

TSVWG
Internet-Draft
Intended status: Informational
Expires: March 29, 2018

G. Fairhurst
University of Aberdeen
C.S. Perkins
University of Glasgow
September 27, 2017

The Impact of Transport Header Encryption on Operation and Evolution of
the Internet
draft-fairhurst-tsvwg-transport-encrypt-04

Abstract

This document describes implications of applying end-to-end encryption at the transport layer. It identifies some in-network uses of transport layer header information that can be used with a transport header integrity check. It reviews the implication of developing encrypted end-to-end transport protocols and examines the implication of developing and deploying encrypted end-to-end transport protocols. Since transport measurement and analysis of the impact of network characteristics have been important to the design of current transport protocols, it also considers some anticipated implications on transport and application evolution.

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 March 29, 2018.

Copyright Notice

Copyright (c) 2017 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	2
1.1. Current uses of Transport Headers within the Network	6
1.1.1. Observing Transport Information in the Network	7
1.1.1.1. Flow Identification	7
1.1.1.2. Metrics derived from Transport Layer Headers	7
1.1.1.3. Metrics derived from Network Layer Headers	10
1.1.2. Transport Measurement	12
1.1.2.1. Point of Measurement	12
1.1.2.2. Use by Operators to Plan and Provision Networks . .	13
1.1.2.3. Service Performance Measurement	13
1.1.2.4. Measuring Transport to Support Network Operations .	13
1.1.3. Use for Network Diagnostics and Troubleshooting	15
1.1.4. Observing Headers to Implement Network Policy	15
2. Encryption and Authentication of Transport Headers	15
2.1. Authenticating the Transport Protocol Header	17
2.2. Encrypting the Transport Payload	17
2.3. Encrypting the Transport Header	18
2.4. Authenticating Transport Information and Selectively Encrypting the Transport Header	18
2.5. Adding Transport Information to Network-Layer Protocol Headers	18
3. Implications of Protecting the Transport Headers	19
3.1. Independent Measurement	19
3.2. Characterising "Unknown" Network Traffic	20
3.3. Accountability and Internet Transport Protocols	20
3.4. Impact on Research, Development and Deployment	21
4. Acknowledgements	21
5. Security Considerations	22
6. IANA Considerations	22
7. References	22
7.1. Normative References	22
7.2. Informative References	22
Appendix A. Revision information	26
Authors' Addresses	27

1. Introduction

This document discusses the implications of end-to-end encryption applied at the transport layer, and examines the impact on transport protocol design, usage, and network operations and management. It also considers anticipated implications on transport and application evolution.

The transport layer provides the first end-to-end interactions across the Internet. Transport protocols layer directly over the network-layer service and are sent in the payload of network-layer packets. They support end-to-end communication between applications, supported by higher-layer protocols, running on the end systems (or transport endpoint). This simple architectural view hides one of the core functions of the transport, however - to discover and adapt to the properties of the Internet path that is currently being used. The design of Internet transport protocols is as much about trying to avoid the unwanted side effects of congestion on a flow and other capacity-sharing flows, avoiding congestion collapse, adapting to changes in the path characteristics, etc., as it is about end-to-end feature negotiation, flow control and optimising for performance of a specific application.

To achieve stable Internet operations the IETF transport community has to date relied heavily on measurement and insights of the network operations community to understand the trade-offs, and to inform selection of select appropriate mechanisms, to ensure a safe, reliable and robust Internet. In turn, the network operations community relies on being able to understand the traffic passing over the Internet, both in aggregate and at the flow level -- inspecting transport layer headers to help understand traffic dynamics.

There are many motivations for deploying encrypted transports, and encryption of transport payloads. The increasing public concerns about the interference with Internet traffic have led to a rapidly expanding deployment of encryption to protect end-user privacy, in protocols like QUIC. At the same time, network operators and access providers, especially in mobile networks, have come to rely on the in-network measurement of transport properties and the functionality provided by middleboxes to both support network operations and enhance performance.

This document considers some implications of working with encrypted transport protocols, and discusses trade-offs around authentication, encryption of transport protocol headers. It describes some of the architectural challenges and considerations in the way transport protocols are designed when using encryption [Measure].

Encryption of the transport layer brings some well-known privacy and security benefits, but also introduces various costs that need to be considered. Specifically, it can impact the following activities that rely on measurement and analysis of traffic flows:

- o Network Operations and Research: Observable transport headers enable operators and the research community to measure and analyse protocol performance, network anomalies, and failure pathologies. This information can help inform capacity planning, and assist in determining the need for equipment and/or configuration changes by network operators. This data also can inform Internet engineering research, and help the develop of new protocols and procedures. Encryption of the entire transport protocol, including header information, will restrict the availability of data, and might lead to the development of alternative, and potentially more intrusive, methods to acquire the needed data. Encrypting the transport payload, but leaving some, or all, of the transport headers unencrypted but authenticated can provide the majority of the privacy and security benefits while allowing some measurement.
- o Network Troubleshooting and diagnostics: Encrypting transport header information eliminates the incentive for operators to troubleshoot what they cannot interpret. A flow experiencing packet loss looks like an unaffected flow when only observing network layer headers (if transport sequence numbers and flow identifiers are obscured). This limits understanding of the impact of packet loss on the flows that share a network segment. Encrypted traffic therefore implies "don't touch", and a likely trouble-shooting response will be "can't help, no trouble found". The additional mechanisms that will need to be introduced to help reconstruct transport-level metrics add complexity and operational costs [I-D.mmm-wg-effect-encrypt].
- o Network Traffic Analysis: The use of encryption can make it harder to determine which transport protocols and features are being used across a network segment. The trends in usage. This could impact the ability for an operator to anticipate the need for network upgrades and roll-out. It can also impact the on-going traffic engineering activities performed by operators. While the impact may, in many cases, be small there are scenarios where operators directly support particular services (e.g., in radio links, or to troubleshoot issues realting to Quality of Service, QoS). The more complex the underlying infrastructure the more important this impact.

- o Open and Verifiable Network Data: The use of transport header encryption reduces the range of actors that can capture useful measurement data. This is, of course, its goal. Doing so, however, limits the information sources available to the Internet community to understand the operation of transport protocols, so preventing access to the information necessary to inform design decisions and standards for new protocols and related operational practices. There are dangers in a model where only endpoints (i.e., at user devices and within service platforms) can observe performance, and this cannot be independently verified. To ensure the health of the standards and research communities, we need independently captured data to develop on the behaviour of the transports. Independently verifiable performance metrics might also be important in order to demonstrate regulatory compliance in some jurisdictions.

The last point leads us to consider the impact of encrypting all the transport headers the specification and development of protocols and standards. It has potential impact on:

- o Understanding Feature Interactions: An appropriate vantage point, coupled with timing information about traffic flows, provides a valuable tool for benchmarking equipment and/or configurations, and to understand complex feature interactions. Transport header encryption limits the ability to diagnose and explore interactions between features at different protocol layers, a side-effect of not allowing a choice of vantage point from which this information is observed.
- o Supporting Common Specifications: The Transmission Control Protocol (TCP) is the predominant transport protocol. Its many variants have broadly consistent approaches to avoiding congestion collapse, and to ensuring the stability of the network. Increased use of transport layer encryption can overcome ossification, allowing deployment of new transports with different types of congestion control. This flexibility can be beneficial, but it comes at the cost of fragmenting the ecosystem. There's little doubt that developers will try to produce high quality transports for their target uses, but it is not clear there are sufficient incentives to ensure good practice that benefits the wide diversity of requirements for the Internet community as a whole. Increased diversity, and the ability to innovate without public scrutiny, risks point solutions that optimise for specific needs, but accidentally disrupt operations of/in different parts of the network. The social compact that maintains the stability of the network relies on accepting common specifications, and on the ability to verify that others also conform.

- o Operational practice: Published transport specifications allow operators to check compliance. This can bring assurance to those operating networks, often avoiding the need to deploy complex techniques that routinely monitor and manage TCP/IP traffic flows (e.g. Avoiding the capital and operational costs of deploying flow rate-limiting and network circuit-breaker methods). This should continue when encrypted transport headers are used, but methods need to confirm that the traffic produced conforms to the expectations of the operator or developer.
- o Restricting research and development: The use of encryption may impede independent research into new mechanisms, measurement of behaviour, and development initiatives. Experience shows that transport protocols are complicated to design and complex to deploy, and that individual mechanisms need to be evaluated while considering other mechanism, across a broad range of network topologies and with attention to the impact on traffic sharing the capacity. Adopting pervasive encryption of transport information could eliminate the independent self-checks that have previously been in place from research and academic contributors (e.g., the role of the IRTF ICCRG, and research publications in reviewing new transport mechanisms and assessing the impact of their experimental deployment).

Pervasive use of transport header encryption can impact the ways that protocols are designed, standardised, deployed, and operated. The choice of whether future transport protocols encrypt their protocol headers therefore needs to be taken based not solely on security and privacy considerations, but also taking into account the impact on operations, standards, and research. A network that is secure but unusable due to persistent congestion collapse is not an improvement, and while that would be an extreme outcome proposals that impose high costs for very limited benefits need to be considered carefully, to ensure the benefits outweigh the costs.

1.1. Current uses of Transport Headers within the Network

The transport layer is the first end-to-end layer in the network stack. Despite headers having end-to-end meaning, some transport headers have come to be used in various ways within the Internet. In response to pervasive monitoring [RFC7624] revelations and the IETF consensus that "Pervasive Monitoring is an Attack" [RFC7258], efforts are underway to increase encryption of Internet traffic, which would prevent visibility of transport headers. This affects on how network protocols are designed and used [I-D.mm-wg-effect-encrypt]. To understand these implications, it is first necessary to understand how transport layer headers are currently observed and/or modified by middleboxes within the network.

Transport protocols can be designed to encrypt or authenticate transport header fields. Authentication methods at the transport layer can be used to detect any changes to an immutable header field that were made by a network device along a path. The intentional modification of transport headers by middleboxes (such as Network Address Translation with Protocol Translation, NAT-PT, or Firewalls) is not considered.

1.1.1.1. Observing Transport Information in the Network

In-network observation of transport protocol headers requires knowledge of the format of the transport header:

- o Flows need to be identified at the level required for monitoring;
- o The protocol and version of the header need to be observable. As protocols evolve over time and there may be a need to introduce new transport headers. This may require interpretation of protocol version information or connection setup information;
- o Location and syntax of any transport headers to be observed. IETF transport protocols specify this information.

The following subsections describe various ways that observable transport information may be utilised.

1.1.1.1.1. Flow Identification

Transport protocol header information can identify a flow and the connection state of the flow, together with the protocol options being used. In some usages, a low-numbered (well-known) port can identify a protocol (although port information alone is not sufficient to guarantee identification of a protocol). Transport protocols, such as TCP and Stream Control Transport Protocol (SCTP) specify a standard base header that includes sequence number information and other data, with the possibility to negotiate additional headers at connection setup, identified by an option number in the transport header. UDP-based protocols can use, but sometimes do not use, well-known ports. Some can instead be identified by signalling protocols or through the use of magic numbers placed in the first byte(s) of the datagram payload.

1.1.1.1.2. Metrics derived from Transport Layer Headers

Some actors have a need to characterise the performance of link/network segments. Passive monitoring uses observed traffic to make inferences from transport headers to derive these measurements. A variety of open source and commercial tools have been deployed that utilise this information. The following metrics can be derived from transport header information:

Traffic Rate and Volume: Header information may allow derivation of volume measures per-application, to characterise the traffic that uses a network segment or the pattern of network usage. This may be measured per endpoint or aggregate of endpoint (e.g., by an operator to assess subscriber usage). It can also be used to trigger measurement-based traffic shaping and to implement QoS support within the network and lower layers. Volume measures can be valuable for capacity planning (providing detail of trends rather than the volume per subscriber).

Loss Rate and Loss Pattern: Flow loss rate may be derived and is often used as a metric for performance assessment and to characterise transport behaviour. Understanding the root cause of loss can help an operator determine whether this requires corrective action.

There are various cause of loss, including: corruption on a link (e.g., interference on a radio link), buffer overflow (e.g., due to congestion), policing (traffic management), buffer management (e.g., Active Queue Management, AQM). Understanding flow loss rate requires either maintaining per flow packet counters or by observing sequence numbers in transport headers. Loss can be monitored at the interface level by devices in the network. It is often important to understand the conditions under which packet loss occurs. This usually requires relating loss to the traffic flowing on the network segment at the time of loss.

Observation of transport feedback information (observing loss reports, e.g., RTP Control Protocol (RTCP), TCP SACK) can increase understanding of the impact of loss and help identify cases where loss may have been wrongly identified, or the transport did not require the lost packet. It is sometimes more important to understand the pattern of loss, than the loss rate - since losses can often occur as bursts, rather than randomly-timed events.

Throughput and Goodput: The throughput observed by a flow can be determined even when a flow is encrypted, providing the individual flow can be identified. Goodput [RFC7928] is a measure of useful data exchanged (the ratio of useful/total volume of traffic sent by a flow), which requires ability to differentiate loss and retransmission of packets (e.g., by observing packet sequence numbers in the TCP or the Real Time Protocol, RTP, headers

[RFC3550])).

Latency: Latency is a key performance metric that impacts application response time and user-perceived response time. It often indirectly impacts throughput and flow completion time. Latency determines the reaction time of the transport protocol itself, impacting flow setup, congestion control, loss recovery, and other transport mechanisms. The observed latency can have many components [Latency]. Of these, unnecessary/unwanted queuing in network buffers has often been observed as a significant factor. Once the cause of unwanted latency has been identified, this can often be eliminated, and determining latency metrics is a key driver in the deployment of AQM [RFC7567], DiffServ [RFC2474], and Explicit Congestion Notification (ECN) [RFC3168] [RFC8087].

To measure latency across a part of the path, an observation point can measure the experienced round trip time (RTT) using packet sequence numbers, and acknowledgements, or by observing header timestamp information. Such information allows an observation point in the network to determine not only the path RTT, but also to measure the upstream and downstream contribution to the RTT. This may be used to locate a source of latency, e.g., by observing cases where the ratio of median to minimum RTT is large for a part of a path.

An example usage of this method could identify excessive buffers to help deploy or configure AQM [RFC7567] [RFC7928] to effectively eliminate unnecessary queuing in routers and other devices. AQM methods need to be deployed at the capacity bottleneck, but are often deployed in combination with other techniques, such as scheduling [RFC7567] [I-D.ietf-aqm-fq-codel] and although parameter-less methods are desired [RFC7567], current methods [I-D.ietf-aqm-fq-codel] [I-D.ietf-aqm-codel] [I-D.ietf-aqm-pie] often cannot scale across all possible deployment scenarios. The service offered by operators can therefore benefit from latency information to understand the impact of deployment and tune deployed services.

Jitter: Some network applications are sensitive to changes in packet timing. For such applications, it can be necessary to measure the jitter observed along a portion of the path. The requirements to measure jitter resemble those for the measurement of latency.

Flow Reordering: Significant flow reordering can impact time-critical applications and can be interpreted as loss by reliable transports. Many transport protocol techniques are impacted by reordering (e.g., triggering TCP retransmission, or re-buffering

of real-time applications). Packet reordering can occur for many reasons (from equipment design to misconfiguration of forwarding rules).

As in the drive to reduce network latency, there is a need for operational tools to detect mis-ordered packet flows and quantify the degree of reordering. Techniques for measuring reordering typically observe packet sequence numbers. Metrics have been defined that evaluate whether a network has maintained packet order on a packet-by-packet basis [RFC4737] and [RFC5236].

There has been initiatives in the IETF transport area to reduce the impact of reordering within a transport flow, possibly leading to reduced the requirements for ordering. These have promise to simplify network equipment design as well as the potential to improve robustness of the transport service. Measurements of reordering can help understand the level of reordering within deployed infrastructure, and inform decisions about how to progress such mechanisms.

Some protocols provide in-built monitoring and reporting functions. Transport fields in the RTP header [RFC3550] [RFC4585] can be observed to derive traffic volume measurements and provide information on the progress and quality of a session using RTP. Key performance indicators are retransmission rate, packet drop rate, sector utilization level, a measure of reordering, peak rate, the CE-marking rate, etc. Metadata is often important to understand the context under which the data was collected, including the time, observation point, and way in which metrics were accumulated. The RTCP protocol directly reports some of this information in a form that can be directly visible in the network. A user of summary measurement data needs to trust the source of this data and the method used to generate the summary information.

When encryption conceals information in packet headers, measurements need to rely on pattern inferences and other heuristics grows, and accuracy suffers [I-D.mmm-wg-effect-encrypt].

1.1.1.3. Metrics derived from Network Layer Headers

Some transport information is made visible in the network-layer protocol header. These header fields are not encrypted and can be used to make flow observations.

Use of IPv6 Network-Layer Flow Label: Endpoints are encouraged expose flow information in the IPv6 Flow Label field of the network-layer header (e.g. [RFC8085]). This can be used to inform network-layer queuing, forwarding (e.g., for equal cost multi-path (ECMP) routing, and Link Aggregation, LAG). This can provide useful information to assign packets to flows in the data collected by measurement campaigns. Although important to characterising a path, it does not directly provide any performance data.

Use Network-Layer Differentiated Services Code Point Point: Applications can expose their delivery expectations to the network by setting the Differentiated Services Code Point (DSCP) field of IPv4 and IPv6 packets. This can be used to inform network-layer queuing and forwarding, and can also provide information on the relative importance of packet information collected by measurement campaigns, but does not directly provide any performance data.

This field provides explicit information that can be used in place of inferring traffic requirements (e.g., by inferring QoS requirements from port information via a multi-field classifier). The DSCP value can therefore impact the quality of experience for a flow. Observations of service performance need to consider this field when a network path has support for differentiated service treatment.

Use of Explicit Congestion Marking: ECN[RFC3168] is an optional transport mechanism that uses a code point in the network-layer header. Use of ECN can offer gains in terms of increased throughput, reduced delay, and other benefits when used over a path that includes equipment that supports an AQM method that performs Congestion Experienced (CE) marking of IP packets [RFC8087].

ECN exposes the presence of congestion on a network path to the transport and network layer. The reception of CE-marked packets can therefore be used to monitor the presence and estimate the level of incipient congestion on the upstream portion of the path from the point of observation (Section 2.5 of [RFC8087]). Because ECN marks carried in the IP protocol header, it is much easier to measure ECN than metering packet loss. However, interpreting the marking behaviour (i.e., assessing congestion and diagnosing faults) requires context from the transport layer (path RTT, visibility of loss - that could be due to queue overflow, congestion response, etc) [RFC7567].

Some ECN-capable network devices can provide richer (more frequent and fine-grained) indication of their congestion state. Setting congestion marks proportional to the level of congestion (e.g., Data Center TCP, DCTP [I-D.ietf-tcpm-dctcp], and Low Latency Low Loss Scalable throughput, L4S, [I-D.ietf-tsvwg-l4s-arch]).

Use of ECN requires feedback a transport to feed back reception information on the path towards the data sender. Exposure of this Transport ECN feedback provides an additional powerful tool to understand ECN-enabled AQM-based networks [RFC8087].

AQM and ECN offer a range of algorithms and configuration options, it is therefore important for tools to be available to network operators and researchers to understand the implication of configuration choices and transport behaviour as use of ECN increases and new methods emerge [RFC7567] [RFC8087]. ECN-monitoring is expected to become important as AQM is deployed that supports ECN [RFC8087].

1.1.2. Transport Measurement

The common language between network operators and application/content providers/users is packet transfer performance at a layer that all can view and analyse. For most packets, this has been transport layer, until the emergence of QUIC, with the obvious exception of VPNs and IPsec. When encryption conceals more layers in a packet, people seeking understanding of the network operation need to rely more on pattern inferences and other heuristics. The accuracy of measurements therefore suffers, as does the ability to investigate and troubleshoot interactions between different anomalies. For example, the traffic patterns between a web server and a browser are dependent on browser supplier and version, even use of the application (e.g., web e-mail access). Even when measurement datasets are made available (e.g., from endpoints) additional metadata, such as the state of the network, is often required to interpret the data. Collecting and coordinating such metadata is more difficult when the observation point is at a different location to the bottleneck/device under evaluation.

Packet sampling techniques can be used to scale the processing involved in observing packets on high rate links. This exports only the packet header information of (randomly) selected packets. The utility of these measurements depends on the type of bearer and number of mechanisms used by network devices. Simple routers are relatively easy to manage, a device with more complexity demands understanding of the choice of many system parameters. This level of complexity exists when several network methods are combined.

This section discusses topics concerning observation of transport flows, with a focus on transport measurement.

1.1.2.1. Point of Measurement

Often measurements can only be understood in the context of the other flows that share a bottleneck. A simple example is monitoring of AQM. For example, FQ-CODEL [I-D.ietf-aqm-fq-codel], combines sub queues (statistically assigned per flow), management of the queue length (CODEL), flow-scheduling, and a starvation prevention mechanism. Usually such algorithms are designed to be self-tuning, but current methods typically employ heuristics that can result in more loss under certain path conditions (e.g., large RTT, effects of multiple bottlenecks [RFC7567]).

In-network measurements can distinguish between upstream and downstream metrics with respect to the measurement point. These are particularly useful for locating the source of problems or to assess the performance of a network segment or a particular device configuration.

By correlating observations at multiple points along the path (e.g., at the ingress and egress of a network segment), an observer can determine the contribution of a portion of the path to an observed metric (to locate a source of delay, jitter, loss, reordering, congestion marking, etc.).

1.1.2.2. Use by Operators to Plan and Provision Networks

Traffic measurements (e.g., traffic volume, loss, latency) is used by operators to help plan deployment of new equipment and configurations in their networks. Data is also important to equipment vendors who need to understand traffic trends and patterns of usage as inputs to decisions about planning products and provisioning for new deployments. This measurement information can also be correlated with billing information when this is also collected by an operator.

A network operator supporting traffic that uses transport header encryption may not have access to per-flow measurement data. Trends in aggregate traffic can be observed and can be related this to the endpoint addresses being used, but it may not be possible to correlate patterns in measurements with changes in transport protocols (e.g., the impact of changes in introducing a new transport protocol mechanism). This increases the dependency on other indirect sources of information to inform planning and provisioning.

1.1.2.3. Service Performance Measurement

Traffic measurements (e.g., traffic volume, loss, latency) can be used by various actors to help analyse the performance available to users of a network segment, and inform operational practice. While active measurements may be used in-network passive measurements can have advantages in terms of eliminating unproductive traffic, reducing the influence of test traffic on the overall traffic mix, and the ability to choose the point of measurement Section 1.1.2.1.

1.1.2.4. Measuring Transport to Support Network Operations

Information provided by tools observing transport headers can help determine whether mechanisms are needed in the network to prevent flows from acquiring excessive network capacity. Operators can implement operational practices to manage traffic flows (e.g., to prevent flows from acquiring excessive network capacity under severe congestion) by deploying rate-limiters, traffic shaping or network transport circuit breakers [RFC8084].

Congestion Control Compliance of Traffic: Congestion control is a key transport function. Many network operators implicitly accept that TCP traffic to comply with a behaviour that is acceptable for use in the shared Internet. TCP algorithms have been continuously improved over decades, and they have reached a level of efficiency and correctness that custom application-layer mechanisms will struggle to easily duplicate [RFC8085].

A standards-compliant TCP stack provides congestion control may therefore be judged safe for use across the Internet. Applications developed on top of well-designed transports can be expected to appropriately control their network usage, reacting when the network experiences congestion, by back-off and reduce the load placed on the network. This is the normal expected behaviour for TCP and SCTP.

However when anomalies are detected, tools can interpret the transport protocol header information to help understand the impact of specific transport protocols (or protocol mechanisms) on the other traffic that shares a network. An observation in the network can gain understanding of the dynamics of a flow and its congestion control behaviour. Analysing observed packet sequence numbers can be used to help build confidence that an application flow backs-off its share of the network load in the face of persistent congestion, and hence to understand whether the behaviour is appropriate for sharing limited network capacity. For example, it is common to visualise plots of TCP sequence numbers versus time for a flow to understand how a flow shares available capacity, deduce its dynamics in response to congestion, etc.

Congestion Control Compliance for UDP Traffic UDP provides a minimal message-passing transport that has no inherent congestion control mechanisms. Because congestion control is critical to the stable operation of the Internet, applications and other protocols that choose to use UDP as an Internet transport are required to employ mechanisms to prevent congestion collapse, avoid unacceptable contributions to jitter/latency, and to establish an acceptable share of capacity with concurrent traffic [RFC8085].

A network operator needs tools to understand if UDP flows comply with congestion control expectations and therefore whether there

is a need to deploy methods such as rate-limiters, transport circuit breakers or other methods to enforce acceptable usage for the offered service.

UDP flows that expose a well-known header by specifying the format of header fields can allow information to be observed to gain understanding of the dynamics of a flow and its congestion control behaviour. For example, tools exist to monitor various aspects of the RTP and RTCP header information of real-time flows (see Section 1.1.1.2).

1.1.1.3. Use for Network Diagnostics and Troubleshooting

Transport header information is useful for a variety of operational tasks [I-D.mm-wg-effect-encrypt]: to diagnose network problems, assess performance, capacity planning, management of denial of service threats, and responding to user performance questions. These tasks seldom involve the need to determine the contents of the transport payload, or other application details.

A network operator supporting traffic that uses transport header encryption can see only encrypted transport headers. This prevents deployment of performance measurement tools that rely on transport protocol information. Choosing to encrypt all information may be expected to reduce the ability for networks to "help" (e.g., in response to tracing issues, making appropriate Quality of Service, QoS, decisions). For some this will be blessing, for others it may be a curse. For example, operational performance data about encrypted flows needs to be determined by traffic pattern analysis, rather than relying on traditional tools. This can impact the ability of the operator to respond to faults, it could require reliance on endpoint diagnostic tools or user involvement in diagnosing and troubleshooting unusual use cases or non-trivial problems. A key need here is that tools need to provide useful information during network anomalies (e.g., significant reordering, high or intermittent loss). Although many network operators utilise transport information as a part of their operational practice, the network will not break because transport headers are encrypted.

1.1.1.4. Observing Headers to Implement Network Policy

Information from the transport protocol can be used by a multi-field classifier as a part of policy framework. Policies are commonly used for QoS management for resource-constrained networks and by firewalls that use the information to implement access rules. Traffic that cannot be classified, will typically receive a default treatment.

2. Encryption and Authentication of Transport Headers

End-to-end encryption can be applied at various protocol layers. It can be applied above the transport to encrypt the transport payload. Encryption methods can hide information from an eavesdropper in the network. Encryption can also help protect the privacy of a user, by hiding data relating to user/device identity or location. Neither an integrity check nor encryption methods prevent traffic analysis, and usage needs to reflect that profiling of users, identification of location and fingerprinting of behaviour can take place even on encrypted traffic flows.

One motive to use encryption is a response to perceptions that the network has become ossified by over-reliance on middleboxes that prevent new protocols and mechanisms from being deployed. This has lead to a common perception that there is too much "manipulation" of protocol headers within the network, and that designing to deploy in such networks is preventing transport evolution. In the light of this, a method that authenticates transport headers may help improve the pace of transport development, by eliminating the need to always consider deployed middleboxes [I-D.trammell-plus-abstract-mech], or potentially to only explicitly enable middlebox use for particular paths with particular middleboxes that are deliberately deployed to realise a useful function for the network and/or users[RFC3135].

Another motivation stems from increased concerns about privacy and surveillance. Some Internet users have valued the ability to protect identity, user location, and defend against traffic analysis, and have used methods such as IPsec ESP and Tor [Tor]. Revelations about the use of pervasive surveillance [RFC7624] have, to some extent, eroded trust in the service offered by network operators, and following the Snowden revelation in the USA in 2013 has led to an increased desire for people to employ encryption to avoid unwanted "eavesdropping" on their communications. Whatever the reasons, there are now activities in the IETF to design new protocols that may include some form of transport header encryption (e.g., QUIC [I-D.ietf-quic-transport]).

Authentication methods (that provide integrity checks of protocols fields) have also been specified at the network layer, and this also protects transport header fields. The network layer itself carries protocol header fields that are increasingly used to help forwarding decisions reflect the need of transport protocols, such the IPv6 Flow Label [RFC6437], the Differentiated Services Code Point (DSCP) [RFC2474] and Explicit Congestion Notification (ECN) [RFC3168].

The use of transport layer authentication and encryption exposes a tussle between middlebox vendors, operators, applications developers and users.

- o On the one hand, future Internet protocols that enable large-scale encryption assist in the restoration of the end-to-end nature of the Internet by returning complex processing to the endpoints, since middleboxes cannot modify what they cannot see.

- o On the other hand, encryption of transport layer header information has implications for people who are responsible for operating networks and researchers and analysts seeking to understand the dynamics of protocols and traffic patterns.

Whatever the motives, a decision to use pervasive of transport header encryption will have implications on the way in which design and evaluation is performed, and which can in turn impact the direction of evolution of the TCP/IP stack.

The next subsections briefly review some security design options for transport protocols.

2.1. Authenticating the Transport Protocol Header

Transport layer header information can be authenticated. An integrity check that protects the immutable transport header fields, but can still expose the transport protocol header information in the clear, allowing in-network devices to observe these fields. An integrity check can not prevent in-network modification, but can avoid a receiving accepting changes and avoid impact on the transport protocol operation.

An example transport authentication mechanism is TCP-Authentication (TCP-AO) [RFC5925]. This TCP option authenticates TCP segments, including the IP pseudo header, TCP header, and TCP data. TCP-AO protects the transport layer, preventing attacks from disabling the TCP connection itself. TCP-AO may interact with middleboxes, depending on their behaviour [RFC3234].

The IPsec Authentication Header (AH) [RFC4302] works at the network layer and authenticates the IP payload. This therefore also authenticates all transport headers, and verifies their integrity at the receiver, preventing in-network modification.

2.2. Encrypting the Transport Payload

The transport layer payload can be encrypted to protect the content of transport segments. This leaves transport protocol header information in the clear. The integrity of immutable transport

header fields could be protected by combining this with an integrity check (Section 2.1).

Examples of encrypting the payload include Transport Layer Security (TLS) over TCP [RFC5246] [RFC7525] or Datagram TLS (DTLS) over UDP [RFC6347] [RFC7525].

2.3. Encrypting the Transport Header

The network layer payload could be encrypted (including the entire transport header and payload). This method does not expose any transport information to devices in the network, which also prevents modification along the network path.

The IPsec Encapsulating Security Payload (ESP) [RFC4303] is an example of encryption at the network layer, it encrypts and authenticates all transport headers, preventing visibility of the headers by in-network devices. Some Virtual Private Network (VPN) methods also encrypt these headers.

2.4. Authenticating Transport Information and Selectively Encrypting the Transport Header

A transport protocol design can encrypt selected header fields, while also choosing to authenticate fields in the transport header. This allows specific transport header fields to be made observable by network devices. End-to end integrity checks can prevent an endpoint from undetected modification of the immutable transport headers.

The choice of which fields to expose and which to encrypt is a design choice for the transport protocol. Any selective encryption method requires trading two conflicting goals for a transport protocol designer to decide which header fields to encrypt. On the one hand, security work typically employs a design technique that seeks to expose only what is needed. On the other hand, there may be performance and operational benefits in exposing selected information to network tools.

Mutable fields in the transport header provide opportunities for middleboxes to modify the transport behaviour (e.g., the extended headers described in [I-D.trammell-plus-abstract-mech]). This considers only immutable fields in the transport headers, that is, fields that may be authenticated end-to-end across a path.

An example of a method that encrypts some, but not all, transport information is GRE-in-UDP [RFC8086] when used with GRE encryption.

2.5. Adding Transport Information to Network-Layer Protocol Headers

The transport information can be made visible in a network-layer header. This has the advantage that this information can then be observed by in-network devices. This has the advantage that a single header can support all transport protocols, but there may also be less desirable implications of separating the operation of the transport protocol from the measurement framework.

Some measurements may be made by adding additional protocol headers carrying operations, administration and management (OAM) information to packets at the ingress to a maintenance domain (e.g., an Ethernet protocol header with timestamps and sequence number information using a method such as 802.1lag) and removing the additional header at the egress of the maintenance domain. This approach enables some types of measurements, but does not cover the entire range of measurements described in this document.

Another example of a network-layer approach is the IPv6 Performance and Diagnostic Metrics (PDM) Destination Option [I-D.ietf-ippm-6man-pdm-option]. This allows a sender to optionally include a destination option that carries header fields that can be used to observe timestamps and packet sequence numbers. This information could be authenticated by receiving transport endpoints when the information is added at the sender and visible at the receiving endpoint, although methods to do this have not currently been proposed. This method needs to be explicitly enabled at the sender.

A drawback of using extension headers is that IPv4 network options are often not supported (or are carried on a slower processing path) and some IPv6 networks are also known to drop packets that set an IPv6 header extension. Another disadvantage is that protocols that separately expose header information do not necessarily have an advantage to expose the information that is utilised by the protocol itself, and could manipulate this header information to gain an advantage from the network.

3. Implications of Protecting the Transport Headers

This section explores key implications of working with encrypted transport protocols.

3.1. Independent Measurement

Independent observation by multiple actors is important for scientific analysis. Encrypting transport header encryption changes the ability for other actors to collect and independently analyse data. Internet transport protocols employ a set of mechanisms. Some

of these need to work in cooperation with the network layer - loss detection and recovery, congestion detection and congestion control, some of these need to work only end-to-end (e.g., parameter negotiation, flow-control).

When encryption conceals information in the transport header, it could be possible for an applications to provide summary data on performance and usage of the network. This data could be made available to other actors. However, this data needs to contain sufficient detail to understand (and possibly reconstruct the network traffic pattern for further testing) and to be correlated with the configuration of the network paths being measured. Sharing information between actors needs also to consider the privacy of the user and the incentives for providing accurate and detailed information. Protocols that expose the state information used by the transport protocol in their header information (e.g., timestamps used to calculate the RTT, packet numbers used to asses congestion and requests for retransmission) provide an incentive for the sending endpoint to provide correct information, increasing confidence that the observer understands the transport interaction with the network. This becomes important when considering changes to transport protocols, changes in network infrastructure, or the emergence of new traffic patterns.

3.2. Characterising "Unknown" Network Traffic

The patterns and types of traffic that share Internet capacity changes with time as networked applications, usage patterns and protocols continue to evolve.

If "unknown" or "uncharacterised" traffic patterns form a small part of the traffic aggregate passing through a network device or segment of the network the path, the dynamics of the uncharacterised traffic may not have a significant collateral impact on the performance of other traffic that shares this network segment. Once the proportion of this traffic increases, the need to monitor the traffic and determine if appropriate safety measures need to be put in place.

Tracking the impact of new mechanisms and protocols requires traffic volume to be measured and new transport behaviours to be identified. This is especially true of protocols operating over a UDP substrate. The level and style of encryption needs to be considered in determining how this activity is performed. On a shorter timescale, information may also need to be collected to manage denial of service attacks against the infrastructure.

3.3. Accountability and Internet Transport Protocols

Information provided by tools observing transport headers can help determine whether mechanisms are needed in the network to prevent flows from acquiring excessive network capacity, and where needed to deploy appropriate tools Section 1.1.2.4. Obfuscating or hiding this information using encryption is expected to lead operators and maintainers of middleboxes (firewalls, etc.) to seek other methods to classify and mechanisms to condition network traffic. A lack of data seems likely to reduce the level of precision with which these mechanisms are applied, and this needs to be considered when evaluating the impact of designs for transport encryption.

3.4. Impact on Research, Development and Deployment

Measurement data is increasingly being used to inform design decisions in networking research, during development of new mechanisms and protocols and in standardisation. Measurement has a critical role in the design of transport protocol mechanisms and their acceptance by the wider community (e.g., as a method to judge the safety for Internet deployment). Observation of pathologies are also important in understanding the interactions between cooperating protocols and network mechanism, the implications of sharing capacity with other traffic and the impact of different patterns of usage.

Attention needs to be paid to the expected scale of deployment of new protocols and protocol mechanisms. Whatever the mechanism, experience has shown that it is often difficult to correctly implement combination of mechanisms [RFC8085]. These mechanisms therefore typically evolve as a protocol matures, or in response to changes in network conditions, changes in network traffic or changes to application usage.

The growth and diversity of applications and protocols using the Internet continues to expand - and there has been recent interest in a wide range of new transport methods, e.g., Larger Initial Window, Proportional Rate Reduction (PRR), congestion control methods based on measuring bottleneck bandwidth and round-trip propagation time, the introduction of AQM techniques and new forms of ECN response (e.g., Data Centre TCP, DCTP [I-D.ietf-tcpm-dctcp], and methods proposed for Low Latency Low Loss Scalable throughput, L4S). For each new method it is desirable to build a body of data reflecting its behaviour under a wide range of deployment scenarios, traffic load, and interactions with other deployed/candidate methods.

Open standards motivate a desire for this evaluation to include independent observation and evaluation of performance data, which in turn suggests control over where and when measurement samples are collected. This requires consideration of the appropriate balance between encrypting all and no transport information.

4. Acknowledgements

The author would like to thank all who have talked to him face-to-face or via email. ...

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688421. The opinions expressed and arguments employed reflect only the authors' view. The European Commission is not responsible for any use that may be made of that information.

5. Security Considerations

This document is about design and deployment considerations for transport protocols. Authentication, confidentiality protection, and integrity protection are identified as Transport Features by RFC8095". As currently deployed in the Internet, these features are generally provided by a protocol or layer on top of the transport protocol; no current full-featured standards-track transport protocol provides these features on its own. Therefore, these features are not considered in this document, with the exception of native authentication capabilities of TCP and SCTP for which the security considerations in RFC4895.

Open data, and accessibility to tools that can help understand trends in application deployment, network traffic and usage patterns can all contribute to understanding security challenges. Standard protocols and understanding of the interactions between mechanisms and traffic patterns can also provide valuable insight into appropriate security design. Like congestion control mechanisms, security mechanisms are difficult to design and implement correctly. It is hence recommended that applications employ well-known standard security mechanisms such as DTLS, TLS or IPsec, rather than inventing their own.

6. IANA Considerations

XX RFC ED - PLEASE REMOVE THIS SECTION XXX

This memo includes no request to IANA.

7. References

7.1. Normative References

[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>>.

7.2. Informative References

[I-D.dolson-plus-middlebox-benefits]
Dolson, D., Snellman, J., Boucadair, M. and C. Jacquenet,
"Beneficial Functions of Middleboxes", Internet-Draft
draft-dolson-plus-middlebox-benefits-03, March 2017.

[I-D.ietf-aqm-codel]

Nichols, K., Jacobson, V., McGregor, A. and J. Jana,
"Controlled Delay Active Queue Management", Internet-Draft
draft-ietf-aqm-codel-00, October 2014.

[I-D.ietf-aqm-fq-codel]
Hoeiland-Joergensen, T., McKenney, P., Taht, D., Gettys,
J. and E. Dumazet, "FlowQueue-Codel", Internet-Draft
draft-ietf-aqm-fq-codel-00, January 2015.

[I-D.ietf-aqm-pie]
Pan, R., Natarajan, P., Baker, F. and G. White, "PIE: A
Lightweight Control Scheme To Address the Bufferbloat
Problem", Internet-Draft draft-ietf-aqm-pie-00, October
2014.

[I-D.ietf-ippm-6man-pdm-option]
Elkins, N., Hamilton, R. and m. mackermann@bcbsm.com,
"IPv6 Performance and Diagnostic Metrics (PDM) Destination
Option", Internet-Draft draft-ietf-ippm-6man-pdm-
option-10, May 2017.

[I-D.ietf-quic-transport]
Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed
and Secure Transport", Internet-Draft draft-ietf-quic-
transport-03, May 2017.

[I-D.ietf-tcpm-accurate-ecn]
Briscoe, B., Kuehlewind, M. and R. Scheffenegger, "More
Accurate ECN Feedback in TCP", Internet-Draft draft-ietf-
tcpm-accurate-ecn-00, December 2015.

[I-D.ietf-tcpm-dctcp]
Bensley, S., Thaler, D., Balasubramanian, P., Eggert, L.
and G. Judd, "Datacenter TCP (DCTCP): TCP Congestion
Control for Datacenters", Internet-Draft draft-ietf-tcpm-
dctcp-06, May 2017.

[I-D.ietf-tsvwg-l4s-arch]
Briscoe, B., Schepper, K. and M. Bagnulo, "Low Latency,
Low Loss, Scalable Throughput (L4S) Internet Service:
Architecture", Internet-Draft draft-ietf-tsvwg-l4s-
arch-00, May 2017.

[I-D.mm-wg-effect-encrypt]
Moriarty, K. and A. Morton, "Effect of Pervasive
Encryption on Operators", Internet-Draft draft-mm-wg-
effect-encrypt-11, April 2017.

[I-D.trammell-plus-abstract-mech]
Trammell, B., "Abstract Mechanisms for a Cooperative Path
Layer under Endpoint Control", Internet-Draft draft-
trammell-plus-abstract-mech-00, September 2016.

[I-D.trammell-plus-statefulness]

Kuehlewind, M., Trammell, B. and J. Hildebrand,
"Transport-Independent Path Layer State Management",
Internet-Draft draft-trammell-plus-statefulness-02,
December 2016.

- [Latency] Briscoe, B., "Reducing Internet Latency: A Survey of Techniques and Their Merits", November 2014.
- [Measure] Fairhurst, G., Kuehlewind, M. and D. Lopez, "Measurement-based Protocol Design", June 2017.
- [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>>.
- [RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G. and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, DOI 10.17487/RFC3135, June 2001, <<http://www.rfc-editor.org/info/rfc3135>>.
- [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, <<http://www.rfc-editor.org/info/rfc3168>>.
- [RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, DOI 10.17487/RFC3234, February 2002, <<http://www.rfc-editor.org/info/rfc3234>>.
- [RFC3449] Balakrishnan, H., Padmanabhan, V., Fairhurst, G. and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry", BCP 69, RFC 3449, DOI 10.17487/RFC3449, December 2002, <<http://www.rfc-editor.org/info/rfc3449>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<http://www.rfc-editor.org/info/rfc3550>>.
- [RFC3819] Karn, P., Ed., Bormann, C., Fairhurst, G., Grossman, D., Ludwig, R., Mahdavi, J., Montenegro, G., Touch, J. and L. Wood, "Advice for Internet Subnetwork Designers", BCP 89, RFC 3819, DOI 10.17487/RFC3819, July 2004, <<http://www.rfc-editor.org/info/rfc3819>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.

- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <<http://www.rfc-editor.org/info/rfc4302>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <<http://www.rfc-editor.org/info/rfc4303>>.
- [RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C. and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, DOI 10.17487/RFC4585, July 2006, <<http://www.rfc-editor.org/info/rfc4585>>.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S. and J. Perser, "Packet Reordering Metrics", RFC 4737, DOI 10.17487/RFC4737, November 2006, <<http://www.rfc-editor.org/info/rfc4737>>.
- [RFC5236] Jayasumana, A., Piratla, N., Banka, T., Bare, A. and R. Whitner, "Improved Packet Reordering Metrics", RFC 5236, DOI 10.17487/RFC5236, June 2008, <<http://www.rfc-editor.org/info/rfc5236>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [RFC5559] Eardley, P., Ed., "Pre-Congestion Notification (PCN) Architecture", RFC 5559, DOI 10.17487/RFC5559, June 2009, <<http://www.rfc-editor.org/info/rfc5559>>.
- [RFC5925] Touch, J., Mankin, A. and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <<http://www.rfc-editor.org/info/rfc5925>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<http://www.rfc-editor.org/info/rfc6347>>.
- [RFC6437] Amante, S., Carpenter, B., Jiang, S. and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<http://www.rfc-editor.org/info/rfc6437>>.
- [RFC6679] Westerlund, M., Johansson, I., Perkins, C., O'Hanlon, P. and K. Carlberg, "Explicit Congestion Notification (ECN) for RTP over UDP", RFC 6679, DOI 10.17487/RFC6679, August 2012, <<http://www.rfc-editor.org/info/rfc6679>>.

- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <<http://www.rfc-editor.org/info/rfc7258>>.
- [RFC7525] Sheffer, Y., Holz, R. and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.
- [RFC7567] Baker, F. Ed., and G. Fairhurst, Ed., "IETF Recommendations Regarding Active Queue Management", BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <<http://www.rfc-editor.org/info/rfc7567>>.
- [RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C. and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <<http://www.rfc-editor.org/info/rfc7624>>.
- [RFC7713] Mathis, M. and B. Briscoe, "Congestion Exposure (ConEx) Concepts, Abstract Mechanism, and Requirements", RFC 7713, DOI 10.17487/RFC7713, December 2015, <<http://www.rfc-editor.org/info/rfc7713>>.
- [RFC7928] Kuhn, N., Ed., Natarajan, P., Ed., Khademi, N. Ed., and D. Ros, "Characterization Guidelines for Active Queue Management (AQM)", RFC 7928, DOI 10.17487/RFC7928, July 2016, <<http://www.rfc-editor.org/info/rfc7928>>.
- [RFC8084] Fairhurst, G., "Network Transport Circuit Breakers", BCP 208, RFC 8084, DOI 10.17487/RFC8084, March 2017, <<http://www.rfc-editor.org/info/rfc8084>>.
- [RFC8085] Eggert, L., Fairhurst, G. and G. Shepherd, "UDP Usage Guidelines", BCP 145, RFC 8085, DOI 10.17487/RFC8085, March 2017, <<http://www.rfc-editor.org/info/rfc8085>>.
- [RFC8086] Yong, L., Ed., Crabbe, E., Xu, X. and T. Herbert, "GRE-in-UDP Encapsulation", RFC 8086, DOI 10.17487/RFC8086, March 2017, <<http://www.rfc-editor.org/info/rfc8086>>.
- [RFC8087] Fairhurst, G. and M. Welzl, "The Benefits of Using Explicit Congestion Notification (ECN)", RFC 8087, DOI 10.17487/RFC8087, March 2017, <<http://www.rfc-editor.org/info/rfc8087>>.
- [Tor] The Tor Project, ., "<<https://www.torproject.org>>", June 2017.

Appendix A. Revision information

- 00 This is an individual draft for the IETF community.
- 01 This draft was a result of walking away from the text for a few days and then reorganising the content.
- 02 This draft fixes textual errors.
- 03 This draft follows feedback from people reading this draft.
- 04 This adds an additional contributor and includes significant reworking to ready this for review by the wider IETF community Colin Perkins joined the author list.

Comments from the community are welcome on the text and recommendations.

Authors' Addresses

Godred Fairhurst
University of Aberdeen
Department of Engineering
Fraser Noble Building
Aberdeen, AB24 3UE
Scotland

Email: gorry@erg.abdn.ac.uk
URI: <http://www.erg.abdn.ac.uk/>

Colin Perkins
University of Glasgow
School of Computing Science
Glasgow, G12 8QQ
Scotland

Email: csp@csp Perkins.org
URI: <https://csp Perkins.org/>