Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 13, 2016

K. Smith
Vodafone Group

May 12, 2016

Network management of encrypted traffic
draft-smith-encrypted-traffic-management-05

Abstract

Encrypted Internet traffic may pose traffic management challenges to network operators. This document recommends approaches to help manage encrypted traffic, without breaching user privacy or security.

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 November 13, 2016.

Table of Contents

1. Introduction

Networks utilise various management techniques to ensure efficient throughput, congestion management, anti-SPAM and security measures. Historically these functions have utilised visibility of the Internet application layer.

This visibility is rapidly diminishing - encrypted Internet traffic is expected to continue its upward trend, driven by increased privacy awareness, uptake by popular services, and advocacy from the [IAB], [RFC7258] and W3C [TAG] .

[IAB], [RFC7258] and [mm-effect-encrypt] recognise that network management functions may be impacted by encryption, and that solutions to persist these management functions must not threaten user security or privacy. Such solutions can ensure the benefits of encryption do not degrade network efficiency.

This document lists such solutions, and points to evolving IETF work addressing the problem.

1.1. Document structure

This document refers to network management functions that may be hindered by traffic encryption, as described in [mm-effect-encrypt]

It then describes the technical details of existing options to fully or partially persist these functions under encryption. The guidance includes existing techniques as well as ongoing IETF work in this area.'Encryption' in this document typically refers to HTTP over TLS [RFC2818]; other forms of encryption are noted where applicable.

Finally, a summary is provided of ongoing IETF work which is investigating how network operators, origin servers and clients may co-operate in efficient traffic delivery without the need for pervasive network monitoring.

The legal, political and commercial aspects of network management are recgnised but not covered in this technical document.

2. Network management functions

Please refer to 'Network Service Provider Monitoring' in [mm-effect-encrypt]

3. Persisting traffic management without breaching encryption

This section involves utilisation of 'Application-based Flow Information Visible to a Network', [mm-effect-encrypt].

3.1. Providing hints to and from the network

The following protocols aim to support a secure and privacy-aware dialogue between client, server and the network elements. These hints can allow information item exchange between the endpoints and the network, to assist queuing mechanisms and traffic pacing that accounts for network congestion and variable connection strength. These relate to the cooperative path to endpoint signalling as discussed at the IAB SEMI [SEMI] and MaRNEW [MaRNEW] workshops, with the network following a more clearly-defined role in encrypted traffic delivery.

3.1.1. DiffServ Code Points (DSCP)

Data packets may be flagged with a traffic class (class of service). Network operators may honour a DiffServ classification [RFC2474] entering their network, or may choose to override it (since it is potentially open to abuse by a service provider that classifies all

its content as high priority). The purpose is to help manage traffic and congestion in the network.

This requires the content provider to flag data packets. This is extra work for the provider, and it has potential for abuse if a content provider simply flags all packets with high priorities. The network would need to know which flags to trust and which to override. The use of DiffServ within the operator network is beneficial where the operator determines the class of service itself; but where content is encrypted then heuristics would be needed to predict the traffic type entering the network. HTTP/2 allows several streams to be multiplexed over a single TCP connection. This means that if a provider decides to send Web pages, videos, chat etc. as individual streams over the same connection, then DiffServ would be useless as it would apply to the TCP/IP connection as a whole. However it may be more efficient for such Web providers to serve each content type from separate, dedicated servers - this will become clearer as HTTP/2 deployments are tuned for optimal delivery.

3.1.2. Explicit Congestion Notification (ECN)

Explicit Congestion Notification [RFC6138] routers can exchange congestion notification headers to ECN compliant endpoints. This is in preference to inferring congestion from dropped packets (e.g. in TCP). The purpose is to help manage traffic and congestion in the network.

This solution is required to be implemented at network and service provider. The service provider will utilise the ECN to reduce throughput until it is notified that congestion has eased.

As with DiffServ, operators may not trust an external entity to mark packets in a fair/consistent manner.

3.1.3. Multiprotocol Label Switching (MPLS)

On entering an MPLS-compliant network [RFC3031], IP packets are flagged with a 'Forward Equivalence Class' (FEC). This allows the network to make packet-forwarding decisions according to their latency requirements. MPLS routers within the network parse and act upon the FEC value. The FEC is set according to the source IP address and port. The purpose is to help managing traffic and congestion in the network. This requires deployment of an MPLS 'backbone' with label-aware switches/ routers.

An up-to-date correspondence table between Websites (or service sites in general) and server IP address must be created. Then, the category(s) of traffic have to be consistently mapped to a set of

MPLS labels ,which entails a significant effort to setup and maintain.

Note: MPLS can specify how OSI Layer 3 (IP layer) traffic can be routed over Layer 2 (Data Link); DiffServ only operates over Layer 3. DiffServ is potentially a less complex integration as it is applied at the network edge servers only.

3.1.4. Substrate Protocol for User Datagrams (SPUD)

SPUD [SPUD] is a prototype to research how network devices on the path between endpoints can share information to improve a flow. The network involvement is outside of the end-to-end context, to minimise any privacy or security breach. The initial prototype involves grouping UDP packets into an explicit 'tube', however support of additional transport layers (such as TCP) will also be investigated.

3.1.5. Mobile throughput Guidance

Mobile Throughput Guidance In-band Signalling [MTG] is a draft proposal to allows the network to inform the server endpoint as to what bandwidth the TCP connection can reasonably expect. This allows the server to adapt their throughput pacing based on dynamic network conditions, which can assist mechanisms such as Adaptive Bitrate Streaming and TCP congestion control.

3.1.6. Port Control Protocol Flowdata options

PCP Flowdata options [PCPFD] defines a mechanism for a host to signal flow characteristics to the network, and the network to signal its ability to accommodate that flow back to the host. This allows certain network flows to receive service that is differentiated from other network flows, and may be used to establish flow priority before connection establishment. PCP Flowdata operates at IPv4/IPv6 level.

3.1.7. IPv6 Flow label

IPv6 includes a flow label header field. [RFC6438] details how this may be used to identify flows for load balancing and multipath routing, which may be of particular use for application-layer encrypted traffic. The flow label field is part of the main header, which means it is not subject to the disadvantages of extension headers (namely their risk of being dropped by intermediary routers). The flow label may also be exposed as part of the outer IP packet in an IP tunnel, thus providing network flow information without compromising the payload.

3.1.8. DISCUSS

Differentiated prIorities and Status Code-points Using Stun Signalling [DISCUSS] describes a mechanism for information exchange between an application and the network, viable only for UDP. As such it can be considered in the same bracket as SPUD

3.1.9. Active Queue Management

The IETF Active Queue Management and Packet Scheduling WG [AQM] works on algorithms to manage network queues, with the aim of reducing packet delay and taming aggressive/misbehaving flows. This includes allowing flow sources to control their sending rates to avoid unnecessary losses (e.g. with [RFC6138]).

3.1.10. Congestion Exposure

The Congestion Exposure WG [CONEX] makes congestion markings (based on congestion experienced in the flow) available to the network via IP headers, in order to drive capacity efficiency. The WG made an IPv6 binding before the group concluded, however it is feasible for the congestion exposure markings to also be transported by another mechanism, such as SPUD.

3.2. Inferred flow information

3.2.1. Heuristics

Heuristics can be used to map given input data to particular conclusions via some heuristic reasoning. Examples of input data to this reasoning include IP destination address, TCP destination port, server name from SNI, and typical traffic patterns (e.g. occurrence of IP packets and TCP segments over time). The accuracy of heuristics depends on whether the observed traffic originates from a source delivering a single service, or a blend of services. In many scenarios, this makes it possible to directly classify the traffic related to a specific server/service even when the traffic is fully encrypted.

If the server/service is co-located on an infrastructure with other services that shares the same IP-address, the encrypted traffic cannot be directly classified. However, commercial traffic classifiers today typically apply heuristic methods, using traffic pattern matching algorithms to be able to identify the traffic. As an example, classifier products are able to identify popular VoIP services using heuristic methods although the traffic is encrypted and mostly peer-to-peer.

3.3. Co-operation on congestion control

One idea from the IAB 'Managing Radio Networks in an Encrypted World' workshop [MaRNEW] was that of better co-operation between 3GPP mobile networks and Internet services on congestion management. . 3GPP networks are concerned with ensuring that all devices attached to a particular cell receive a fair share of radio resources. This is critical, since these resources are constrained to various licenced spectrum bands, and volatile due to signal strength variation/cell handover/interference etc. The resource sharing process occurs independently to TCP congestion management performed between the client and server connected via the mobile network: the result is that TCP may wrongly infer congestion and react accordingly, or attempt to accelerate throughput without consideration of the available radio resources. Therefore the notion is to investigate co-operation between radio and TCP congestion controls to better manage connection throughput.

4. Acknowledgements

The editor would like to thank the GSMA Web Working Group for their contributions, in particular to the technical solutions and network management functions; the contributions via the SAAG mailing list (Panos Kampanakis, Brian Carpenter); and Kathleen Moriarty and Al Morton for their guidance in aligning this draft with [mm-effect-encrypt]

5. IANA Considerations

There are no IANA considerations.

6. Security Considerations

The intention of this document is to consider how to persist network management of encrypted traffic, without breaching user privacy or end-to-end security. In particular this document does not recommend any approach that intercepts or modifies client-server Transport Layer Security.

7. References

7.1. Normative References

[RFC2474] Nichols, K., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, Dec 1998.

[RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.

[RFC3031] Rosen, E., "Multiprotocol Label Switching Architecture", RFC 3031, Jan 2001.

[RFC6138] Ramakrishnan, K., "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 6138, Sep 2001.

[RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, 2011.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, May 2014.

7.2. Informative References

[AQM] IETF, "Active Queue Management and Packet Scheduling (IETF WG)", 2016, <https://tools.ietf.org/wg/aqm/charters>.

[CONEX] IETF, "Congestion Exposure (concluded IETF WG)", 2015, <https://datatracker.ietf.org/wg/conex/charter/>.

[DISCUSS] Cisco, "Differentiated prIorities and Status Code-points Using Stun Signalling", 2015, <https://tools.ietf.org/html/draft-martinsen-tram-discuss-02>.

[IAB] IAB, "IAB statement on Internet confidentiality", n.d., <https://www.iab.org/2014/11/14/iab-statement-on-internet-confidentiality/>.

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

[mm-effect-encrypt] IETF, "Effect of Ubiquitous Encryption", n.d., <https://datatracker.ietf.org/doc/draft-mm-wg-effect-encrypt/>.

[MTG] IETF, "Mobile Throughput Guidance Inband Signaling Protocol", n.d., <https://datatracker.ietf.org/doc/draft-flinck-mobile-throughput-guidance/>.

[PCPFD] Cisco, "PCP Flowdata option", 2013, <https://tools.ietf.org/html/draft-wing-pcp-flowdata-00>.

[SEMI] IAB, "IAB workshop, 'Stack Evolution in a Middlebox Internet'", n.d., <https://www.iab.org/activities/workshops/semi/>.

[SPUD] IETF, "Substrate Protocol for User Datagrams", n.d., <https://tools.ietf.org/html/draft-hildebrand-spud-prototype-03>.

[TAG] W3C, "Securing the Web", n.d., <https://w3ctag.github.io/web-https/>.

Author's Address

Kevin Smith
Vodafone Group

Email: kevin.smith@vodafone.com