CDNs, Network Services and Encrypted Traffic – MaRNEW 2015

1. Introduction

In 2014, concerns over privacy of communication on the Internet accelerated the trend of websites serving all their content over HTTPS. In fact, the opinion that all the web traffic should be end-to-end encrypted and it is estimated that by early 2015, over 40% of web traffic was being served encrypted [1]. At the same time, some communities pointed out the collateral damage that would be caused if all the traffic used such encryption technology. Network operators use intermediate proxies for various services such as transparent caching, transcoding/transrating, parental guidance filtering, real-time network intelligence sharing, etc. It is feared that most such services would either cease to function in the encrypted communication environment or become very difficult to implement. Many operators believe that this would have a severe impact on their resource management practices. How to reconcile the desire to use end-to-end encryption while not making worthwhile network services entirely obsolete continues to be a topic of debate at the time of writing this paper.

One could classify intermediate proxy-based network services into two broad categories based on their data access needs:

1. **Content-aware services:** Services such as transparent caching, transcoding/transrating and parental guidance filtering fall in this first category. Without any end-to-end encryption, they are able to intercept and read/modify/write the payload of communication between the user and servers (e.g., HTTP requests and responses).
2. **Content-unaware services:** This category of services deals with the metadata related to HTTP requests and responses rather than the payload of the communication itself. Sending network conditions intelligence such as real-time throughput and congestion guidance to the client and server is a widely used example of a proxy-based service of this category. While the metadata is agnostic to the payload, the implementation of this type of service may choose to use an in-band technique such as inserting HTTP request/response headers in the communication as the vehicle to communicate the metadata to its intended endpoint. There are also discussions in progress about the creation of blind caching mechanisms that could operate in a content-unaware fashion.

If end-to-end encryption is used for all (or a large fraction) of web traffic, then the first class of services is rendered useless. The second class of services need not be, but if their vehicle becomes unavailable, then the services of that class are unable to function in their current form. At the time of writing, it is unclear if there will be a broad consensus on solutions to re-enable both classes of services for end-to-end encrypted traffic or just the second class. It is also unclear if unified solutions do exist that will solve the issue for both classes or if there will need to be multiple limited solution spaces that are addressed with specific point solutions. A good survey of potential solutions along with their pros and cons can be found in [2].

We believe that network services that can provide metadata for the optimization of content and its delivery can be very valuable, and we would like to see such services continue to develop even as more and more Internet traffic uses end-to-end encryption. Ideally, such network services will be designed to operate in a manner that neither weakens the privacy of communication nor leaks sensitive information as a side effect. We would like to encourage a collaborative approach towards that goal. A collaborative approach will allow CDNs and content providers to be aware of the features/optimizations that get applied to their traffic in the network, be able to approve/disapprove the application of these features, and work with other parties in the ecosystem to make the best use of these features.

In this paper, we first provide in Section 2 some examples of services that we consider to be of value to CDNs and content providers. In Section 3, we present some techniques that are under discussion at Akamai to make such services work with end-to-end encryption without reducing the privacy of encrypted communication or compromising sensitive user data. Solutions that entail sharing customers' SSL certificates with intermediate third parties or exposing content unencrypted from end-to-end encrypted sessions are out of scope for this paper. Section 4 concludes the paper.

2. Useful Metadata-Based Network Services

Guidance to Edge servers for media quality and delivery optimization

This is an example of the case in which a network sends metadata to content servers (CDN and otherwise) located on the public Internet so that the content and delivery can be optimized. The two main use-cases we see here are:

(a) As is well known, in mobile networks, users experience highly varying radio conditions over time which gives rise to rapidly fluctuating achievable throughput on their connections. This factor significantly hurts media QoE for mobile users. Various network elements in the radio network may have real-time bandwidth estimates for every mobile user, and they may be able to provide this bandwidth guidance to both client and server. This metadata can then be used in the media delivery optimization algorithm which includes both the choice of media bitrate (application layer), and flow control (transport layer). The MPEG-DASH SAND initiative [6] is a good example which needs an ability to exchange metadata to make network elements DASH aware.

(b) Mobile end-users typically have limited data plans. Watching long high-resolution videos can exhaust the data plan very quickly. Providing the user's data-plan info and status to media servers may allow them to adapt the quality of media delivered to suit the users' data plans. Media content providers that make media delivery sensitive to an user's data-plan status stand to gain in user engagement, as users are likely to watch more video on these sites instead of competitive sites that are data-plan unaware.

Cell-id information for media servers

This is another use-case in which a mobile network sends metadata to media servers located on the public Internet in order to optimize network resource usage and deliver a better QoE to its subscribers.

3GPP specifications on radio networks offer an on-demand operation of the multicast service (eMBMS), which is termed as MOOD [4]. This facility is meant for webservers to identify if a large user population in a mobile network is simultaneously consuming the same content and, if suitable, move them over to multicast delivery. Since multicast is highly spectrum efficient, such an opportunistic move helps not only the network conserve its radio resources, but it may also help end-users get better performance.

The media server knows the URLs of popular content being served by it, and clients' IP addresses. But, in order to initiate MOOD, the webserver needs to identify the cells where the users may be advised to switch to multicast delivery. If the network provides the users cell-identifier metadata with the requests, then the media server can perform clustering of the users, identify if any cells qualify, and invoke MOOD in those cells. The cell-id here only needs to be an opaque handle, which need not reveal any location information of the users.

TCP parameters hints for optimization proxies

This is an example of the case in which a web-server on the public Internet wants to send metadata to a proxy located inside the private network space of a mobile network.

The path from the server to the end-user may be thought of as two segments – a wired network segment from the server to the radio transmitter, and a *last-mile* radio segment. The two segments have very different performance characteristics. Segmenting the end-to-end TCP into two pieces corresponding to these network segments and using different flow-control parameters on the two segments is desirable from a performance perspective. Such TCP optimization proxies are commercially available, and may be available in the future as MEC applications, which extend some of the servers' functionality on the radio transmitter itself. Web-servers may want to provide these proxies with hints on the TCP flow-control parameters that should be used for the radio segment.

Parental guidance filtering

This is an example of an alternate implementation of a content-aware proxy service, which cannot function if the traffic is end-to-end encrypted.

Many ISPs sell parental filtering products to end-users. They provide the functionality of filtering URLs for child-safe browsing and virus/malware avoidance. Such services need no software installation on users' devices and the filtering is performed by a network service. It is implemented as an intermediate network element examining traffic flow from users subscribed to the service, identifying the URLs in HTTP requests, looking up their rating with a ratings database lookup, and then determining the appropriate filtering action based on the user's filter settings. For HTTPS traffic, the requested URLs cannot be examined by network elements in the middle. However, since the network elements can filter DNS lookups on hostname or observe the TLS Server Name Indicator, a host-granularity filter is possible for ratings based on hostname.

An alternate implementation for this service may be conceived by converting the filtering service from a content-examining service to a metadata-based service. This could allow the service to function at URL-level granularity with trustworthy collaborating websites or collaborating client software. Under this implementation, networks and web-servers collaborate on the filtering decision with metadata exchange. We present two implementation possibilities below:

1. The filtering network service sends the users preferences to webservers as metadata. The server uses its own ratings database to determine whether the content should be served to the user or not,
2. The webserver sends the rating of the content to the filtering network services as metadata, before the response body is served. The filtering decision may then be taken by the service based on the users preference and the rating received from the webserver,

3. Mechanisms to Make Metadata Work with Encryption

Various approaches for network-server metadata exchange for end-to-end encrypted traffic are in discussion at Akamai. We present three approaches below. Mechanisms have also been proposed for client devices to communicate network and device information to webservers on the Internet (for example, HTTP Client Hints [8]), however we do not focus on those in this paper.

API calls for communicating metadata

Commonly, Akamai CDN webservers that serve content to mobile users are located close to the mobile gateways. The network may host application servers in the Gi-LAN segment that support an API for 2-way metadata exchange with the Akamai's (or any other CDN's) webservers. The webservers may use these APIs for sending/receiving the required metadata out-of-band from traffic flow by providing a flow-identifier (the IP + port of the two endpoints, and protocol-id). The applications hosted by the carrier are responsible for translating the external flow-id to internal flow-id (if necessary), and implementing the necessary functionality to retrieve and process the metadata. Akamai's AME platform takes this approach to metadata exchange [4].

TLS-AUX: Injecting auxiliary data into TLS connections

The TLS-AUX proposal [5] is for an in-band technique to inject metadata directly into TLS flows that does not interfere with end-to-end encryption in any way. TLS-AUX is designed as a software layer located between the TLS layer and TCP sockets and is implemented as a user-level library. A new type of TLS record, called the AUX record, is introduced, which is allowed to exist only at the TLS-AUX layer and not higher. The AUX records are for the communication of metadata between network services and either endpoint of the TLS connection. Using TLS-AUX library calls, network services and webservers may directly insert their metadata into the TLS connections in the form of AUX records, and retrieve them. In this manner, this mechanism is an analog of sorts to the insertion/removal of request and response headers directly into HTTP flows. The related SPUD proposal is a prototype that allows for communication from network services to servers in the headers of UDP-based protocols. [7]

ICMP messages carrying metadata

In this proposal for out-of-band metadata communication, a new ICMP message type (called network-metadata type), is defined. When any network element wants to send metadata about a flow to its terminating webserver on the Internet, it generates an ICMP network-metadata message, populates it with the metadata and an identifying flow-id, and sends it to the webserver. At the webserver end, the ICMP message is received by an ICMP network-metadata listener service loaded into the OS kernel. This service retrieves the metadata, and makes it available to the webserver application through a local API. As flow-ids visible to network's internal services may differ from those visible to external webservers due to CGNs, it is anticipated that this method gains more appeal as IPv6 addressing becomes more prevalent and CGNs disappear.

4. Conclusion

In this paper, we considered the issue of some network services becoming obsolete from the rapid rise in the use of end-to-end encryption for web traffic. We expect to see research and development activity in this area resulting in proposals for new solutions that allow some network services to continue to work with encrypted traffic but without compromising on end-to-end privacy needs. In this paper, we presented examples of services that could be built on the basis of an exchange of metadata between the network and web- and media servers on the Internet. We also presented some approaches that are being researched at Akamai on how to keep these services running despite end-to-end encryption. We believe that collaborative approaches like these could go a long way towards alleviating network operators' concerns about the lack of visibility and the value of network management techniques. At the same time, the availability of data from the network, the potential for utilizing network-based services, and improved QoE could meaningfully incent content providers to work with the other parties in the ecosystem to better optimize content delivery and network costs.

References

[1] D. Naylor, A. Finamore, I. Leontiadis, et al. The Cost of the "S" in HTTPS. In Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies, CoNEXT '14, pages 133–140, New York, NY, USA, 2014. ACM.
[2] Network Management of Encrypted Traffic, position paper of the GSMA, 28 February, 2015, http://www.gsma.com/newsroom/wp-content/uploads/WWG-04-v1-0.pdf
[3] Multimedia Broadcast/Multicast Service (MBMS) improvements; MBMS operation on demand, http://www.3gpp.org/DynaReport/26849.htm
[4] Akamai AME, https://www.akamai.com/us/en/about/news/press/2015-press/akamai-unveils-emerging-mobile-business-unit-to-address-needs-of-mobile-network-operators-and-mobile-device-manufacturers.jsp
[5] TLS-AUX: injecting auxiliary data into TLS connections, https://www.akamai.com/us/en/multimedia/documents/secure/tls-aux.pdf
[6] MPEG DASH Requirements for a webpush Protocol, https://tools.ietf.org/html/draft-begen-webpush-dash-reqs-00
[7] SPUD Prototype: https://tools.ietf.org/html/draft-hildebrand-spud-prototype-03
[8] HTTP Client Hints: http://igrigorik.github.io/http-client-hints/