Internet Engineering Task Force (IETF)
Request for Comments: 8404
Category: Informational
ISSN: 2070-1721

K. Moriarty, Ed.
Dell EMC
A. Morton, Ed.
AT&T Labs

July 2018

Effects of Pervasive Encryption on Operators

Abstract

Pervasive monitoring attacks on the privacy of Internet users are of serious concern to both user and operator communities. RFC 7258 discusses the critical need to protect users' privacy when developing IETF specifications and also recognizes that making networks unmanageable to mitigate pervasive monitoring is not an acceptable outcome: an appropriate balance is needed. This document discusses current security and network operations as well as management practices that may be impacted by the shift to increased use of encryption to help guide protocol development in support of manageable and secure networks.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc8404.

Table of Contents

1. Introduction

In response to pervasive monitoring revelations and the IETF consensus that pervasive monitoring is an attack [RFC7258], efforts are underway to increase encryption of Internet traffic. Pervasive monitoring is of serious concern to users, operators, and application providers. RFC 7258 discusses the critical need to protect users' privacy when developing IETF specifications and also recognizes that making networks unmanageable to mitigate pervasive monitoring is not an acceptable outcome; rather, an appropriate balance would emerge over time.

This document describes practices currently used by network operators to manage, operate, and secure their networks and how those practices may be impacted by a shift to increased use of encryption. It provides network operators' perspectives about the motivations and objectives of those practices as well as effects anticipated by operators as use of encryption increases. It is a summary of

concerns of the operational community as they transition to managing networks with less visibility. This document does not endorse the use of the practices described herein, nor does it aim to provide a comprehensive treatment of the effects of current practices, some of which have been considered controversial from a technical or business perspectives or contradictory to previous IETF statements (e.g., [RFC1958], [RFC1984], and [RFC2804]). The following RFCs consider the end-to-end (e2e) architectural principle to be a guiding principle for the development of Internet protocols [RFC2775] [RFC3724] [RFC7754].

This document aims to help IETF participants understand network operators' perspectives about the impact of pervasive encryption, both opportunistic and strong end-to-end encryption, on operational practices. The goal is to help inform future protocol development to ensure that operational impact is part of the conversation. Perhaps new methods could be developed to accomplish some of the goals of current practices despite changes in the extent to which cleartext will be available to network operators (including methods that rely on network endpoints where applicable). Discussion of current practices and the potential future changes is provided as a prerequisite to potential future cross-industry and cross-layer work to support the ongoing evolution towards a functional Internet with pervasive encryption.

Traditional network management, planning, security operations, and performance optimization have been developed on the Internet where a large majority of data traffic flows without encryption. While unencrypted traffic has made information that aids operations and troubleshooting at all layers accessible, it has also made pervasive monitoring by unseen parties possible. With broad support and increased awareness of the need to consider privacy in all aspects across the Internet, it is important to catalog existing management, operational, and security practices that have depended upon the availability of cleartext to function and to explore if critical operational practices can be met by less-invasive means.

This document refers to several different forms of Service Providers (SPs). For example, network service providers (or network operators) provide IP-packet transport primarily, though they may bundle other services with packet transport. Alternatively, application service providers primarily offer systems that participate as an endpoint in communications with the application user and hosting service providers lease computing, storage, and communications systems in data centers. In practice, many companies perform two or more service provider roles but may be historically associated with one.

This document includes a sampling of current practices and does not attempt to describe every nuance. Some sections cover technologies used over a broad spectrum of devices and use cases.

1.1. Additional Background on Encryption Changes

Pervasive encryption in this document refers to all types of session encryption including Transport Layer Security (TLS), IP Security (IPsec), TCPcrypt [TCPcrypt], QUIC [QUIC] (IETF's specification of Google's QUIC), and others that are increasingly deployed. It is well understood that session encryption helps to prevent both passive and active attacks on transport protocols; more on pervasive monitoring can be found in "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement" [RFC7624]. Active attacks have long been a motivation for increased encryption, and preventing pervasive monitoring became a focus just a few years ago. As such, the Internet Architecture Board (IAB) released a statement advocating for increased use of encryption in November 2014 (see <https://www.iab.org/2014/11/14/iab-statement-on-internet-confidentiality/>). Perspectives on encryption paradigms have shifted over time to make ease of deployment a high priority and to balance that against providing the maximum possible level of security, regardless of deployment considerations.

One such shift is documented in Opportunistic Security (OS) [RFC7435], which suggests that when use of authenticated encryption is not possible, cleartext sessions should be upgraded to unauthenticated session encryption, rather than no encryption. OS encourages upgrading from cleartext but cannot require or guarantee such upgrades. Once OS is used, it allows for an evolution to authenticated encryption. These efforts are necessary to improve an end user's expectation of privacy, making pervasive monitoring cost prohibitive. With OS in use, active attacks are still possible on unauthenticated sessions. OS has been implemented as NULL Authentication with IPsec [RFC7619], and there are a number of infrastructure use cases such as server-to-server encryption where this mode is deployed. While OS is helpful in reducing pervasive monitoring by increasing the cost to monitor, it is recognized that risk profiles for some applications require authenticated and secure session encryption as well prevention of active attacks. IPsec, and other session encryption protocols, with authentication has many useful applications, and usage has increased for infrastructure applications such as for virtual private networks between data centers. OS, as well as other protocol developments like the Automated Certificate Management Environment (ACME), have increased the usage of session encryption on the Internet.

Risk profiles vary and so do the types of session encryption deployed. To understand the scope of changes in visibility, a few examples are highlighted. Work continues to improve the implementation, development, and configuration of TLS and DTLS sessions to prevent active attacks used to monitor or intercept session data. The changes from TLS 1.2 to 1.3 enhance the security of TLS, while hiding more of the session negotiation and providing less visibility on the wire. The Using TLS in Applications (UTA) Working Group has been publishing documentation to improve the security of TLS and DTLS sessions. They have documented the known attack vectors in [RFC7457], have documented best practices for TLS and DTLS in [RFC7525], and have other documents in development. The recommendations from these documents were built upon for TLS 1.3 to provide a more inherently secure end-to-end protocol.

In addition to encrypted website access (HTTP over TLS), there are other well-deployed application-level transport encryption efforts such as MTA-to-MTA (mail transfer agent) session encryption transport for email (SMTP over TLS) and gateway-to-gateway for instant messaging (the Extensible Messaging and Presence Protocol (XMPP) over TLS). Although this does provide protection from transport-layer attacks, the servers could be a point of vulnerability if user-to-user encryption is not provided for these messaging protocols. User-to-user content encryption schemes, such as S/MIME and Pretty Good Privacy (PGP) for email and Off-the-Record (OTR) encryption for XMPP are used by those interested in protecting their data as it crosses intermediary servers, preventing transport-layer attacks by providing an end-to-end solution. User-to-user schemes are under review, and additional options will emerge to ease the configuration requirements, making this type of option more accessible to non-technical users interested in protecting their privacy.

Increased use of encryption, either opportunistic or authenticated, at the transport, network, or application layer, impacts how networks are operated, managed, and secured. In some cases, new methods to operate, manage, and secure networks will evolve in response. In other cases, currently available capabilities for monitoring or troubleshooting networks could become unavailable. This document lists a collection of functions currently employed by network operators that may be impacted by the shift to increased use of encryption. This document does not attempt to specify responses or solutions to these impacts; it documents the current state.

1.2. Examples of Attempts to Preserve Functions

Following the Snowden [Snowden] revelations, application service providers (Yahoo, Google, etc.) responded by encrypting traffic between their data centers (IPsec) to prevent passive monitoring from taking place unbeknownst to them. Infrastructure traffic carried over the public Internet has been encrypted for some time; this change for universal encryption was specific to their private backbones. Large mail service providers also began to encrypt session transport (TLS) to hosted mail services. This and other increases in the use of encryption had the immediate effect of providing confidentiality and integrity for protected data, but it created a problem for some network-management functions. Operators could no longer gain access to some session streams resulting in actions by several to regain their operational practices that previously depended on cleartext data sessions.

The Electronic Frontier Foundation (EFF) reported [EFF2014] several network service providers using a downgrade attack to prevent the use of SMTP over TLS by breaking STARTTLS (Section 3.2 of [RFC7525]), essentially preventing the negotiation process resulting in fallback to the use of cleartext. There have already been documented cases of service providers preventing STARTTLS to avoid session encryption negotiation on some sessions. Doing so allows them to inject a super cookie that enables advertisers to track users; these actions are also considered an attack. These serve as examples of undesirable behavior that could be prevented through upfront discussions in protocol work for operators and protocol designers to understand the implications of such actions. In other cases, some service providers and enterprises have relied on middleboxes having access to cleartext for load-balancing, monitoring for attack traffic, meeting regulatory requirements, or other purposes. The implications for enterprises that own the data on their networks or that have explicit agreements that permit the monitoring of user traffic are very different from those for service providers who may be accessing content in a way that violates privacy considerations. Additionally, service provider equipment is designed for accessing only the headers exposed for the data-link, network, and transport layers. Delving deeper into packets is possible, but there is typically a high degree of accuracy from the header information and packet sizes when limited to header information from these three layers. Service providers also have the option of adding routing overlay protocols to traffic. These middlebox implementations, performing functions either considered legitimate by the IETF or not, have been impacted by increases in encrypted traffic. Only methods keeping with the goal of balancing network management and pervasive monitoring mitigation as discussed in [RFC7258] should be considered in work toward a solution resulting from this document.

It is well known that national surveillance programs monitor traffic for criminal activities [JNSLP] [RFC2804] [RFC7258]. Governments vary on their balance between monitoring versus the protection of user privacy, data, and assets. Those that favor unencrypted access to data ignore the real need to protect users' identities, financial transactions, and intellectual property (which require security and encryption to prevent crime). A clear understanding of technology, encryption, and monitoring goals will aid in the development of solutions as work continues towards finding an appropriate balance that allows for management while protecting user privacy with strong encryption solutions.

2. Network Service Provider Monitoring Practices

Service providers, for this definition, include the backbone ISPs as well as those providing infrastructure at scale for core Internet use (hosted infrastructure and services such as email).

Network service providers use various techniques to operate, manage, and secure their networks. The following subsections detail the purpose of several techniques as well as which protocol fields are used to accomplish each task. In response to increased encryption of these fields, some network service providers may be tempted to undertake undesirable security practices in order to gain access to the fields in unencrypted data flows. To avoid this situation, new methods could be developed to accomplish the same goals without service providers having the ability to see session data.

2.1. Passive Monitoring

2.1.1. Traffic Surveys

Internet traffic surveys are useful in many pursuits, such as input for studies of the Center for Applied Internet Data Analysis (CAIDA) [CAIDA], network planning, and optimization. Tracking the trends in Internet traffic growth, from earlier peer-to-peer communication to the extensive adoption of unicast video streaming applications, has relied on a view of traffic composition with a particular level of assumed accuracy, based on access to cleartext by those conducting the surveys.

Passive monitoring makes inferences about observed traffic using the maximal information available and is subject to inaccuracies stemming from incomplete sampling (of packets in a stream) or loss due to monitoring-system overload. When encryption conceals more layers in each packet, reliance on pattern inferences and other heuristics grows and accuracy suffers. For example, the traffic patterns between server and browser are dependent on browser supplier and

version, even when the sessions use the same server application (e.g., web email access). It remains to be seen whether more complex inferences can be mastered to produce the same monitoring accuracy.

2.1.2. Troubleshooting

Network operators use protocol-dissecting analyzers when responding to customer problems, to identify the presence of attack traffic, and to identify root causes of the problem such as misconfiguration. In limited cases, packet captures may also be used when a customer approves of access to their packets or provides packet captures close to the endpoint. The protocol dissection is generally limited to supporting protocols (e.g., DNS and DHCP), network and transport (e.g., IP and TCP), and some higher-layer protocols (e.g., RTP and the RTP Control Protocol (RTCP)). Troubleshooting will move closer to the endpoint with increased encryption and adjustments in practices to effectively troubleshoot using a 5-tuple may require education. Packet-loss investigations, and those where access is limited to a 2-tuple (IPsec tunnel mode), rely on network and transport-layer headers taken at the endpoint. In this case, captures on intermediate nodes are not reliable as there are far too many cases of aggregate interfaces and alternate paths in service provider networks.

Network operators are often the first ones called upon to investigate application problems (e.g., "my HD video is choppy"), to first rule out network and network services as a cause for the underlying issue. When diagnosing a customer problem, the starting point may be a particular application that isn’t working. The ability to identify the problem application’s traffic is important, and packet capture provided from the customer close to the edge may be used for this purpose; IP address filtering is not useful for applications using Content Delivery Networks (CDNs) or cloud providers. After identifying the traffic, an operator may analyze the traffic characteristics and routing of the traffic. This diagnostic step is important to help determine the root cause before exploring if the issue is directly with the application.

For example, by investigating packet loss (from TCP sequence and acknowledgement numbers), Round-Trip Time (RTT) (from TCP timestamp options or application-layer transactions, e.g., DNS or HTTP response time), TCP receive-window size, packet corruption (from checksum verification), inefficient fragmentation, or application-layer problems, the operator can narrow the problem to a portion of the network, server overload, client or server misconfiguration, etc. Network operators may also be able to identify the presence of attack

traffic as not conforming to the application the user claims to be using. In many instances, the exposed packet header is sufficient for this type of troubleshooting.

One way of quickly excluding the network as the bottleneck during troubleshooting is to check whether the speed is limited by the endpoints. For example, the connection speed might instead be limited by suboptimal TCP options, the sender's congestion window, the sender temporarily running out of data to send, the sender waiting for the receiver to send another request, or the receiver closing the receive window. All this information can be derived from the cleartext TCP header.

Packet captures and protocol-dissecting analyzers have been important tools. Automated monitoring has also been used to proactively identify poor network conditions, leading to maintenance and network upgrades before user experience declines. For example, findings of loss and jitter in Voice over IP (VoIP) traffic can be a predictor of future customer dissatisfaction (supported by metadata from RTP/RTCP) [RFC3550], or increases in DNS response time can generally make interactive web browsing appear sluggish. But, to detect such problems, the application or service stream must first be distinguished from others.

When increased encryption is used, operators lose a source of data that may be used to debug user issues. For example, IPsec obscures TCP and RTP header information, while TLS and the Secure Real-time Transport Protocol (SRTP) do not. Because of this, application-server operators using increased encryption might be called upon more frequently to assist with debugging and troubleshooting; thus, they may want to consider what tools can be put in the hands of their clients or network operators.

Further, the performance of some services can be more efficiently managed and repaired when information on user transactions is available to the service provider. It may be possible to continue transaction-monitoring activities without cleartext access to the application layers of interest; however, inaccuracy will increase and efficiency of repair activities will decrease. For example, an application-protocol error or failure would be opaque to network troubleshooters when transport encryption is applied, making root cause location more difficult and, therefore, increasing the time to repair. Repair time directly reduces the availability of the service, and most network operators have made availability a key metric in their Service Level Agreements (SLAs) and/or subscription rebates. Also, there may be more cases of user-communication failures when the additional encryption processes are introduced (e.g., key management at large scale), leading to more customer

service contacts and (at the same time) less information available to network-operation repair teams.

In mobile networks, knowledge about TCP's stream transfer progress (by observing ACKs, retransmissions, packet drops, and the Sector Utilization Level, etc.) is further used to measure the performance of network segments (sector, eNodeB (eNB), etc.). This information is used as key performance indicators (KPIs) and for the estimation of user/service key quality indicators at network edges for circuit emulation (CEM) as well as input for mitigation methods. If the makeup of active services per user and per sector are not visible to a server that provides Internet Access Point Names (APNs), it cannot perform mitigation functions based on network segment view.

It is important to note that the push for encryption by application providers has been motivated by the application of the described techniques. Although network operators have noted performance improvements with network-based optimization or enhancement of user traffic (otherwise, deployment would not have occurred), application providers have likewise noted some degraded performance and/or user experience, and such cases may result in additional operator troubleshooting. Further, encrypted application streams might avoid outdated optimization or enhancement techniques, where they exist.

A gap exists for vendors where built-in diagnostics and serviceability are not adequate to provide detailed logging and debugging capabilities that, when possible, could be accessed with cleartext network parameters. In addition to traditional logging and debugging methods, packet tracing and inspection along the service path provides operators the visibility to continue to diagnose problems reported both internally and by their customers. Logging of service path upon exit for routing overlay protocols will assist with policy management and troubleshooting capabilities for traffic flows on encrypted networks. Protocol trace logging and protocol data unit (PDU) logging should also be considered to improve visibility to monitor and troubleshoot application-level traffic. Additional work on this gap would assist network operators to better troubleshoot and manage networks with increasing amounts of encrypted traffic.

2.1.3. Traffic-Analysis Fingerprinting

Fingerprinting is used in traffic analysis and monitoring to identify traffic streams that match certain patterns. This technique can be used with both cleartext and encrypted sessions. Some Distributed Denial-of-Service (DDoS) prevention techniques at the network-provider level rely on the ability to fingerprint traffic in order to mitigate the effect of this type of attack. Thus, fingerprinting may be an aspect of an attack or part of attack countermeasures.

A common, early trigger for DDoS mitigation includes observing uncharacteristic traffic volumes or sources, congestion, or degradation of a given network or service. One approach to mitigate such an attack involves distinguishing attacker traffic from legitimate user traffic. The ability to examine layers and payloads above transport provides an increased range of filtering opportunities at each layer in the clear. If fewer layers are in the clear, this means that there are reduced filtering opportunities available to mitigate attacks. However, fingerprinting is still possible.

Passive monitoring of network traffic can lead to invasion of privacy by external actors at the endpoints of the monitored traffic. Encryption of traffic end to end is one method to obfuscate some of the potentially identifying information. For example, browser fingerprints are comprised of many characteristics, including User Agents, HTTP Accept headers, browser plug-in details, screen size and color details, system fonts, and time zones. A monitoring system could easily identify a specific browser, and by correlating other information, identify a specific user.

2.2. Traffic Optimization and Management

2.2.1. Load Balancers

A standalone load balancer is a function one can take off the shelf, place in front of a pool of servers, and configure appropriately, and it will balance the traffic load among servers in the pool. This is a typical setup for load balancers. Standalone load balancers rely on the plainly observable information in the packets they are forwarding and industry-accepted standards in interpreting the plainly observable information. Typically, this is a 5-tuple of the connection. This type of configuration terminates TLS sessions at the load balancer, making it the endpoint instead of the server. Standalone load balancers are considered middleboxes, but they are an integral part of server infrastructure that scales.

In contrast, an integrated load balancer is developed to be an integral part of the service provided by the server pool behind that load balancer. These load balancers can communicate state with their pool of servers to better route flows to the appropriate servers. They rely on non-standard, system-specific information and operational knowledge shared between the load balancer and its servers.

Both standalone and integrated load balancers can be deployed in pools for redundancy and load sharing. For high availability, it is important that when packets belonging to a flow start to arrive at a

different load balancer in the load-balancer pool, the packets continue to be forwarded to the original server in the server pool. The importance of this requirement increases as the chance of such a load balancer change event increases.

Mobile operators deploy integrated load balancers to assist with maintaining connection state as devices migrate. With the proliferation of mobile connected devices, there is an acute need for connection-oriented protocols that maintain connections after a network migration by an endpoint. This connection persistence provides an additional challenge for multihomed anycast-based services typically employed by large content owners and CDNs. The challenge is that a migration to a different network in the middle of the connection greatly increases the chances of the packets routed to a different anycast point of presence (POP) due to the new network's different connectivity and Internet peering arrangements. The load balancer in the new POP, potentially thousands of miles away, will not have information about the new flow and would not be able to route it back to the original POP.

To help with the endpoint network migration challenges, anycast service operations are likely to employ integrated load balancers that, in cooperation with their pool servers, are able to ensure that client-to-server packets contain some additional identification in plainly observable parts of the packets (in addition to the 5-tuple). As noted in Section 2 of [RFC7258], careful consideration in protocol design to mitigate pervasive monitoring is important, while ensuring manageability of the network.

An area for further research includes end-to-end solutions that would provide a simpler architecture and that may solve the issue with CDN anycast. In this case, connections would be migrated to a CDN unicast address.

Current protocols, such as TCP, allow the development of stateless integrated load balancers by availing such load balancers of additional plaintext information in client-to-server packets. In case of TCP, such information can be encoded by having server-generated sequence numbers (that are ACKed by the client), segment values, lengths of the packet sent, etc. The use of some of these mechanisms for load balancing negates some of the security assumptions associated with those primitives (e.g., that an off-path attacker guessing valid sequence numbers for a flow is hard). Another possibility is a dedicated mechanism for storing load-balancer state, such as QUIC's proposed connection ID to provide visibility to the load balancer. An identifier could be used for tracking purposes, but this may provide an option that is an improvement from bolting it on to an unrelated transport signal.

This method allows for tight control by one of the endpoints and can be rotated to avoid roving client linkability: in other words, being a specific, separate signal, it can be governed in a way that is finely targeted at that specific use case.

Some integrated load balancers have the ability to use additional plainly observable information even for today's protocols that are not network-migration tolerant. This additional information allows for improved availability and scalability of the load-balancing operation. For example, BGP reconvergence can cause a flow to switch anycast POPs, even without a network change by any endpoint. Additionally, a system that is able to encode the identity of the pool server in plaintext information available in each incoming packet is able to provide stateless load balancing. This ability confers great reliability and scalability advantages, even if the flow remains in a single POP, because the load-balancing system is not required to keep state of each flow. Even more importantly, there's no requirement to continuously synchronize such state among the pool of load balancers. An integrated load balancer repurposing limited existing bits in transport-flow state must maintain and synchronize per-flow state occasionally: using the sequence number as a cookie only works for so long given that there aren't that many bits available to divide across a pool of machines.

Mobile operators apply 3GPP Self-Organizing Networks (SONs) for intelligent workflows such as content-aware Mobility Load Balancing (MLB). Where network load balancers have been configured to route according to application-layer semantics, an encrypted payload is effectively invisible. This has resulted in practices of intercepting TLS in front of load balancers to regain that visibility, but at a cost to security and privacy.

In future Network Function Virtualization (NFV) architectures, load-balancing functions are likely to be more prevalent (deployed at locations throughout operators' networks). NFV environments will require some type of identifier (IPv6 flow identifiers, the proposed QUIC connection ID, etc.) for managing traffic using encrypted tunnels. The shift to increased encryption will have an impact on visibility of flow information and will require adjustments to perform similar load-balancing functions within an NFV.

2.2.2. Differential Treatment Based on Deep Packet Inspection (DPI)

Data transfer capacity resources in cellular radio networks tend to be more constrained than in fixed networks. This is a result of variance in radio signal strength as a user moves around a cell, the rapid ingress and egress of connections as users hand off between adjacent cells, and temporary congestion at a cell. Mobile networks

alleviate this by queuing traffic according to its required bandwidth and acceptable latency: for example, a user is unlikely to notice a 20 ms delay when receiving a simple web page or email, or an instant message response, but will very likely notice a rebuffering pause in a video playback or a VoIP call de-jitter buffer. Ideally, the scheduler manages the queue so that each user has an acceptable experience as conditions vary, but inferences of the traffic type have been used to make bearer assignments and set scheduler priority.

Deep Packet Inspection (DPI) allows identification of applications based on payload signatures, in contrast to trusting well-known port numbers. Application- and transport-layer encryption make the traffic type estimation more complex and less accurate; therefore, it may not be effectual to use this information as input for queue management. With the use of WebSockets [RFC6455], for example, many forms of communications (from isochronous/real-time to bulk/elastic file transfer) will take place over HTTP port 80 or port 443, so only the messages and higher-layer data will make application differentiation possible. If the monitoring system sees only "HTTP port 443", it cannot distinguish application streams that would benefit from priority queuing from others that would not.

Mobile networks especially rely on content-/application-based prioritization of Over-the-Top (OTT) services -- each application type or service has different delay/loss/throughput expectations, and each type of stream will be unknown to an edge device if encrypted. This impedes dynamic QoS adaptation. An alternate way to achieve encrypted application separation is possible when the User Equipment (UE) requests a dedicated bearer for the specific application stream (known by the UE), using a mechanism such as the one described in Section 6.5 of 3GPP TS 24.301 [TS3GPP]. The UE’s request includes the Quality Class Indicator (QCI) appropriate for each application, based on their different delay/loss/throughput expectations. However, UE requests for dedicated bearers and QCI may not be supported at the subscriber’s service level, or in all mobile networks.

These effects and potential alternative solutions have been discussed at the accord BoF [ACCORD] at IETF 95.

This section does not consider traffic discrimination by service providers related to Net Neutrality, where traffic may be favored according to the service provider’s preference as opposed to the user’s preference. These use cases are considered out of scope for this document as controversial practices.

2.2.3. Network-Congestion Management

For 3GPP User Plane Congestion Management (UPCON) [UPCON], the ability to understand content and manage networks during periods of congestion is the focus. Mitigating techniques such as deferred download, off-peak acceleration, and outbound roamers are a few examples of the areas explored in the associated 3GPP documents. The documents describe the issues, describe the data utilized in managing congestion, and make policy recommendations.

2.2.4. Performance-Enhancing Proxies

Performance-enhancing TCP proxies may perform local retransmission at the network edge; this also applies to mobile networks. In TCP, duplicated ACKs are detected and potentially concealed when the proxy retransmits a segment that was lost on the mobile link without involvement of the far end (see Section 2.1.1 of [RFC3135] and Section 3.5 of [MIDDLEBOXES]).

Operators report that this optimization at network edges improves real-time transmission over long-delay Internet paths or networks with large capacity variation (such as mobile/cellular networks). However, such optimizations can also cause problems with performance, for example, if the characteristics of some packet streams begin to vary significantly from those considered in the proxy design.

In general, some operators have stated that performance-enhancing proxies have a lower RTT to the client; therefore, they determine the responsiveness of flow control. A lower RTT makes the flow-control loop more responsive to changes in the mobile-network conditions and enables faster adaptation in a delay- and capacity-varying network due to user mobility.

Further, some use service-provider-operated proxies to reduce the control delay between the sender and a receiver on a mobile network where resources are limited. The RTT determines how quickly a user’s attempt to cancel a video is recognized and, therefore, how quickly the traffic is stopped, thus keeping unwanted video packets from entering the radio-scheduler queue. If impacted by encryption, performance-enhancing proxies could make use of routing overlay protocols to accomplish the same task, but this results in additional overhead.

An application-type-aware network edge (middlebox) can further control pacing, limit simultaneous HD videos, or prioritize active videos against new videos, etc. Services at this more granular level are limited with the use of encryption.

Performance-enhancing proxies are primarily used on long-delay links (satellite) with access to the TCP header to provide an early ACK and make the long-delay link of the path seem shorter. With some specific forms of flow control, TCP can be more efficient than alternatives such as proxies. The editors cannot cite research on this point specific to the performance-enhancing proxies described, but they agree this area could be explored to determine if flow-control modifications could preserve the end-to-end performance on long-delay path sessions where the TCP header is exposed.

2.2.5. Caching and Content Replication near the Network Edge

The features and efficiency of some Internet services can be augmented through analysis of user flows and the applications they provide. For example, network caching of popular content at a location close to the requesting user can improve delivery efficiency (both in terms of lower request response times and reduced use of links on the international Internet when content is remotely located), and service providers through an authorized agreement acting on their behalf use DPI in combination with content-distribution networks to determine if they can intervene effectively. Encryption of packet contents at a given protocol layer usually makes DPI processing of that layer and higher layers impossible. That being said, it should be noted that some content providers prevent caching to control content delivery through the use of encrypted end-to-end sessions. CDNs vary in their deployment options of end-to-end encryption. The business risk of losing control of content is a motivation outside of privacy and pervasive monitoring that is driving end-to-end encryption for these content providers.

It should be noted that caching was first supported in [RFC1945] and continued in the recent update of "Hypertext Transfer Protocol (HTTP/1.1): Caching" [RFC7234]. Some operators also operate transparent caches that neither the user nor the origin opt-in. The use of these caches is controversial within the IETF and is generally precluded by the use of HTTPS.

Content replication in caches (for example, live video and content protected by Digital Rights Management (DRM)) is used to most efficiently utilize the available limited bandwidth and thereby maximize the user’s Quality of Experience (QoE). Especially in mobile networks, duplicating every stream through the transit network increases backhaul cost for live TV. 3GPP Enhanced Multimedia Broadcast/Multicast Services (eMBMS) utilize trusted edge proxies to facilitate delivering the same stream to different users, using either unicast or multicast depending on channel conditions to the user. There are ongoing efforts to support multicast inside carrier networks while preserving end-to-end security: Automatic Multicast

Tunneling (AMT), for instance, allows CDNs to deliver a single (potentially encrypted) copy of a live stream to a carrier network over the public Internet and for the carrier to then distribute that live stream as efficiently as possible within its own network using multicast.

Alternate approaches are in the early phase of being explored to allow caching of encrypted content. These solutions require cooperation from content owners and fall outside the scope of what is covered in this document. Content delegation allows for replication with possible benefits, but any form of delegation has the potential to affect the expectation of client-server confidentiality.

2.2.6. Content Compression

In addition to caching, various applications exist to provide data compression in order to conserve the life of the user's mobile data plan or make delivery over the mobile link more efficient. The compression proxy access can be built into a specific user-level application, such as a browser, or it can be available to all applications using a system-level application. The primary method is for the mobile application to connect to a centralized server as a transparent proxy (user does not opt-in), with the data channel between the client application and the server using compression to minimize bandwidth utilization. The effectiveness of such systems depends on the server having access to unencrypted data flows.

Aggregated data stream content compression that spans objects and data sources that can be treated as part of a unified compression scheme (e.g., through the use of a shared segment store) is often effective at providing data offload when there is a network element close to the receiver that has access to see all the content.

2.2.7. Service Function Chaining

Service Function Chaining (SFC) is defined in RFC 7665 [RFC7665] and RFC 8300 [RFC8300]. As discussed in RFC 7498 [RFC7498], common SFC deployments may use classifiers to direct traffic into VLANs instead of using a Network Service Header (NSH), as defined in RFC 8300 [RFC8300]. As described in RFC 7665 [RFC7665], the ordered steering of traffic to support specific optimizations depends upon the ability of a classifier to determine the microflows. RFC 2474 [RFC2474] defines the following:

> Microflow: a single instance of an application-to-application flow of packets which is identified by source address, destination address, protocol id, and source port, destination port (where applicable).

SFC currently depends upon a classifier to at least identify the microflow. As the classifier's visibility is reduced from a 5-tuple to a 2-tuple, or if information above the transport layer becomes inaccessible, then the SFC classifier is not able to perform its job, and the service functions of the path may be adversely affected.

There are also mechanisms provided to protect security and privacy. In the SFC case, the layer below a network service header can be protected with session encryption. A goal is protecting end-user data, while retaining the intended functions of RFC 7665 [RFC7665] at the same time.

2.3. Content Filtering, Network Access, and Accounting

Mobile networks and many ISPs operate under the regulations of their licensing government authority. These regulations include Lawful Intercept, adherence to Codes of Practice on content filtering, and application of court order filters. Such regulations assume network access to provide content filtering and accounting, as discussed below. As previously stated, the intent of this document is to document existing practices; the development of IETF protocols follows the guiding principles of [RFC1984] and [RFC2804] and explicitly does not support tools and methods that could be used for wiretapping and censorship.

2.3.1. Content Filtering

There are numerous reasons why service providers might block content: to comply with requests from law enforcement or regulatory authorities, to effectuate parental controls, to enforce content-based billing, or for other reasons, possibly considered inappropriate by some. See RFC 7754 [RFC7754] for a survey of Internet filtering techniques and motivations and the IAB consensus on those mechanisms. This section is intended to document a selection of current content-blocking practices by operators and the effects of encryption on those practices. Content blocking may also happen at endpoints or at the edge of enterprise networks, but those scenarios are not addressed in this section.

In a mobile network, content filtering usually occurs in the core network. With other networks, content filtering could occur in the core network or at the edge. A proxy is installed that analyzes the transport metadata of the content users are viewing and filters content based on either a blacklist of sites or the user's predefined profile (e.g., for age-sensitive content). Although filtering can be done by many methods, one commonly used method involves a trigger based on the proxy identifying a DNS lookup of a host name in a URL that appears on a blacklist being used by the operator. The

subsequent requests to that domain will be rerouted to a proxy that checks whether the full URL matches a blocked URL on the list, and it will return a 404 if a match is found. All other requests should complete. This technique does not work in situations where DNS traffic is encrypted (e.g., by employing [RFC7858]). This method is also used by other types of network providers enabling traffic inspection, but not modification.

Content filtering via a proxy can also utilize an intercepting certificate where the client's session is terminated at the proxy enabling for cleartext inspection of the traffic. A new session is created from the intercepting device to the client's destination; this is an opt-in strategy for the client, where the endpoint is configured to trust the intercepting certificate. Changes to TLS 1.3 do not impact this more invasive method of interception, which has the potential to expose every HTTPS session to an active man in the middle (MITM).

Another form of content filtering is called parental control, where some users are deliberately denied access to age-sensitive content as a feature to the service subscriber. Some sites involve a mixture of universal and age-sensitive content and filtering software. In these cases, more-granular (application-layer) metadata may be used to analyze and block traffic. Methods that accessed cleartext application-layer metadata no longer work when sessions are encrypted. This type of granular filtering could occur at the endpoint or as a proxy service. However, the lack of ability to efficiently manage endpoints as a service reduces network service providers' ability to offer parental control.

2.3.2. Network Access and Data Usage

Approved access to a network is a prerequisite to requests for Internet traffic.

However, there are cases (beyond parental control) when a network service provider currently redirects customer requests for content (affecting content accessibility):

1. The network service provider is performing the accounting and billing for the content provider, and the customer has not (yet) purchased the requested content.

2. Further content may not be allowed as the customer has reached their usage limit and needs to purchase additional data service, which is the usual billing approach in mobile networks.

Currently, some network service providers redirect the customer using HTTP redirect to a captive portal page that explains to those customers the reason for the blockage and the steps to proceed. [RFC6108] describes one viable web notification system. When the HTTP headers and content are encrypted, this appropriately prevents mobile carriers from intercepting the traffic and performing an HTTP redirect. As a result, some mobile carriers block customer's encrypted requests, which impacts customer experience because the blocking reason must be conveyed by some other means. The customer may need to call customer care to find out the reason and/or resolve the issue, possibly extending the time needed to restore their network access. While there are well-deployed alternate SMS-based solutions that do not involve out-of-specification protocol interception, this is still an unsolved problem for non-SMS users.

Further, when the requested service is about to consume the remainder of the user's plan limits, the transmission could be terminated and advance notifications may be sent to the user by their service provider to warn the user ahead of the exhausted plan. If web content is encrypted, the network provider cannot know the data transfer size at request time. Lacking this visibility of the application type and content size, the network would continue the transmission and stop the transfer when the limit was reached. A partial transfer may not be usable by the client wasting both network and user resources, possibly leading to customer complaints. The content provider does not know a user's service plans or current usage and cannot warn the user of plan exhaustion.

In addition, some mobile network operators sell tariffs that allow free-data access to certain sites, known as 'zero rating'. A session to visit such a site incurs no additional cost or data usage to the user. For some implementations, zero rating is impacted if encryption hides the details of the content domain from the network.

2.3.3. Application Layer Gateways (ALGs)

Application Layer Gateways (ALGs) assist applications to set connectivity across Network Address Translators (NATs), firewalls, and/or load balancers for specific applications running across mobile networks. Section 2.9 of [RFC2663] describes the role of ALGs and their interaction with NAT and/or application payloads. ALGs are deployed with an aim to improve connectivity. However, it is an IETF best common practice recommendation that ALGs for UDP-based protocols be turned off [RFC4787].

One example of an ALG in current use is aimed at video applications that use the Real-Time Streaming Protocol (RTSP) [RFC7826] primary stream as a means to identify related RTP/RTCP [RFC3550] flows at setup. The ALG in this case relies on the 5-tuple flow information derived from RTSP to provision NAT or other middleboxes and provide connectivity. Implementations vary, and two examples follow:

1. Parse the content of the RTSP stream and identify the 5-tuple of the supporting streams as they are being negotiated.

2. Intercept and modify the 5-tuple information of the supporting media streams as they are being negotiated on the RTSP stream, which is more intrusive to the media streams.

When RTSP-stream content is encrypted, the 5-tuple information within the payload is not visible to these ALG implementations; therefore, they cannot provision their associated middleboxes with that information.

The deployment of IPv6 may well reduce the need for NAT and the corresponding requirement for ALGs.

2.3.4. HTTP Header Insertion

Some mobile carriers use HTTP header insertion (see Section 3.2.1 of [RFC7230]) to provide information about their customers to third parties or to their own internal systems [Enrich]. Third parties use the inserted information for analytics, customization, advertising, cross-site tracking of users, customer billing, or selectively allowing or blocking content. HTTP header insertion is also used to pass information internally between a mobile service provider's sub-systems, thus keeping the internal systems loosely coupled. When HTTP connections are encrypted to protect user privacy, mobile network service providers cannot insert headers to accomplish the, sometimes considered controversial, functions above.

Guidance from the Internet Architecture Board has been provided in "Design Considerations for Metadata Insertion" [RFC8165]. The guidance asserts that designs that share metadata only by explicit actions at the host are preferable to designs in which middleboxes insert metadata. Alternate notification methods that follow this and other guidance would be helpful to mobile carriers.

3. Encryption in Hosting and Application SP Environments

Hosted environments have had varied requirements in the past for encryption, with many businesses choosing to use these services primarily for data and applications that are not business or privacy sensitive. A shift prior to the revelations on surveillance/passive monitoring began where businesses were asking for hosted environments to provide higher levels of security so that additional applications and service could be hosted externally. Businesses understanding the threats of monitoring in hosted environments increased that pressure to provide more secure access and session encryption to protect the management of hosted environments as well as the data and applications.

3.1. Management-Access Security

Hosted environments may have multiple levels of management access, where some may be strictly for the Hosting service provider (infrastructure that may be shared among customers), and some may be accessed by a specific customer for application management. In some cases, there are multiple levels of hosting service providers, further complicating the security of management infrastructure and the associated requirements.

Hosting service provider management access is typically segregated from other traffic with a control channel and may or may not be encrypted depending upon the isolation characteristics of the management session. Customer access may be through a dedicated connection, but discussion for that connection method is out of scope for this document.

In overlay networks (e.g., Virtual eXtensible Local Area Network (VXLAN), Geneve, etc.) that are used to provide hosted services, management access for a customer to support application management may depend upon the security mechanisms available as part of that overlay network. While overlay-network data encapsulations may be used to indicate the desired isolation, this is not sufficient to prevent deliberate attacks that are aware of the use of the overlay network. [GENEVE-REQS] describes requirements to handle attacks. It is possible to use an overlay header in combination with IPsec or other encrypted traffic sessions, but this adds the requirement for authentication infrastructure and may reduce packet transfer performance. The use of an overlay header may also be deployed as a mechanism to manage encrypted traffic streams on the network-by-network service providers. Additional extension mechanisms to provide integrity and/or privacy protections are being investigated for overlay encapsulations. Section 7 of [RFC7348] describes some of

the security issues possible when deploying VXLAN on Layer 2 networks. Rogue endpoints can join the multicast groups that carry broadcast traffic, for example.

3.1.1. Monitoring Customer Access

Hosted applications that allow some level of customer-management access may also require monitoring by the hosting service provider. Monitoring could include access-control restrictions such as authentication, authorization, and accounting for filtering and firewall rules to ensure they are continuously met. Customer access may occur on multiple levels, including user-level and administrative access. The hosting service provider may need to monitor access through either session monitoring or log evaluation to ensure security SLAs for access management are met. The use of session encryption to access hosted environments limits access restrictions to the metadata described below. Monitoring and filtering may occur at a:

2-tuple: IP level with source and destination IP addresses alone, or

5-tuple: IP and protocol level with a source IP address, destination IP address, protocol number, source port number, and destination port number.

Session encryption at the application level, for example, TLS, currently allows access to the 5-tuple. IP-level encryption, such as IPsec in tunnel mode, prevents access to the original 5-tuple and may limit the ability to restrict traffic via filtering techniques. This shift may not impact all hosting service provider solutions as alternate controls may be used to authenticate sessions, or access may require that clients access such services by first connecting to the organization before accessing the hosted application. Shifts in access may be required to maintain equivalent access-control management. Logs may also be used for monitoring that access-control restrictions are met, but would be limited to the data that could be observed due to encryption at the point of log generation. Log analysis is out of scope for this document.

3.1.2. SP Content Monitoring of Applications

The following observations apply to any IT organization that is responsible for delivering services, whether to third parties, for example, as a web-based service, or to internal customers in an enterprise, e.g., a data-processing system that forms a part of the enterprise’s business.

Organizations responsible for the operation of a data center have many processes that access the contents of IP packets (passive methods of measurement, as defined in [RFC7799]). These processes are typically for service assurance or security purposes as part of their data-center operations.

Examples include:

- Network-Performance Monitoring / Application-Performance Monitoring
- Intrusion defense/prevention systems
- Malware detection
- Fraud monitoring
- Application DDOS protection
- Cyber-attack investigation
- Proof of regulatory compliance
- Data leakage prevention

Many application service providers simply terminate sessions to/from the Internet at the edge of the data center in the form of SSL/TLS offload in the load balancer. Not only does this reduce the load on application servers, it simplifies the processes to enable monitoring of the session content.

However, in some situations, encryption deeper in the data center may be necessary to protect personal information or in order to meet industry regulations, e.g., those set out by the Payment Card Industry (PCI). In such situations, various methods have been used to allow service assurance and security processes to access unencrypted data. These include SSL/TLS decryption in dedicated units, which then forward packets to SP-controlled tools, or real-time or post-capture decryption in the tools themselves. A number of these tools provide passive decryption by providing the monitoring device with the server’s private key. The move to increased use of the forward-secret key exchange mechanism impacts the use of these techniques.

Operators of data centers may also maintain packet recordings in order to be able to investigate attacks, breaches of internal processes, etc. In some industries, organizations may be legally required to maintain such information for compliance purposes.

Investigations of this nature have used access to the unencrypted contents of the packet. Alternate methods to investigate attacks or breaches of process will rely on endpoint information, such as logs. As previously noted, logs often lack complete information, and this is seen as a concern resulting in some relying on session access for additional information.

Application service providers may offer content-level monitoring options to detect intellectual property leakage or other attacks. In service provider environments where Data Loss Prevention (DLP) has been implemented on the basis of the service provider having cleartext access to session streams, the use of encrypted streams prevents these implementations from conducting content searches for the keywords or phrases configured in the DLP system. DLP is often used to prevent the leakage of Personally Identifiable Information (PII) as well as financial account information, Personal Health Information (PHI), and PCI. If session encryption is terminated at a gateway prior to accessing these services, DLP on session data can still be performed. The decision of where to terminate encryption to hosted environments will be a risk decision made between the application service provider and customer organization according to their priorities. DLP can be performed at the server for the hosted application and on an end user's system in an organization as alternate or additional monitoring points of content; however, this is not frequently done in a service provider environment.

Application service providers, by their very nature, control the application endpoint. As such, much of the information gleaned from sessions is still available on that endpoint. However, when a gap exists in the application's logging and debugging capabilities, it has led the application service provider to access data in transport for monitoring and debugging.

3.2. Hosted Applications

Organizations are increasingly using hosted applications rather than in-house solutions that require maintenance of equipment and software. Examples include Enterprise Resource Planning (ERP) solutions, payroll service, time and attendance, travel and expense reporting, among others. Organizations may require some level of management access to these hosted applications and will typically require session encryption or a dedicated channel for this activity.

In other cases, hosted applications may be fully managed by a hosting service provider with SLA expectations for availability and performance as well as for security functions including malware detection. Due to the sensitive nature of these hosted environments, the use of encryption is already prevalent. Any impact may be

similar to an enterprise with tools being used inside of the hosted environment to monitor traffic. Additional concerns were not reported in the call for contributions.

3.2.1. Monitoring Managed Applications

Performance, availability, and other aspects of an SLA are often collected through passive monitoring. For example:

- Availability: ability to establish connections with hosts to access applications and to discern the difference between network- or host-related causes of unavailability.
- Performance: ability to complete transactions within a target response time and to discern the difference between network- or host-related causes of excess response time.

Here, as with all passive monitoring, the accuracy of inferences is dependent on the cleartext information available, and encryption would tend to reduce the information and, therefore, the accuracy of each inference. Passive measurement of some metrics will be impossible with encryption that prevents inferring-packet correspondence across multiple observation points, such as for packet-loss metrics.

Application logging currently lacks detail sufficient to make accurate inferences in an environment with increased encryption, and so this constitutes a gap for passive performance monitoring (which could be closed if log details are enhanced in the future).

3.2.2. Mail Service Providers

Mail (application) service providers vary in what services they offer. Options may include a fully hosted solution where mail is stored external to an organization’s environment on mail service provider equipment or the service offering may be limited to monitor incoming mail to remove spam (Section 5.1), phishing attacks (Section 5.3), and malware (Section 5.6) before mail is directed to the organization’s equipment. In both of these cases, content of the messages and headers is monitored to detect and remove messages that are undesirable or that may be considered an attack.

STARTTLS should have zero effect on anti-spam efforts for SMTP traffic. Anti-spam services could easily be performed on an SMTP gateway, eliminating the need for TLS decryption services. The impact to anti-spam service providers should be limited to a change in tools, where middleboxes were deployed to perform these functions.

Many efforts are emerging to improve user-to-user encryption, including promotion of PGP and newer efforts such as Dark Mail [DarkMail]. Of course, content-based spam filtering will not be possible on encrypted content.

3.3. Data Storage

Numerous service offerings exist that provide hosted storage solutions. This section describes the various offerings and details the monitoring for each type of service and how encryption may impact the operational and security monitoring performed.

Trends in data storage encryption for hosted environments include a range of options. The following list is intentionally high-level to describe the types of encryption used in coordination with data storage that may be hosted remotely, meaning the storage is physically located in an external data center requiring transport over the Internet. Options for monitoring will vary with each encryption approach described below. In most cases, solutions have been identified to provide encryption while ensuring management capabilities were maintained through logging or other means.

3.3.1. Object-Level Encryption

For higher security and/or privacy of data and applications, options that provide end-to-end encryption of the data from the user's desktop or server to the storage platform may be preferred. This description includes any solution that encrypts data at the object level, not the transport level. Encryption of data may be performed with libraries on the system or at the application level, which includes file-encryption services via a file manager. Object-level encryption is useful when data storage is hosted or scenarios when the storage location is determined based on capacity or based on a set of parameters to automate decisions. This could mean that large datasets accessed infrequently could be sent to an off-site storage platform at an external hosting service, data accessed frequently may be stored locally, or the decision of where to store datasets could be based on the transaction type. Object-level encryption is grouped separately for the purpose of this document since data may be stored in multiple locations including off-site remote storage platforms. If session encryption is also used, the protocol is likely to be TLS.

Impacts to monitoring may include access to content inspection for data-leakage prevention and similar technologies, depending on their placement in the network.

3.3.1.1. Monitoring for Hosted Storage

Monitoring of hosted storage solutions that use host-level (object) encryption is described in this subsection. Host-level encryption can be employed for backup services and occasionally for external storage services (operated by a third party) when internal storage limits are exceeded.

Monitoring of data flows to hosted storage solutions is performed for security and operational purposes. The security monitoring may be to detect anomalies in the data flows that could include changes to destination, the amount of data transferred, or alterations in the size and frequency of flows. Operational considerations include capacity and availability monitoring.

3.3.2. Disk Encryption, Data at Rest (DAR)

There are multiple ways to achieve full disk encryption for stored data. Encryption may be performed on data to be stored while in transit close to the storage media with solutions like Controller Based Encryption (CBE) or in the drive system with Self-Encrypting Drives (SEDs). Session encryption is typically coupled with encryption of these data at rest (DAR) solutions to also protect data in transit. Transport encryption is likely via TLS.

3.3.2.1. Monitoring Session Flows for DAR Solutions

Monitoring for transport of data-to-storage platforms, where object-level encryption is performed close to or on the storage platform, is similar to that described in Section 3.3.1.1. The primary difference for these solutions is the possible exposure of sensitive information, which could include privacy-related data, financial information, or intellectual property if session encryption via TLS is not deployed. Session encryption is typically used with these solutions, but that decision would be based on a risk assessment. There are use cases where DAR or disk-level encryption is required. Examples include preventing exposure of data if physical disks are stolen or lost. In the case where TLS is in use, monitoring and the exposure of data is limited to a 5-tuple.

3.3.3. Cross-Data-Center Replication Services

Storage services also include data replication, which may occur between data centers and may leverage Internet connections to tunnel traffic. The traffic may use an Internet Small Computer System Interface (iSCSI) [RFC7143] or Fibre Channel over TCP/IP (FCIP) [RFC7146] encapsulated in IPsec. Either transport or tunnel mode may be used for IPsec depending upon the termination points of the IPsec

session, if it is from the storage platform itself or from a gateway device at the edge of the data center, respectively.

3.3.3.1. Monitoring IPsec for Data Replication Services

The monitoring of data flows between data centers (for data replication) may be performed for security and operational purposes and would typically concentrate more on operational aspects since these flows are essentially virtual private networks (VPNs) between data centers. Operational considerations include capacity and availability monitoring. The security monitoring may be to detect anomalies in the data flows, similar to what was described in Section 3.3.1.1. If IPsec tunnel mode is in use, monitoring is limited to a 2-tuple; with transport mode, it’s limited to a 5-tuple.

4. Encryption for Enterprises

Encryption of network traffic within the private enterprise is a growing trend, particularly in industries with audit and regulatory requirements. Some enterprise-internal networks are almost completely TLS and/or IPsec encrypted.

For each type of monitoring, different techniques and access to parts of the data stream are part of current practice. As we transition to an increased use of encryption, alternate methods of monitoring for operational purposes may be necessary to reduce the practice of breaking encryption (other policies may apply in some enterprise settings).

4.1. Monitoring Practices of the Enterprise

Large corporate enterprises are the owners of the platforms, data, and network infrastructure that provide critical business services to their user communities. As such, these enterprises are responsible for all aspects of the performance, availability, security, and quality of experience for all user sessions. In many such enterprises, users are required to consent to the enterprise monitoring all their activities as a condition of employment. Subsections of Section 4 discuss techniques that access data beyond the data-link, network, and transport-level headers typically used in service provider networks since the corporate enterprise owns the data. These responsibilities break down into three basic areas:

1. Security Monitoring and Control

2. Application-Performance Monitoring and Reporting

3. Network Diagnostics and Troubleshooting

In each of the above areas, technical support teams utilize collection, monitoring, and diagnostic systems. Some organizations currently use attack methods such as replicated TLS server RSA private keys to decrypt passively monitored copies of encrypted TLS packet streams.

For an enterprise to avoid costly application down time and deliver expected levels of performance, protection, and availability, some forms of traffic analysis, sometimes including examination of packet payloads, are currently used.

4.1.1. Security Monitoring in the Enterprise

Enterprise users are subject to the policies of their organization and the jurisdictions in which the enterprise operates. As such, proxies may be in use to:

1. intercept outbound session traffic to monitor for intellectual property leakage (by users, malware, and trojans),

2. detect viruses/malware entering the network via email or web traffic,

3. detect malware/trojans in action, possibly connecting to remote hosts,

4. detect attacks (cross-site scripting and other common web-related attacks),

5. track misuse and abuse by employees,

6. restrict the types of protocols permitted to/from the entire corporate environment, and

7. detect and defend against Internet DDoS attacks, including both volumetric and Layer 7 attacks.

A significant portion of malware hides its activity within TLS or other encryption protocols. This includes lateral movement, Command and Control (C&C), and Data Exfiltration.

The impact to a fully encrypted internal network would include cost and possible loss of detection capabilities associated with the transformation of the network architecture and tools for monitoring. The capabilities of detection through traffic fingerprinting, logging, host-level transaction monitoring, and flow analysis would vary depending on access to a 2-tuple or 5-tuple in the network as well.

Security monitoring in the enterprise may also be performed at the endpoint with numerous current solutions that mitigate the same problems as some of the above-mentioned solutions. Since the software agents operate on the device, they are able to monitor traffic before it is encrypted, monitor for behavior changes and lock down devices to use only the expected set of applications. Session encryption does not affect these solutions. Some might argue that scaling is an issue in the enterprise, but some large enterprises have used these tools effectively.

Use of bring-your-own-device (BYOD) policies within organizations may limit the scope of monitoring permitted with these alternate solutions. Network endpoint assessment (NEA) or the use of virtual hosts could help to bridge the monitoring gap.

4.1.2. Monitoring Application Performance in the Enterprise

There are two main goals of monitoring:

1. Assess traffic volume on a per-application basis for billing, capacity planning, optimization of geographical location for servers or proxies, and other goals.

2. Assess performance in terms of application response time and user-perceived response time.

Network-based application-performance monitoring tracks application response time by user and by URL, which is the information that the application owners and the lines of business request. CDNs add complexity in determining the ultimate endpoint destination. By their very nature, such information is obscured by CDNs and encrypted protocols, adding a new challenge for troubleshooting network and application problems. URL identification allows the application support team to do granular, code-level troubleshooting at multiple tiers of an application.

New methodologies to monitor user-perceived response time and to separate network from server time are evolving. For example, the IPv6 Destination Option Header (DOH) implementation of Performance and Diagnostic Metrics (PDM) [RFC8250] will provide this. Using PDM with IPsec Encapsulating Security Payload (ESP) Transport Mode requires placement of the PDM DOH within the ESP-encrypted payload to avoid leaking timing and sequence number information that could be useful to an attacker. Use of PDM DOH also may introduce some security weaknesses, including a timing attack, as described in Section 4 of [RFC8250]. For these and other reasons, [RFC8250]

requires that the PDM DOH option be explicitly turned on by administrative action in each host where this measurement feature will be used.

4.1.3. Diagnostics and Troubleshooting for Enterprise Networks

One primary key to network troubleshooting is the ability to follow a transaction through the various tiers of an application in order to isolate the fault domain. A variety of factors relating to the structure of the modern data center and multi-tiered application have made it difficult to follow a transaction in network traces without the ability to examine some of the packet payload. Alternate methods, such as log analysis, need improvement to fill this gap.

4.1.3.1. Address Sharing (NAT)

CDNs, NATs, and Network Address and Port Translators (NAPTs) obscure the ultimate endpoint designation (see [RFC6269] for types of address sharing and a list of issues). Troubleshooting a problem for a specific end user requires finding information such as the IP address and other identifying information so that their problem can be resolved in a timely manner.

NAT is also frequently used by lower layers of the data-center infrastructure. Firewalls, load balancers, web servers, app servers, and middleware servers all regularly NAT the source IP of packets. Combine this with the fact that users are often allocated randomly by load balancers to all these devices, and the network troubleshooter is often left with very few options in today’s environment due to poor logging implementations in applications. As such, network troubleshooting is used to trace packets at a particular layer, decrypt them, and look at the payload to find a user session.

This kind of bulk packet capture and bulk decryption is frequently used when troubleshooting a large and complex application. Endpoints typically don’t have the capacity to handle this level of network packet capture, so out-of-band networks of robust packet brokers and network sniffers that use techniques such as copies of TLS RSA private keys accomplish this task today.

4.1.3.2. TCP Pipelining / Session Multiplexing

TCP pipelining / session multiplexing used mainly by middleboxes today allows for multiple end-user sessions to share the same TCP connection. This raises several points of interest with an increased use of encryption. TCP session multiplexing should still be possible when TLS or TCPcrypt is in use since the TCP header information is exposed, leaving the 5-tuple accessible. The use of TCP session

multiplexing of an IP-layer encryption, e.g., IPsec, that only exposes a 2-tuple would not be possible. Troubleshooting capabilities with encrypted sessions from the middlebox may limit troubleshooting to the use of logs from the endpoints performing the TCP multiplexing or from the middleboxes prior to any additional encryption that may be added to tunnel the TCP multiplexed traffic.

Increased use of HTTP/2 will likely further increase the prevalence of session multiplexing, both on the Internet and in the private data center. HTTP pipelining requires both the client and server to participate; visibility of packets once encrypted will hide the use of HTTP pipelining for any monitoring that takes place outside of the endpoint or proxy solution. Since HTTP pipelining is between a client and server, logging capabilities may require improvement in some servers and clients for debugging purposes if this is not already possible. Visibility for middleboxes includes anything exposed by TLS and the 5-tuple.

4.1.3.3. HTTP Service Calls

When an application server makes an HTTP service call to back-end services on behalf of a user session, it uses a completely different URL and a completely different TCP connection. Troubleshooting via network trace involves matching up the user request with the HTTP service call. Some organizations do this today by decrypting the TLS packet and inspecting the payload. Logging has not been adequate for their purposes.

4.1.3.4. Application-Layer Data

Many applications use text formats such as XML to transport data or application-level information. When transaction failures occur and the logs are inadequate to determine the cause, network and application teams work together, each having a different view of the transaction failure. Using this troubleshooting method, the network packet is correlated with the actual problem experienced by an application to find a root cause. The inability to access the payload prevents this method of troubleshooting.

4.2. Techniques for Monitoring Internet-Session Traffic

Corporate networks commonly monitor outbound session traffic to detect or prevent attacks as well as to guarantee service-level expectations. In some cases, alternate options are available when encryption is in use through a proxy or a shift to monitoring at the endpoint. In both cases, scaling is a concern, and advancements to support this shift in monitoring practices will assist the deployment of end-to-end encryption.

Some DLP tools intercept traffic at the Internet gateway or proxy services with the ability to MITM encrypted session traffic (HTTP/TLS). These tools may monitor for key words important to the enterprise including business-sensitive information such as trade secrets, financial data, PII, or PHI. Various techniques are used to intercept HTTP/TLS sessions for DLP and other purposes and can be misused as described in "Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS)" [RFC7457] (see Section 2.8). Note: many corporate policies allow access to personal financial and other sites for users without interception. Another option is to terminate a TLS session prior to the point where monitoring is performed. Aside from exposing user information to the enterprise, MITM devices often are subject to severe security defects, which can lead to exposure of user data to attackers outside the enterprise user data [UserData]. In addition, implementation errors in middleboxes have led to major difficulties in deploying new versions of security protocols such as TLS [Ben17a] [Ben17b] [Res17a] [Res17b].

Monitoring traffic patterns for anomalous behavior such as increased flows of traffic that could be bursty at odd times or flows to unusual destinations (small or large amounts of traffic) is common. This traffic may or may not be encrypted, and various methods of encryption or just obfuscation may be used.

Web-filtering devices are sometimes used to allow only access to well-known sites found to be legitimate and free of malware on last check by a web-filtering service company. One common example of web filtering in a corporate environment is blocking access to sites that are not well known to these tools for the purpose of blocking malware; this may be noticeable to those in research who are unable to access colleagues' individual sites or new websites that have not yet been screened. In situations where new sites are required for access, they can typically be added after notification by the user or log alerts and review. Account access for personal mail may be blocked in corporate settings to prevent another vector for malware from entering as well as to prevent intellectual property leaks out of the network. This method remains functional with increased use of encryption and may be more effective at preventing malware from entering the network. Some enterprises may be more aggressive in their filtering and monitoring policy, causing undesirable outcomes. Web-filtering solutions monitor and potentially restrict access based on the destination URL (when available), server name, IP address, or DNS name. A complete URL may be used in cases where access restrictions vary for content on a particular site or for the sites hosted on a particular server. In some cases, the enterprise may use a proxy to access this additional information based on their policy. This type of restriction is intended to be transparent to users in a

corporate setting as the typical corporate user does not access sites that are not well known to these tools. However, the mechanisms that these web filters use to do monitoring and enforcement have the potential to cause access issues or other user-visible failures.

Desktop DLP tools are used in some corporate environments as well. Since these tools reside on the desktop, they can intercept traffic before it is encrypted and may provide a continued method for monitoring leakage of intellectual property from the desktop to the Internet or attached devices.

DLP tools can also be deployed by network service providers, as they have the vantage point of monitoring all traffic paired with destinations off the enterprise network. This makes an effective solution for enterprises that allow "bring-your-own" devices when the traffic is not encrypted and for devices outside the desktop category (such as mobile phones) that are used on corporate networks nonetheless.

Enterprises may wish to reduce the traffic on their Internet access facilities by monitoring requests for within-policy content and caching it. In this case, repeated requests for Internet content spawned by URLs in email trade newsletters or other sources can be served within the enterprise network. Gradual deployment of end-to-end encryption would tend to reduce the cacheable content over time, owing to concealment of critical headers and payloads. Many forms of enterprise-performance management may be similarly affected. It should be noted that transparent caching is considered an anti-pattern.

5. Security Monitoring for Specific Attack Types

Effective incident response today requires collaboration at Internet scale. This section will only focus on efforts of collaboration at Internet scale that are dedicated to specific attack types. They may require new monitoring and detection techniques in an increasingly encrypted Internet. As mentioned previously, some service providers have been interfering with STARTTLS to prevent session encryption to be able to perform functions they are used to (injecting ads, monitoring, etc.). By detailing the current monitoring methods used for attack detection and response, this information can be used to devise new monitoring methods that will be effective in the changed Internet via collaboration and innovation.

Changes to improve encryption or to deploy OS methods have little impact on the detection of malicious actors. Malicious actors have had access to strong encryption for quite some time. Incident responders, in many cases, have developed techniques to locate

malicious traffic within encrypted sessions. The following section will note some examples where detection and mitigation of such traffic has been successful.

5.1. Mail Abuse and Spam

The largest operational effort to prevent mail abuse is through the Messaging, Malware, Mobile Anti-Abuse Working Group (M3AAWG) [M3AAWG]. Mail abuse is combatted directly with mail administrators who can shut down or stop continued mail abuse originating from large-scale providers that participate in using the Abuse Reporting Format (ARF) agents standardized in the IETF [RFC5965] [RFC6430] [RFC6590] [RFC6591] [RFC6650] [RFC6651] [RFC6652]. The ARF agent directly reports abuse messages to the appropriate service provider who can take action to stop or mitigate the abuse. Since this technique uses the actual message, the use of SMTP over TLS between mail gateways will not affect its usefulness. As mentioned previously, SMTP over TLS only protects data while in transit, and the messages may be exposed on mail servers or mail gateways if a user-to-user encryption method is not used. Current user-to-user message encryption methods on email (S/MIME and PGP) do not encrypt the email header information used by ARF and the service provider operators in their efforts to mitigate abuse.

Another effort, "Domain-based Message Authentication, Reporting, and Conformance (DMARC)" [RFC7489], is a mechanism for policy distribution that enables increasingly strict handling of messages that fail authentication checks, ranging from no action, through altered delivery, up to message rejection. DMARC is also not affected by the use of STARTTLS.

5.2. Denial of Service

Responses to Denial-of-Service (DoS) attacks are typically coordinated by the service provider community with a few key vendors who have tools to assist in the mitigation efforts. Traffic patterns are determined from each DoS attack to stop or rate limit the traffic flows with patterns unique to that DoS attack.

Data types used in monitoring traffic for DDoS are described in the documents in development by the DDoS Open Threat Signaling (DOTS) [DOTS] Working Group. The impact of encryption can be understood from their documented use cases [DDOS-USECASE].

Data types used in DDoS attacks have been detailed in the Incident Object Description Exchange Format (IODEF) Guidance document (see [RFC8274], Appendix B.2) with the help of several members of the service provider community. The examples provided are intended to

help identify the useful data in detecting and mitigating these attacks independent of the transport and protocol descriptions in the documents.

5.3. Phishing

Investigations and responses to phishing attacks follow well-known patterns, requiring access to specific fields in email headers as well as content from the body of the message. When reporting phishing attacks, the recipient has access to each field as well as the body to make content reporting possible, even when end-to-end encryption is used. The email header information is useful to identify the mail servers and accounts used to generate or relay the attack messages in order to take the appropriate actions. The content of the message often includes an embedded attack that may be in an infected file or may be a link that results in the download of malware to the user's system.

Administrators often find it helpful to use header information to track down similar messages in their mail queue or in users' inboxes to prevent further infection. Combinations of To:, From:, Subject:, and Received: from header information might be used for this purpose. Administrators may also search for document attachments of the same name or size or that contain a file with a matching hash to a known phishing attack. Administrators might also add URLs contained in messages to block lists locally, or this may also be done by browser vendors through larger-scale efforts like that of the Anti-Phishing Working Group (APWG). See "Coordinating Attack Response at Internet Scale (CARIS) Workshop Report" [RFC8073] for additional information and pointers to the APWG's efforts on anti-phishing.

A full list of the fields used in phishing attack incident responses can be found in RFC 5901. Future plans to increase privacy protections may limit some of these capabilities if some email header fields are encrypted, such as the To:, From:, and Subject: header fields. This does not mean that those fields should not be encrypted, only that we should be aware of how they are currently used.

Some products protect users from phishing by maintaining lists of known phishing domains (such as misspelled bank names) and blocking access. This can be done by observing DNS, cleartext HTTP, or Server Name Indication (SNI) in TLS, in addition to analyzing email. Alternate options to detect and prevent phishing attacks may be needed. More recent examples of data exchanged in spear phishing attacks has been detailed in the IODEF Guidance document (see [RFC8274], Appendix B.3).

5.4. Botnets

Botnet detection and mitigation is complex as botnets may involve hundreds or thousands of hosts with numerous C&C servers. The techniques and data used to monitor and detect each may vary. Connections to C&C servers are typically encrypted; therefore, a move to an increasingly encrypted Internet may not affect the detection and sharing methods used.

5.5. Malware

Techniques for the detection and monitoring of malware vary. As mentioned in Section 4, malware monitoring may occur at gateways to the organization analyzing email and web traffic. These services can also be provided by service providers, changing the scale and location of this type of monitoring. Additionally, incident responders may identify attributes unique to types of malware to help track down instances by their communication patterns on the Internet or by alterations to hosts and servers.

Data types used in malware investigations have been summarized in an example of the IODEF Guidance document (see [RFC8274], Appendix B.3).

5.6. Spoofed-Source IP Address Protection

The IETF has reacted to spoofed-source IP address-based attacks, recommending the use of network ingress filtering in BCP 38 [RFC2827] and of the unicast Reverse Path Forwarding (uRPF) mechanism [RFC3704]. But uRPF suffers from limitations regarding its granularity: a malicious node can still use a spoofed IP address included inside the prefix assigned to its link. Source Address Validation Improvement (SAVI) mechanisms try to solve this issue. Basically, a SAVI mechanism is based on the monitoring of a specific address assignment/management protocol (e.g., Stateless Address Autoconfiguration (SLAAC) [RFC4862], Secure Neighbor Discovery (SEND) [RFC3971], and DHCPv4/v6 [RFC2131][RFC3315]) and, according to this monitoring, sets up a filtering policy allowing only the IP flows with a correct source IP address (i.e., any packet with a source IP address from a node not owning it is dropped). The encryption of parts of the address assignment/management protocols, critical for SAVI mechanisms, can result in a dysfunction of the SAVI mechanisms.

5.7. Further Work

Although incident response work will continue, new methods to prevent system compromise through security automation and continuous monitoring [SACM] may provide alternate approaches where system security is maintained as a preventative measure.

6. Application-Based Flow Information Visible to a Network

This section describes specific techniques used in monitoring applications that are visible to the network if a 5-tuple is exposed and as such can potentially be used as input for future network-management approaches. It also includes an overview of IP Flow Information Export (IPFIX), a flow-based protocol used to export information about network flows.

6.1. IP Flow Information Export

Many of the accounting, monitoring, and measurement tasks described in this document, especially in Sections 2.3.2, 3.1.1, 4.1.3, 4.2, and 5.2, use the IPFIX protocol [RFC7011] for export and storage of the monitored information. IPFIX evolved from the widely deployed NetFlow protocol [RFC3954], which exports information about flows identified by 5-tuple. While NetFlow was largely concerned with exporting per-flow byte and packet counts for accounting purposes, IPFIX's extensible Information Model [RFC7012] provides a variety of Information Elements (IEs) [IPFIX-IANA] for representing information above and below the traditional network-layer flow information. Enterprise-specific IEs allow exporter vendors to define their own non-standard IEs as well, and many of these are driven by header and payload inspection at the Metering Process.

While the deployment of encryption has no direct effect on the use of IPFIX, certain defined IEs may become unavailable when the Metering Process observing the traffic cannot decrypt former cleartext information. For example, HTTPS renders HTTP header analysis impossible, so IEs derived from the header (e.g., httpContentType, httpUserAgent) cannot be exported.

The collection of IPFIX data itself, of course, provides a point of centralization for information that is potentially business and privacy critical. The IPFIX File Format specification [RFC5655] recommends encryption for this data at rest, and the IP Flow Anonymization specification [RFC6235] defines a metadata format for describing the anonymization functions applied to an IPFIX dataset, if anonymization is employed for data sharing of IPFIX information between enterprises or network operators.

6.2. TLS Server Name Indication

When initiating the TLS handshake, the client may provide an extension field (server_name) that indicates the server to which it is attempting a secure connection. TLS SNI was standardized in 2003 to enable servers to present the "correct TLS certificate" to clients in a deployment of multiple virtual servers hosted by the same server

infrastructure and IP address. Although this is an optional extension, it is today supported by all modern browsers, web servers, and developer libraries. Akamai [Nygren] reports that many of their customers see client TLS SNI usage over 99%. It should be noted that HTTP/2 introduces the Alt-SVC method for upgrading the connection from HTTP/1 to either unencrypted or encrypted HTTP/2. If the initial HTTP/1 request is unencrypted, the destination alternate service name can be identified before the communication is potentially upgraded to encrypted HTTP/2 transport. HTTP/2 requires the TLS implementation to support the SNI extension (see Section 9.2 of [RFC7540]). It is also worth noting that [RFC7838] "allows an origin server to nominate additional means of interacting with it on the network", while [RFC8164] allows for a URI to be accessed with HTTP/2 and TLS using Opportunistic Security (on an experimental basis).

This information is only available if the client populates the SNI extension. Doing so is an optional part of the TLS standard, and as stated above, this has been implemented by all major browsers. Due to its optional nature, though, existing network filters that examine a TLS ClientHello for an SNI extension cannot expect to always find one. "SNI Encryption in TLS Through Tunneling" [SNI-TLS] has been adopted by the TLS Working Group, which provides solutions to encrypt SNI. As such, there will be an option to encrypt SNI in future versions of TLS. The per-domain nature of SNI may not reveal the specific service or media type being accessed, especially where the domain is of a provider offering a range of email, video, web pages, etc. For example, certain blog or social network feeds may be deemed "adult content", but the SNI will only indicate the server domain rather than a URL path.

There are additional issues for identification of content using SNI: [RFC7540] includes connection coalescing, [RFC8336] defines the ORIGIN frame, and the proposal outlined in [HTTP2-CERTS] will increase the difficulty of passive monitoring.

6.3. Application-Layer Protocol Negotiation (ALPN)

ALPN is a TLS extension that may be used to indicate the application protocol within the TLS session. This is likely to be of more value to the network where it indicates a protocol dedicated to a particular traffic type (such as video streaming) rather than a multi-use protocol. ALPN is used as part of HTTP/2 'h2', but will not indicate the traffic types that may make up streams within an HTTP/2 multiplex. ALPN is sent cleartext in the ClientHello, and the server returns it in Encrypted Extensions in TLS 1.3.

6.4. Content Length, Bitrate, and Pacing

The content length of encrypted traffic is effectively the same as that of the cleartext. Although block ciphers utilize padding, this makes a negligible difference. Bitrate and pacing are generally application specific and do not change much when the content is encrypted. Multiplexed formats (such as HTTP/2 and QUIC [QUIC]) may, however, incorporate several application streams over one connection, which makes the bitrate/pacing no longer application specific. Also, packet padding is available in HTTP/2, TLS 1.3, and many other protocols. Traffic analysis is made more difficult by such countermeasures.

7. Effect of Encryption on the Evolution of Mobile Networks

Transport header encryption prevents the use of transit proxies in the center of the network and the use of some edge proxies by preventing the proxies from taking action on the stream. It may be that the claimed benefits of such proxies could be achieved by end-to-end client and server optimizations, distribution using CDNs, plus the ability to continue connections across different access technologies (across dynamic user IP addresses). The following aspects should be considered in this approach:

1. In a wireless mobile network, the delay and channel capacity per user and sector varies due to coverage, contention, user mobility, scheduling balances fairness, capacity, and service QoE. If most users are at the cell edge, the controller cannot use more-complex Quadrature Amplitude Modulation (QAM), thus reducing total cell capacity; similarly, if a Universal Mobile Telecommunications System (UMTS) edge is serving some number of CS-Voice Calls, the remaining capacity for packet services is reduced.

2. Mobile wireless networks service inbound roamers (users of Operator A in the foreign network of Operator B) by backhauling their traffic through the network (from Operator B to Operator A) and then serving them through the P-Gateway (PGW), General Packet Radio Service (GPRS) Support Node (GGSN), CDN, etc., of Operator A (the user’s home operator). Increasing window sizes to compensate for the path RTT will have the limitations outlined earlier for TCP. The outbound roamer scenario has a similar TCP performance impact.

3. Issues in deploying CDNs in Radio Access Networks (RANs) include decreasing the client-server control loop that requires deploying CDNs / Cloud functions that terminate encryption closer to the edge. In Cellular RAN, the user IP traffic is encapsulated into

GPRS Tunneling Protocol-User Plane (GTP-U in UMTS and LTE) tunnels to handle user mobility; the tunnels terminate in APN/GGSN/PGW that are in central locations. One user's traffic may flow through one or more APN's (for example, Internet APN, Roaming APN for Operator X, Video-Service APN, OnDeckAPN, etc.). The scope of operator private IP addresses may be limited to specific APNs. Since CDNs generally operate on user IP flows, deploying them would require enhancing them with tunnel translation, tunnel-management functions, etc.

4. While CDNs that decrypt flows or split connection proxies (similar to split TCP) could be deployed closer to the edges to reduce control-loop RTT, with transport header encryption, such CDNs perform optimization functions only for partner client flows. Therefore, content from some Small-Medium Businesses (SMBs) would not get such CDN benefits.

8. Response to Increased Encryption and Looking Forward

As stated in [RFC7258], "an appropriate balance [between network management and pervasive monitoring mitigations] will emerge over time as real instances of this tension are considered." Numerous operators made it clear in their response to this document that they fully support strong encryption and providing privacy for end users; this is a common goal. Operators recognize that not all the practices documented need to be supported going forward, either because of the risk to end-user privacy or because alternate technologies and tools have already emerged. This document is intended to support network engineers and other innovators to work toward solving network and security management problems with protocol designers and application developers in new ways that facilitate adoption of strong encryption rather than preventing the use of encryption. By having the discussions on network and security management practices with application developers and protocol designers, each side of the debate can understand each other's goals, work toward alternate solutions, and disband with practices that should no longer be supported. A goal of this document is to assist the IETF in understanding some of the current practices so as to identify new work items for IETF-related use cases that can facilitate the adoption of strong session encryption and support network and security management.

9. Security Considerations

There are no additional security considerations as this is a summary and does not include a new protocol or functionality.

10. IANA Considerations

This document has no IANA actions.

11. Informative References

[ACCORD] IETF, "Alternatives to Content Classification for Operator Resource Deployment (accord) (BOF)", IETF-95 Proceedings, April 2016, <https://www.ietf.org/proceedings/95/accord.html>.

[Ben17a] Benjamin, D., "Chrome Data", Presentation before the TLS WG at IETF 100, November 2017, <https://datatracker.ietf.org/meeting/100/materials/slides-100-tls-sessa-tls13/>.

[Ben17b] Benjamin, D., "Subject: Additional TLS 1.3 results from Chrome", message to the TLS mailing list, 18 December 2017, <https://www.ietf.org/mail-archive/web/tls/current/msg25168.html>.

[CAIDA] CAIDA, "The CAIDA USCD Anonymized Internet Traces 2016 Dataset", <http://www.caida.org/data/passive/passive_2016_dataset.xml>.

[DarkMail] "Dark Mail Technical Alliance", <https://darkmail.info/>.

[DDOS-USECASE] Dobbins, R., Migault, D., Fouant, S., Moskowitz, R., Teague, N., Xia, L., and K. Nishizuka, "Use cases for DDoS Open Threat Signaling", Work in Progress, draft-ietf-dots-use-cases-16, July 2018.

[DOTS] IETF, "DDoS Open Threat Signaling (dots)", <https://datatracker.ietf.org/wg/dots/charter>.

[EFF2014] Hoffman-Andrews, J., "ISPs Removing Their Customers' Email Encryption", November 2014, <https://www.eff.org/deeplinks/2014/11/starttls-downgrade-attacks>.

[Enrich] Narseo Vallina-Rodriguez, N., Sundaresan, S., Kreibich, C., and V. Paxson, "Header Enrichment or ISP Enrichment: Emerging Privacy Threats in Mobile Networks", Proceedings of the ACM SIGCOMM Workshop on Hot Topics in Middleboxes and Network Function Virtualization, pp. 23-30, DOI 10.1145/2785989.2786002, August 2015.

[GENEVE-REQS]
Migault, D., Boutros, S., Wing, D., and S. Krishnan, "Geneve Protocol Security Requirements", Work in Progress, draft-mglt-nvo3-geneve-security-requirements-03, February 2018.

[HTTP2-CERTS]
Bishop, M., Sullivan, N., and M. Thomson, "Secondary Certificate Authentication in HTTP/2", Work in Progress, draft-ietf-httpbis-http2-secondary-certs-02, June 2018.

[IPFIX-IANA]
IANA, "IP Flow Information Export (IPFIX) Entities", <https://www.iana.org/assignments/ipfix/>.

[JNSLP] Eskens, S., "10 Standards for Oversight and Transparency of National Intelligence Services", Surveillance, Vol. 8, No. 3, July 2016, <http://jnslp.com/?s=10+Standards+for+Oversight+and+Transparency+of+National>.

[M3AAWG] M3AAWG, "Messaging, Malware and Mobile Anti-Abuse Working Group (M3AAWG)", <https://www.maawg.org/>.

[MIDDLEBOXES]
Dolson, D., Snellman, J., Boucadair, M., and C. Jacquenet, "An Inventory of Transport-centric Functions Provided by Middleboxes", Work in Progress, draft-dolson-transport-middlebox-03, June 2018.

[Nygren] Nygren, E., "Reaching toward Universal TLS SNI", Akamai Technologies, March 2017, <https://blogs.akamai.com/2017/03/reaching-toward-universal-tls-sni.html>.

[QUIC] IETF, "QUIC (quic)", <https://datatracker.ietf.org/wg/quic/charter/>.

[Res17a] Rescorla, E., "Subject: Preliminary data on Firefox TLS 1.3 Middlebox experiment", message to the TLS mailing list, 5 December 2017, <https://www.ietf.org/mail-archive/web/tls/current/msg25091.html>.

[Res17b] Rescorla, E., "Subject: More compatibility measurement results", message to the TLS mailing list, 22 December 2017, <https://www.ietf.org/mail-archive/web/tls/current/msg25179.html>.

[RFC1945] Berners-Lee, T., Fielding, R., and H. Frystyk, "Hypertext Transfer Protocol -- HTTP/1.0", RFC 1945, DOI 10.17487/RFC1945, May 1996, <https://www.rfc-editor.org/info/rfc1945>.

[RFC1958] Carpenter, B., Ed., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <https://www.rfc-editor.org/info/rfc1958>.

[RFC1984] IAB and IESG, "IAB and IESG Statement on Cryptographic Technology and the Internet", BCP 200, RFC 1984, DOI 10.17487/RFC1984, August 1996, <https://www.rfc-editor.org/info/rfc1984>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[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, <https://www.rfc-editor.org/info/rfc2474>.

[RFC2663] Srisuresh, P. and M. Holdrege, "IP Network Address Translator (NAT) Terminology and Considerations", RFC 2663, DOI 10.17487/RFC2663, August 1999, <https://www.rfc-editor.org/info/rfc2663>.

[RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, DOI 10.17487/RFC2775, February 2000, <https://www.rfc-editor.org/info/rfc2775>.

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

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <https://www.rfc-editor.org/info/rfc2827>.

[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, <https://www.rfc-editor.org/info/rfc3135>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <https://www.rfc-editor.org/info/rfc3315>.

[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, <https://www.rfc-editor.org/info/rfc3550>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <https://www.rfc-editor.org/info/rfc3704>.

[RFC3724] Kempf, J., Ed., Austein, R., Ed., and IAB, "The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture", RFC 3724, DOI 10.17487/RFC3724, March 2004, <https://www.rfc-editor.org/info/rfc3724>.

[RFC3954] Claise, B., Ed., "Cisco Systems NetFlow Services Export Version 9", RFC 3954, DOI 10.17487/RFC3954, October 2004, <https://www.rfc-editor.org/info/rfc3954>.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <https://www.rfc-editor.org/info/rfc3971>.

[RFC4787] Audet, F., Ed. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <https://www.rfc-editor.org/info/rfc4787>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <https://www.rfc-editor.org/info/rfc4862>.

[RFC5655] Trammell, B., Boschi, E., Mark, L., Zseby, T., and A. Wagner, "Specification of the IP Flow Information Export (IPFIX) File Format", RFC 5655, DOI 10.17487/RFC5655, October 2009, <https://www.rfc-editor.org/info/rfc5655>.

[RFC5965] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, DOI 10.17487/RFC5965, August 2010, <https://www.rfc-editor.org/info/rfc5965>.

[RFC6108] Chung, C., Kasyanov, A., Livingood, J., Mody, N., and B. Van Lieu, "Comcast's Web Notification System Design", RFC 6108, DOI 10.17487/RFC6108, February 2011, <https://www.rfc-editor.org/info/rfc6108>.

[RFC6235] Boschi, E. and B. Trammell, "IP Flow Anonymization Support", RFC 6235, DOI 10.17487/RFC6235, May 2011, <https://www.rfc-editor.org/info/rfc6235>.

[RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <https://www.rfc-editor.org/info/rfc6269>.

[RFC6430] Li, K. and B. Leiba, "Email Feedback Report Type Value: not-spam", RFC 6430, DOI 10.17487/RFC6430, November 2011, <https://www.rfc-editor.org/info/rfc6430>.

[RFC6455] Fette, I. and A. Melnikov, "The WebSocket Protocol", RFC 6455, DOI 10.17487/RFC6455, December 2011, <https://www.rfc-editor.org/info/rfc6455>.

[RFC6590] Falk, J., Ed. and M. Kucherawy, Ed., "Redaction of Potentially Sensitive Data from Mail Abuse Reports", RFC 6590, DOI 10.17487/RFC6590, April 2012, <https://www.rfc-editor.org/info/rfc6590>.

[RFC6591] Fontana, H., "Authentication Failure Reporting Using the Abuse Reporting Format", RFC 6591, DOI 10.17487/RFC6591, April 2012, <https://www.rfc-editor.org/info/rfc6591>.

[RFC6650] Falk, J. and M. Kucherawy, Ed., "Creation and Use of Email Feedback Reports: An Applicability Statement for the Abuse Reporting Format (ARF)", RFC 6650, DOI 10.17487/RFC6650, June 2012, <https://www.rfc-editor.org/info/rfc6650>.

[RFC6651] Kucherawy, M., "Extensions to DomainKeys Identified Mail (DKIM) for Failure Reporting", RFC 6651, DOI 10.17487/RFC6651, June 2012, <https://www.rfc-editor.org/info/rfc6651>.

[RFC6652] Kitterman, S., "Sender Policy Framework (SPF) Authentication Failure Reporting Using the Abuse Reporting Format", RFC 6652, DOI 10.17487/RFC6652, June 2012, <https://www.rfc-editor.org/info/rfc6652>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <https://www.rfc-editor.org/info/rfc7011>.

[RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <https://www.rfc-editor.org/info/rfc7012>.

[RFC7143] Chadalapaka, M., Satran, J., Meth, K., and D. Black, "Internet Small Computer System Interface (iSCSI) Protocol (Consolidated)", RFC 7143, DOI 10.17487/RFC7143, April 2014, <https://www.rfc-editor.org/info/rfc7143>.

[RFC7146] Black, D. and P. Koning, "Securing Block Storage Protocols over IP: RFC 3723 Requirements Update for IPsec v3", RFC 7146, DOI 10.17487/RFC7146, April 2014, <https://www.rfc-editor.org/info/rfc7146>.

[RFC7230] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, DOI 10.17487/RFC7230, June 2014, <https://www.rfc-editor.org/info/rfc7230>.

[RFC7234] Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Caching", RFC 7234, DOI 10.17487/RFC7234, June 2014, <https://www.rfc-editor.org/info/rfc7234>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/info/rfc7258>.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <https://www.rfc-editor.org/info/rfc7348>.

[RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", RFC 7435, DOI 10.17487/RFC7435, December 2014, <https://www.rfc-editor.org/info/rfc7435>.

[RFC7457] Sheffer, Y., Holz, R., and P. Saint-Andre, "Summarizing Known Attacks on Transport Layer Security (TLS) and Datagram TLS (DTLS)", RFC 7457, DOI 10.17487/RFC7457, February 2015, <https://www.rfc-editor.org/info/rfc7457>.

[RFC7489] Kucherawy, M., Ed. and E. Zwicky, Ed., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)", RFC 7489, DOI 10.17487/RFC7489, March 2015, <https://www.rfc-editor.org/info/rfc7489>.

[RFC7498] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for Service Function Chaining", RFC 7498, DOI 10.17487/RFC7498, April 2015, <https://www.rfc-editor.org/info/rfc7498>.

[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, <https://www.rfc-editor.org/info/rfc7525>.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <https://www.rfc-editor.org/info/rfc7540>.

[RFC7619] Smyslov, V. and P. Wouters, "The NULL Authentication Method in the Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 7619, DOI 10.17487/RFC7619, August 2015, <https://www.rfc-editor.org/info/rfc7619>.

[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, <https://www.rfc-editor.org/info/rfc7624>.

[RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <https://www.rfc-editor.org/info/rfc7665>.

[RFC7754] Barnes, R., Cooper, A., Kolkman, O., Thaler, D., and E. Nordmark, "Technical Considerations for Internet Service Blocking and Filtering", RFC 7754, DOI 10.17487/RFC7754, March 2016, <https://www.rfc-editor.org/info/rfc7754>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <https://www.rfc-editor.org/info/rfc7799>.

[RFC7826] Schulzrinne, H., Rao, A., Lanphier, R., Westerlund, M., and M. Stiemerling, Ed., "Real-Time Streaming Protocol Version 2.0", RFC 7826, DOI 10.17487/RFC7826, December 2016, <https://www.rfc-editor.org/info/rfc7826>.

[RFC7838] Nottingham, M., McManus, P., and J. Reschke, "HTTP Alternative Services", RFC 7838, DOI 10.17487/RFC7838, April 2016, <https://www.rfc-editor.org/info/rfc7838>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8073] Moriarty, K. and M. Ford, "Coordinating Attack Response at Internet Scale (CARIS) Workshop Report", RFC 8073, DOI 10.17487/RFC8073, March 2017, <https://www.rfc-editor.org/info/rfc8073>.

[RFC8164] Nottingham, M. and M. Thomson, "Opportunistic Security for HTTP/2", RFC 8164, DOI 10.17487/RFC8164, May 2017, <https://www.rfc-editor.org/info/rfc8164>.

[RFC8165] Hardie, T., "Design Considerations for Metadata Insertion", RFC 8165, DOI 10.17487/RFC8165, May 2017, <https://www.rfc-editor.org/info/rfc8165>.

[RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <https://www.rfc-editor.org/info/rfc8250>.

[RFC8274] Kampanakis, P. and M. Suzuki, "Incident Object Description Exchange Format Usage Guidance", RFC 8274, DOI 10.17487/RFC8274, November 2017, <https://www.rfc-editor.org/info/rfc8274>.

[RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <https://www.rfc-editor.org/info/rfc8300>.

[RFC8336] Nottingham, M. and E. Nygren, "The ORIGIN HTTP/2 Frame", RFC 8336, DOI 10.17487/RFC8336, March 2018, <https://www.rfc-editor.org/info/rfc8336>.

[SACM] IETF, "Security Automation and Continuous Monitoring (sacm)", <https://datatracker.ietf.org/wg/sacm/charter/>.

[SNI-TLS] Huitema, C. and E. Rescorla, "Issues and Requirements for SNI Encryption in TLS", Work in Progress, draft-ietf-tls-sni-encryption-03, May 2018.

[Snowden] Verble, J., "The NSA and Edward Snowden: Surveillance in the 21st Century", SIGCAS Computer & Society, Vol. 44, No. 3, DOI 10.1145/2684097.2684101, September 2014, <http://www.jjsylvia.com/bigdatacourse/wp-content/uploads/2016/04/p14-verble-1.pdf>.

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

[TS3GPP] 3GPP, "Non-Access-Stratum (NAS) protocol for Evolved Packet System (EPS); Stage 3", 3GPP TS 24.301, version 15.2.0, March 2018.

[UPCON] 3GPP, "User Plane Congestion Management", 3GPP Rel-13, September 2014, <http://www.3gpp.org/DynaReport/FeatureOrStudyItemFile-570029.htm>.

[UserData] Durumeric, Z., Ma, Z., Springall, D., Barnes, R., Sullivan, N., Bursztein, E., Bailey, M., Alex Halderman, J., and V. Paxson, "The Security Impact of HTTPS Interception", Network and Distributed Systems Symposium, February 2017, <http://dx.doi.org/10.14722/ndss.2017.23456>.

Acknowledgements

Thanks to our reviewers, Natasha Rooney, Kevin Smith, Ashutosh Dutta, Brandon Williams, Jean-Michel Combes, Nalini Elkins, Paul Barrett, Badri Subramanyan, Igor Lubashev, Suresh Krishnan, Dave Dolson, Mohamed Boucadair, Stephen Farrell, Warren Kumari, Alia Atlas, Roman Danyliw, Mirja Kuehlewind, Ines Robles, Joe Clarke, Kyle Rose, Christian Huitema, and Chris Morrow for their editorial and content suggestions. Surya K. Kovvali provided material for Section 7. Chris Morrow and Nik Teague provided reviews and updates specific to the DoS fingerprinting text. Brian Trammell provided the IPFIX text.

Authors' Addresses

Kathleen Moriarty (editor)
Dell EMC
176 South St
Hopkinton, MA
United States of America

Email: Kathleen.Moriarty@dell.com

Al Morton (editor)
AT&T Labs
200 Laurel Avenue South
Middletown, NJ 07748
United States of America

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acm@research.att.com