a short period of time, masking aspects of data that contain user-identifying information, reducing the accuracy of collected data, and aggregating data.

2.1.4. Discussion

The intents and goals of users, application developers, and network operators align in some cases, but not in others. One of the recurring challenges that was discussed was the lack of a clear way to understand or to communicate intents and requirements. Both traffic classification and traffic obfuscation attempt to change the visibility of traffic without cooperation of other parties: traffic classification is an attempt by the network to inspect application traffic without coordination from applications, and traffic obfuscation is an attempt by the application to hide that same traffic as it transits a network.

Traffic adaptation and prioritization is one dimension in which the incentives for cooperation seem most clear. Even if an application is trying to prevent the leaking of metadata, it could benefit from signals from the network about sudden capacity changes that can help it adapt its application quality, such as bitrates and codecs. Such signaling may not be appropriate for the most privacy-sensitive applications, like Tor, but could be applicable for many others. There are existing protocols that involve explicit signaling between applications and networks, such as Explicit Congestion Notification (ECN) [RFC3168], but that has yet to see wide adoption.

Managed networks (such as private corporate networks) were brought up in several comments as particularly challenging for meeting management requirements while maintaining encryption and privacy. These networks can have legal and regulated requirements for detection of specific fraudulent or malicious traffic.

Personal networks that enable managed parental controls have similar complications with encrypted traffic and user privacy. In these scenarios, the parental controls that are operated by the network may be as simple as a DNS filter, which can be made ineffective by a device routing traffic to an alternate DNS resolver.

2.2. "Where We Want to Go" - Collaboration Principles

The second day of the workshop focused on the emerging techniques for analyzing, managing, or monitoring encrypted traffic. Presentations covered advanced classification and identification, including machine-learning techniques, for the purposes of managing network flows or monitoring or monetizing usage.

After an introduction that covered the goals of the workshop and the starting questions (as described in Section 1), there were three presentations, followed by open discussion.

2.2.1. First-Party Collaboration for Network Management

It is the intent of end-to-end encryption of traffic to create a barrier between entities inside the communication channel and everyone else, including network operators. Therefore, any attempt to overcome that intentional barrier requires collaboration between the inside and outside entities. At a minimum, those entities must agree on the benefits of overcoming the barrier (or solving the problem), agree that costs are proportional to the benefits, and agree to additional limitations or safeguards against bad behavior by collaborators including other non-insiders [BARNES].

The Internet is designed interoperably, which means an outside entity wishing to collaborate with the inside might be any number of intermediaries and not, say, a specific person that can be trusted in the human sense. Additionally, the use of encryption, especially network-layer or transport-layer encryption, introduces dynamic or opportunistic or perfunctory discoverability. These realities point to a need to ask why an outside entity might make an engineering case to collaborate with the user of a network with encrypted traffic and to ask whether the trade-offs and potential risks are worth it to the user.

However, the answers cannot be specific, and the determinations or guidance need to be general as the encryption boundary is inevitably an application used by many people. Trade-offs must make sense to users who are unlikely to be thinking about network management considerations. Harms need to be preemptively reduced because, in general terms, few users would choose network management benefits over their own privacy if given the choice.

Some have found that there appears to be little, if any, evidence that encryption causes network problems that are meaningful to the user. Since alignment on problem solving is a prerequisite to collaboration on a solution, it does not seem that collaboration across the encryption boundary is called for.

2.2.2. Second- and Third-Party Collaboration for Network Management

Even with the wide-scale deployment of encryption in new protocols and of techniques that prevent passive observers of network traffic from knowing the content of exchanged communications, important information, such as which parties communicate and sometimes even which services have been requested, may still be able to be deduced. The future is to conceal more data and metadata from passive observers and also to minimize information exposure to second parties (where the user is the first party) by, maybe counterintuitively, introducing third-party relay services to intermediate communications. As discussed in [KUEHLEWIND], the relay is a mechanism that uses additional levels of encryption to separate two important pieces of information: knowledge of the identity of the person accessing a service is separated from knowledge about the service being accessed. By contrast, a VPN uses only one level of encryption and does not separate identity (first party) and service (second party) metadata.

Relay mechanisms are termed "oblivious", there is a future for specifications in privacy-preserving measurement (PPM), and protocols like Multiplexed Application Substrate over QUIC Encryption (MASQUE) are discussed in the IETF. In various schemes, users are ideally able to share their identity only with the entity they have identified as a trusted one. That data is not shared with the service provider. However, this is more complicated for network management, but there may be opportunities for better collaboration between the network and, say, the application or service at the endpoint.

A queriable relay mechanism could preserve network management functions that are disrupted by encryption, such as TCP optimization, quality of service, zero-rating, parental controls, access control, redirection, content enhancement, analytics, and fraud prevention. Instead of encrypting communication between only two ends with passive observation by all on-path elements, intermediate relays could be introduced as trusted parties that get to see limited information for the purpose of collaboration between in-network intermediary services.

2.2.3. Visible, Optional Network Management

Out of all of the possible network management functions that might be ameliorated by proxying, the ability to control congestion in encrypted communications has been researched in depth. These techniques are realized based on TCP performance-enhancing proxies (PEPs) that either entirely intercept a TCP connection or interfere with the transport information in the TCP header. However, despite the challenge that the new encrypted protocol will limit any such in-

network interference, these techniques can also have a negative impact on the evolvability of these protocols. Therefore, a new approach was presented where, instead of manipulating existing information, additional information is sent using a so-called sidecar protocol independent of the main transport protocol that is used end to end [WELZL]. For example, sidecar information can contain additional acknowledgments to enable in-network local retransmission or faster end-to-end retransmission by reducing the signaling round-trip time.

Taking user privacy benefits for granted, there is a need to investigate the comparable performance outputs of various encrypted traffic configurations such as the use of an additional sidecar protocol, or explicit encrypted and trusted network communication using MASQUE in relation to existing techniques such as TCP PEPs, etc.

2.2.4. Discussion

One size fits all? On the issue of trust, different networks or devices will have different trust requirements for devices, users, or each other, and vice versa. For example, imagine two networks with really different security requirements, like a home network with a requirement to protect its child users versus a national security institution's network. How could one network architecture solve the needs of all use cases?

Does our destination have consequences? It seems sometimes that there may be future consequences caused by the ubiquitous, strong encryption of network traffic because it will cause intermediaries to poke holes in what are supposed to be long-term solutions for user privacy and security.

Can we bring the user along? While there has been a focus on the good reasons why people might collaborate across the encryption barrier, there will always be others who want to disrupt that in order to exploit the data for their own gain, and sometimes exploitation is called innovation. High-level policy mitigations have exposed how powerless end users are against corporate practices of data harvesting. And yet interfaces to help users understand these lower-layer traffic flows to protect their financial transactions or privacy haven't been achieved yet. That means that engineers must make inferences about user wants. Instead, we should make these relationships and trade-offs more visible.

2.3. "How We Get There" - Collaboration Use Cases

The third day focused on techniques that could be used to improve the management of encrypted networks.
The potential paths forward described in the presentations included some element of collaboration between the networks and the subscribing clients that simultaneously want both privacy and protection. Thus, the central theme of the third day became negotiation and collaboration.

2.3.1. Establishing Expected Contracts to Enable Security Management

For enterprise networks where client behavior is potentially managed, [COLLINS] proposes "Improving network monitoring through contracts", where contracts describe different states of network behavior.

Because network operators have a limited amount of time to focus on problems and process alerts, contracts and states let the operator focus on a particular aspect of a current situation or problem. The current estimate for the number of events a Security Operations Center (SOC) operator can handle is about 10 per hour. Operators must work within the limits imposed by their organization and must pick among options that frequently only frustrate attackers -- preventing attacks entirely is potentially impossible. Finally, operators must prioritize and manage the most events possible.

Validating which alerts are true positives is challenging because lots of weird traffic creates many anomalies, and not all anomalies are malicious events. Identifying which anomalous traffic is rooted in malicious activity with any level of certainty is extremely challenging. Unfortunately, applying the latest machine-learning techniques has produced mixed results. To make matters worse, the large amounts of Internet-wide scanning has resulted in endless traffic that is technically malicious but only creates an information overload and challenges event prioritization. Any path forward must free up analyst time to concentrate on the more challenging events.

The proposed contract solution is to define a collection of acceptable behaviors that comprises different states that might include IP addresses, domain names, and indicators of compromise. Deviation from a contract might indicate that a system is acting outside a normal mode of behavior or even that a normal mode of behavior is suddenly missing. An example contract might be "this system is expected to update its base OS once a day". If this doesn't occur, then this expectation has not been met, and the system should be checked as it failed to call home to look for (potentially security-related) updates.

Within the IETF, the Manufacturer Usage Description Specification (MUD) [RFC8520] is one subset of contracts. Note that contracts are likely to succeed only in a constrained, expected environment maintained by operational staff and may not work in an open Internet environment where end users drive all network connections.

2.3.2. Zero-Knowledge Middleboxes

The world is not only shifting to increased encrypted traffic but is also encrypting more and more of the metadata (e.g., DNS queries and responses). This makes network policy enforcement by middleboxes significantly more challenging. The result is a significant tension between security enforcement and privacy protection.

Goals for solving this problem should include enabling networks to enforce their policies, but should not include the weakening of encryption nor the deployment of new server software. Existing solutions fail to meet at least one of these points.

A cryptographic principle of a "zero-knowledge proof" (ZKP) [GRUBBS] may be one path forward to consider. A ZKP allows a third party to verify that a statement is true without revealing what the statement actually is. Applying this to network traffic has been shown to allow a middlebox to verify that traffic to a web server is compliant with a policy without revealing the actual contents. This solution meets the three criteria listed above. Using ZKP within TLS 1.3 traffic turns out to be plausible.

An example engine using encrypted DNS was built to test ZKP. Clients were able to create DNS requests that were not listed within a DNS block list. Middleboxes could verify, without knowing the exact request, that the client's DNS request was not on the prohibited list. Although the result was functional, the computational overhead was still too slow, and future work will be needed to decrease the ZKP-imposed latencies.

2.3.3. Red Rover - a Collaborative Approach to Content Filtering

The principle challenge being studied is how to handle the inherent conflict between filtering and privacy. Network operators need to implement policies and regulations that can originate from many locations (e.g., security, governmental, parental, etc.). Conversely, clients need to protect their users' privacy and security.

Safe browsing, originally created by Google, is one example of a mechanism that tries to meet both sides of this conflict. It would be beneficial to standardize this and other similar mechanisms. Operating systems could continually protect their users by ensuring that malicious destinations are not being reached. This would require some coordination between cooperating clients and servers offering protection services. These collaborative solutions may be the best compromise to resolve the tension between privacy services and protection services [PAULY].

3. Conclusions

Looking forward, the workshop participants identified that solving the entire problem space with a single approach will be challenging for several reasons:

- The scalability of many solutions will likely be an issue as some solutions are complex or expensive to implement.
- Collaboration between multiple parties will be required for many mechanisms to function, and the sets of parties required for different use cases might be disjoint.
- There is an unanswered question of whether or not network operators are willing to participate by allowing new encryption technologies into their environment in exchange for technologies that prove their clients are being good net-citizens. If so, some of these solutions might be required to exist before networks allow a certain type of increased encryption; consider the example of TLS Encrypted Client Hello being blocked by some network operators.

The breadth of the problem space itself is another complicating factor. There is a wide variety of network architectures, and each has different requirements for both data encryption and network management. Each problem space will have multiple, different encumbrances: for example, technical, legal, data ownership, and regulatory concerns. New network architectures might be needed to solve this problem at a larger scope, which would in turn require interoperability support from network product vendors. Education about various solutions will be required in order to ensure regulation and policy organizations can understand and thus support the deployment of developed solutions.

After new technologies and related standards are developed and deployed, unintended consequences can emerge. These lead to effects in multiple directions: on one hand, exposed protocol values not intended for network management might be used by networks to differentiate traffic; on the other hand, changes to a protocol that break existing use cases might have an impact on private network deployments. While making decisions on technology direction and protocol design, it is important to consider the impact on various kinds of network deployments and their unique requirements. When protocols change to make different network management functions easier or harder, the impact on various deployment models ought to be considered and documented.

4. Informative References

[BARNES]   Barnes, R., "What's In It For Me? Revisiting the reasons people collaborate", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Barnes-Whats-In-It-For-Me-Revisiting-the-reasons-people-collaborate.pdf>.

[CASAS]   Casas, P., "Monitoring User-Perceived Quality in an Encrypted Internet - AI to the Rescue", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Casas-AI-driven-real-time-QoE-monitoring-encrypted-traffic.pdf>.

[COLLINS]   Collins, M., "Improving Network Monitoring Through Contracts", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Collins-Improving-Network-Monitoring-Through-Contracts.pdf>.

[DERI]   Deri, L., "nDPI Research Proposal", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Deri-nDPI-Research-Proposal.pdf>.

[DITTO]   Meier, R., Lenders, V., and L. Vanbever, "ditto: WAN Traffic Obfuscation at Line Rate", Network and Distributed Systems Security (NDSS) Symposium, DOI 10.14722/ndss.2022.24056, April 2022, <https://doi.org/10.14722/ndss.2022.24056>.

[ELKINS]   Elkins, N., Ackermann, M., Tahiliani, M., Dhody, D., and T. Pecorella, "Performance Monitoring in Encrypted Networks: PDMv2", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Elkins-Performance-Monitoring-in-Encrypted-Networks-PDMv2.pdf>.

[GRUBBS]   Grubbs, P., Arun, A., Zhang, Y., Bonneau, J., and M. Walfish, "Zero-Knowledge Middleboxes", 31st USENIX Security Symposium (USENIX Security 22), August 2022, <https://www.usenix.org/conference/usenixsecurity22/presentation/grubbs>.

[HARDAKER]   Hardaker, W., "Network Flow Management by Probability", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Hardaker-Encrypted-Traffic-Estimation.pdf>.

[JIANG]   Jiang, X., Liu, S., Naama, S., Bronzino, F., Schmitt, P., and N. Feamster, "Towards Designing Robust and Efficient Classifiers for Encrypted Traffic in the Modern Internet", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Jiang-Towards-Designing-Robust-and-Efficient-Classifiers-for-Encrypted-Traffic-in-the-Modern-Internet.pdf>.

[KNODEL]   Knodel, M., "(Introduction) 'Guidelines for Performing Safe Measurement on the Internet'", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Knodel-Guidelines-for-Performing-Safe-Measurement-on-the-Internet.pdf>.

[KUEHLEWIND]   Kuehlewind, M., Westerlund, M., Sarker, Z., and M. Ihlar, "Relying on Relays: The future of secure communication", June 2022, <https://www.ericsson.com/en/blog/2022/6/relays-and-online-user-privacy>.

[LEARMONTH]   Learmonth, I. R., Grover, G., and M. Knodel, "Guidelines for Performing Safe Measurement on the Internet", Work in Progress, Internet-Draft, draft-irtf-pearg-safe-internet-measurement-08, 10 July 2023, <https://datatracker.ietf.org/doc/html/draft-irtf-pearg-safe-internet-measurement-08>.

[LEI]   Lei, Y., Wu, J., Sun, X., Zhang, L., and Q. Wu, "Encrypted Traffic Classification Through Deep Learning", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Lei-Encrypted-Traffic-Classification-Through-Deep-Learning.pdf>.

[PAULY]   Pauly, T. and R. Barnes, "Red Rover: A collaborative approach to content filtering", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Pauly-Red-Rover-A-collaborative-approach-to-content-filtering.pdf>.

[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>.

[RFC8520]   Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <https://www.rfc-editor.org/info/rfc8520>.

[WELZL]   Welzl, M., "The Sidecar: 'Opting in' to PEP Functions", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Welzl-The-Sidecar-Opting-in-to-PEP-Functions.pdf>.

[WU]   Wu, Q., Wu, J., and Q. Ma, "Network Management of Encrypted Traffic: Detect it don't decrypt it", August 2022, <https://www.iab.org/wp-content/IAB-uploads/2023/11/Wu-mten-taxonomy.pdf>.

Appendix A. Position Papers

Interested participants were openly invited to submit position papers on the workshop topics, including Internet-Drafts, relevant academic papers, or short abstracts explaining their interest. The papers below constitute the inputs that were considered relevant for workshop attendees and that focused the discussions themselves. The program committee grouped the papers by theme.

A.1. Motivations and Principles

Richard Barnes. "What's In It For Me? Revisiting the reasons people collaborate." [BARNES]

Mallory Knodel. "(Introduction) 'Guidelines for Performing Safe Measurement on the Internet'." (Additional rationale) [KNODEL]

Qin Wu, Jun Wu, Qiufang Ma. "Network Management of Encrypted Traffic: Detect it don't decrypt it." [WU]

A.2. Classification and Identification of Encrypted Traffic

Luca Deri. "nDPI Research Proposal." [DERI]

Wes Hardaker. "Network Flow Management by Probability." [HARDAKER]

Xi Jiang, Shinan Liu, Saloua Naama, Francesco Bronzino, Paul Schmitt, Nick Feamster. "Towards Designing Robust and Efficient Classifiers for Encrypted Traffic in the Modern Internet." [JIANG]

Yupeng Lei, Jun Wu, Xudong Sun, Liang Zhang, Qin Wu. "Encrypted Traffic Classification Through Deep Learning." [LEI]

A.3. Ideas for Collaboration and Coordination between Devices and Networks

Michael Collins. "Improving Network Monitoring Through Contracts." [COLLINS]

Paul Grubbs, Arasu Arun, Ye Zhang, Joseph Bonneau, Michael Walfish. "Zero-Knowledge Middleboxes." [GRUBBS]

Mirja Kuehlewind, Magnus Westerlund, Zaheduzzaman Sarker, Marcus Ihlar. "Relying on Relays: The future of secure communication." [KUEHLEWIND]

Tommy Pauly, Richard Barnes. "Red Rover: A collaborative approach to content filtering." [PAULY]

Michael Welzl. "The Sidecar: 'Opting in' to PEP Functions." [WELZL]

A.4. Other Background Material

Pedro Casas. "Monitoring User-Perceived Quality in an Encrypted Internet - AI to the Rescue." [CASAS]

Nalini Elkins, Mike Ackermann, Mohit P. Tahiliani, Dhruv Dhody, Prof. Tommaso Pecorella. "Performance Monitoring in Encrypted Networks: PDMv2." [ELKINS]

Appendix B.   Workshop Participants

The workshop participants were Cindy Morgan, Colin Perkins, Cullen Jennings, Deborah Brungard, Dhruv Dhody, Éric Vyncke, Georg Carle, Ivan Nardi, Jari Arkko, Jason Livingood, Jiankang Yao, Karen O'Donoghue, Keith Winstein, Lars Eggert, Laurent Vanbever, Luca Deri, Mallory Knodel, Marcus Ihlar, Matteo, Michael Collins, Michael Richardson, Michael Welzl, Mike Ackermann, Mirja Kühlewind, Mohit P. Tahiliani, Nalini Elkins, Patrick Tarpey, Paul Grubbs, Pedro Casas, Qin Wu, Qiufang Ma, Richard Barnes, Rob Wilton, Russ White, Saloua Naama, Shinan Liu, Srinivas C, Toerless Eckert, Tommy Pauly, Wes Hardaker, Xi Chase Jiang, Zaheduzzaman Sarker, and Zhenbin Li.

Appendix C.   Program Committee

The workshop program committee members were Wes Hardaker (IAB, USC/ISI), Mallory Knodel (IAB, Center for Democracy and Technology), Mirja Kühlewind (IAB, Ericsson), Tommy Pauly (IAB, Apple), Russ White (IAB, Juniper), Qin Wu (IAB, Huawei).

IAB Members at the Time of Approval

Internet Architecture Board members at the time this document was approved for publication were:

> Dhruv Dhody
>
> Lars Eggert
>
> Wes Hardaker
>
> Cullen Jennings
>
> Mallory Knodel
>
> Suresh Krishnan
>
> Mirja Kühlewind
>
> Tommy Pauly
>
> Alvaro Retana
>
> David Schinazi
>
> Christopher Wood

Qin Wu

Jiankang Yao

Acknowledgments

Authors' Addresses

Mallory Knodel
Email: mknodel@cdt.org

Wes Hardaker
Email: ietf@hardakers.net

Tommy Pauly
Email: tpauly@apple.com