June 2023
Geoff Huston

RIPE 86 Bites – Encryption and Active Network Management

RIPE held a community meeting in May in Rotterdam. There were a number of presentations that sparked my interest, but rather than write my impressions in a single lengthy note, I thought I would just take a couple of topics and use a shorter, and hopefully more readable bite-sized format.

Here's the third of these bite-sized notes from the RIPE 86 meeting, on the topic of the implications of an encrypted network on active network management.

Change is hard, and the larger the system the slower the pace of change. There are just so many systems that need to change their behaviours and the motivations of users, vendors, service providers, content generators and many others all vary. Getting all of us to change some aspect of our technology, platform or application set is hard, if not impossible, to orchestrate such that it happens at the same time. We've become comfortable to an Internet which is the opposite of the Internet of the early 1990's. Change is now resisted.

Most of the time.

When change does occur over a short period of time, then it can catch us unawares. The shift to use content distribution networks (CDNs) was rapid. Not overnight, but certainly within a decade or so, when between 2005 to 2015 a significant proportion of the Internet's content shifted from dedicated servers into CDNs. These days most of us no longer use the long-haul public Internet to reach distant content and services, as CDNs bring them closer to us. Transit, as an essential component of the Internet service, is diminishing in relative value as CDNs replace transit with localised service delivery. The average *packet miles* of delivered traffic to users for a retail ISP today has reduced as significant content volume moves into the CDN and as the CDN replicate their points of presence closer to increasing populations of users. Its faster, cheaper and represents a more efficient use of network resources.

Another change which has also occurred over a short period of time is application encryption. This topic was the subject of a presentation by Cisco's Bart van de Velde at RIPE 86 (https://ripe86.ripe.net/wp-content/uploads/presentations/49-The_New_Encyrpyted_Stack_RIPE.pdf). The efforts of the browser vendors to strongly encourage the use Transport Layer Security (TLS) has combined with the adoption of freely available domain name certificates and the outsourcing of content provision of CDNs to facilitate the relatively ubiquitous use of HTTPS. Cisco's network measurement programs show a picture taken from a large-scale US carrier in 2020 where the CDN cloud services and TLS appear to dominate (Figure 1).

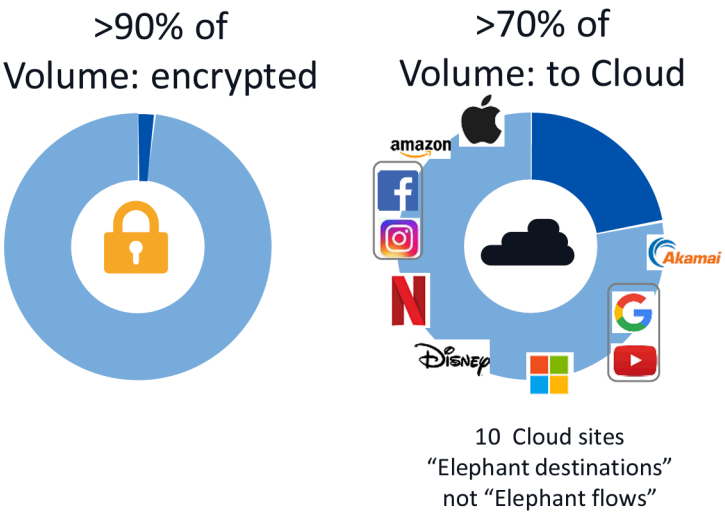

Figure 1 – Traffic Characteristics of a US ISP in 2020 - Bart van de Velde, RIPE 86 Presentation

There are some details in this figure that merit a little further mention.

While 90% of the volume of traffic is encrypted, the majority of this traffic is TLS session payload encryption layered above a TCP substrate. This implies that the outer IP wrapper and the TCP control parameters are still exposed to the network. The two communicating parties are still exposed in the IP address header, and while one end of the network transaction might be a shared CDN service point, the Server Name Indication in TLS setup is still passed over the network in the clear, so the service point being accessed by the user is still exposed. So, while the exact nature of the transaction is opaque to the network observer, the metadata of the transaction, including the service name and port, the client IP address, the time of day, session profile and duration, is still available to the network observer.

The second is that volume can be a misleading measurement. Video streams are a significant part of the volume of Internet content these days, yet there are other traffic elements that are smaller in terms of volume yet larger in terms of the assumed value of the transaction. Online payments, for example, or search are low volume, high value network transactions. There is a distinction that can be drawn between traffic volume and counts of network transactions, or *flows*. Both measurements are important in terms of understanding network behaviour.

The picture this figure paints is a situation where much of the payload volume on the Internet is now encrypted, but a network operator is still in a position to observe, and potentially "manage" the various traffic streams that they carry.

The next figure from Bart's presentation that I would like to highlight illustrates the current level of use of the QUIC transport protocol (Figure 2). By volume the use of QUIC is now at some 40% of traffic by volume, and some 35% of flows. It constitutes some 90% of YouTube traffic and 86% of Facebook traffic.

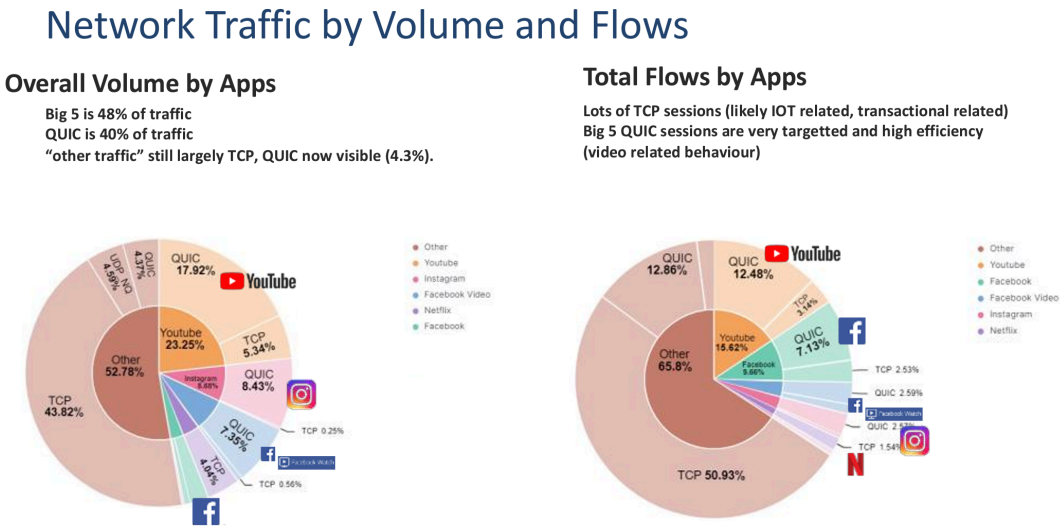

Figure 2 – Use of QUIC in Traffic Characteristics 2021 - Bart van de Velde, RIPE 86 Presentation

QUIC has a more encompassing encryption profile. All the flow control parameters, the equivalent of the TCP window values and sequence numbers, sit within the encrypted part of QUIC payloads. Even the concept of what constitutes a *flow* is opaque in QUIC, as the protocol is address agile. Unlike TCP-based TLS, QUIC already encrypts the initial packets of a connection. Although the keys of this initial encryption are visible to observers of the connection, the observer has to do the additional work of extracting the initial encryption keys and decrypting the payload information in order to eavesdrop on the information contained in the initial packets, including the server name. It is a future option to use the Encrypted Client Hello functionality of TLS 1.3 into QUIC to close off this peephole.

At this point the desires of the content distributors and those of the network operators are not aligned very well. QUIC offers content distributors greater efficiency through integration of transport session and encryption startup. The QUIC transport service offers applications a broader set of services without the issues of TCP head of line blocking, its address agile and it has a larger encryption profile, so it offers an improved privacy profile. Network operators appear to prefer better visibility of traffic to the level of applications and services and ideally would like to exert management control over the resource demands of these various traffic flows to optimise the network's response to the most valued applications. Vendors of equipment to network operators are faced with a challenge of how to instrument their equipment to achieve these various traffic management objectives while having no direct visibility into the flows being passed through the network, and no ability to directly alter the flow control parameters as could be done with TCP. About the only network-level tool available to identify traffic flows in an all-QUIC world is flow profiling, and the only way to impose shaping upon a flow is through packet drop and selective delay, both of which are relatively crude and approximate responses.

The application that delivers content and service is asserting a primacy of role in today's Internet. At this level the application has the users' interests (and their own interest of course) at heart and are not particularly motived to respond to the interests of the network operator.

The idea behind the work on Quality of Service (QoS), where the application was supposed to signal to the network its anticipated flow profile and the network would then attempt to meet the parameters of the flow, has proved to be a failure in the public Internet. I suspect that such a level of cooperation and mutual reliance between application and network is unthinkable in today's environment.

In its place we are seeing the rise of the paranoid application, which has a strong desire not to expose any information whatsoever to any other party, and certainly not to the network. This paranoia includes not even wishing to expose the application's flow control parameters to the network. From the application's perspective the network is now regarded as a hostile environment where all gratuitous data exposure is to be avoided, and all forms of attempts by the network to distort the flow characteristics are

to be resisted. The combination of QUIC, as a form of obscuring both the payload and the control parameters from the network, and BBR as a more resilient form of flow control management point to the application layer gaining the upper hand in this particular tussle.

The role of the network operator as an allocator of a scarce common resource between competing demands only made a whole lot of sense when the level of demand contention far exceeded available resources. But scarcity is no longer an abiding feature of much of the Internet's communications system. Fibre optic technologies have not only made communications capacity abundant, but the continued evolution of silicon systems has made computing and storage abundant as well. There is no enduring role for the network operator to act as a resource manager in this environment, and the continuing desires of such operators to persist in such a role seem to be somewhat anachronistic these days. In addition, applications are reinforcing this change of role of the network by using encryption to deliberately obscure the network's view of the application's role and behaviour.

The task of actively managing the mix of traffic supported by the network, and biasing network behaviour to favour certain services and content seems to be a task that is bring frustrated by the major shift of applications to obscure their behaviour behind the veil of encryption. And having gone there it's highly unlikely that applications are ever going to go back.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net