


Whitepaper Ethernet Traffic Shaping

Throughput optimization on
TCP-based WAN connections

August, 2020
Erwin Louwers, Commercial Product Manager



This document is targeted to all users of Ethernet WAN services for TCP traffic where the committed bandwidth of the hired service is less than the port speed.



1. More and more Ethernet service providers find themselves in difficult technical discussions with users of their services on the perceived throughput of TCP traffic over Layer 2 Ethernet or Layer 3 IP services. Tests show that the TCP protocol (OSI Layer 4) is very sensitive to badly aligned shapers and policers at the lower OSI layers.
2. In order to optimize the overall set of shapers and policers for maximum TCP throughput both the user and the Ethernet service provider must be aware of the principles and need for properly designed shapers, policers and buffers.
3. Correct shaping of your outgoing traffic will significantly reduce packet loss on TCP sessions.

This Whitepaper Ethernet Traffic Shaping seeks to provide a solution for users to prevent packet loss on TCP traffic using Ethernet WAN services. If no solution is provided, then the problem may continue to persist and cause inconveniences in the form of reduced throughput and packet loss.

Content

Traffic Shaping versus Traffic Policing _____	4
TCP and UDP _____	5
Traffic Policing and Shaping on Layer 2 and Layer 3 services _____	6
Conclusion _____	7
Curious about what we can do for your organization? _____	8



Traffic Shaping versus Traffic Policing

Traffic shaping is a method of limiting the bandwidth going through an interface. Typically, traffic shaping will have a committed information rate (CIR) measured in bits/second (bps). For example, if you purchase a 100Mbps circuit, the CIR is 100,000,000 bps. With traffic shaping, any traffic that exceeds the committed bit rate will be buffered. With traffic policing, any traffic that exceeds the committed bit rate will be dropped or re-marked for discard eligibility, depending on how the policer is configured.

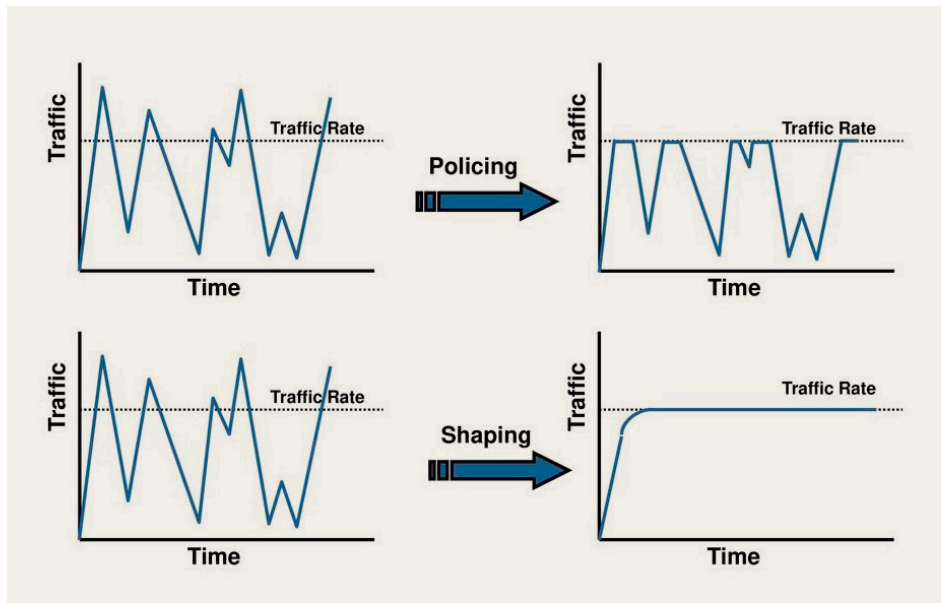


Figure 1 Policing versus shaping



TCP and UDP

The most commonly used protocols in OSI Layer 4 (Transport layer) are TCP and UDP. TCP stands for Transport Control Protocol and it provides a reliable session (connection oriented) between both communicating hosts. UDP stands for User Datagram Protocol and this is a very efficient and lightweight connectionless transport protocol.

TCP establishes a session between both hosts, including a handshake protocol in order to provide reliable packet delivery. TCP will detect packet loss during transport and will retransmit the lost packet. UDP has no such mechanism and will not retransmit any lost packets; it will not even detect it. UDP relies on the upper layer protocols to handle the occasional lost data.

UDP is mostly used for real-time applications where the retransmission of lost data has no use because of the real-time character of the application. TCP is used for applications where assured packet delivery is important. TCP-based applications show a complete different behavior if policers start dropping frames/packets.

Once a TCP session is established between the two communicating hosts the slow-start algorithm will start sending data at a low bitrate. Once the receipt of the transmitted packets have been acknowledged by the remote host the transmitting host will send the next batch of packets. The amount of data in transit is called the 'Window size', basically this is the amount of data (bytes) that can be in transit without being acknowledged. The receiver will slowly increase this window (sliding window algorithm) to allow the transmitter to increase the bitrate.

This works fine until the rate of acknowledgements being received by the transmitter drops. The transmitter will detect that packets have been lost during transmission and must retransmit the lost packets. This means that the effective bitrate is decreased in order to deal with the packet loss. This results in the notorious sawtooth throughput diagram depicted below.

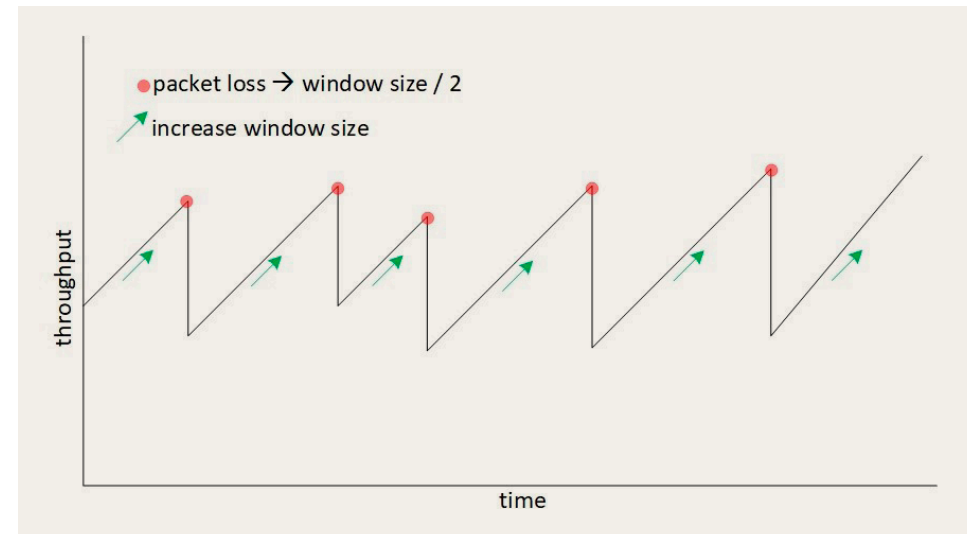


Figure 2 Effective throughput of TCP traffic without proper shaping on layer 2

As can be seen in figure 2 the effective throughput (the surface under the sawtooth) is seriously affected by the packet loss and consequent reduced sending rate. This is what most users complain about when they experience low bitrates and badly performing applications.

Traffic Policing and Shaping on Layer 2 and Layer 3 services

There are several technical solutions on limiting the bandwidth of Layer-2 and Layer-3 services. But if these measures do not match the TCP settings of Layer-4 the user of the Ethernet service might experience poor throughput. The question is: who is responsible for Layer-4?

If a service provider delivers a Layer-2 Ethernet EVC or a Layer-3 IP service which is policed at the committed bandwidth at that layer, can that provider also be fully responsible for the TCP layer?

Or, is it the responsibility of the user who orders a Layer-2 or Layer-3 service to make sure the Layer-4 TCP performs properly?

In practice it is the responsibility of both the user and the service provider to make sure the TCP layer performs properly; it is a shared responsibility, as depicted below.

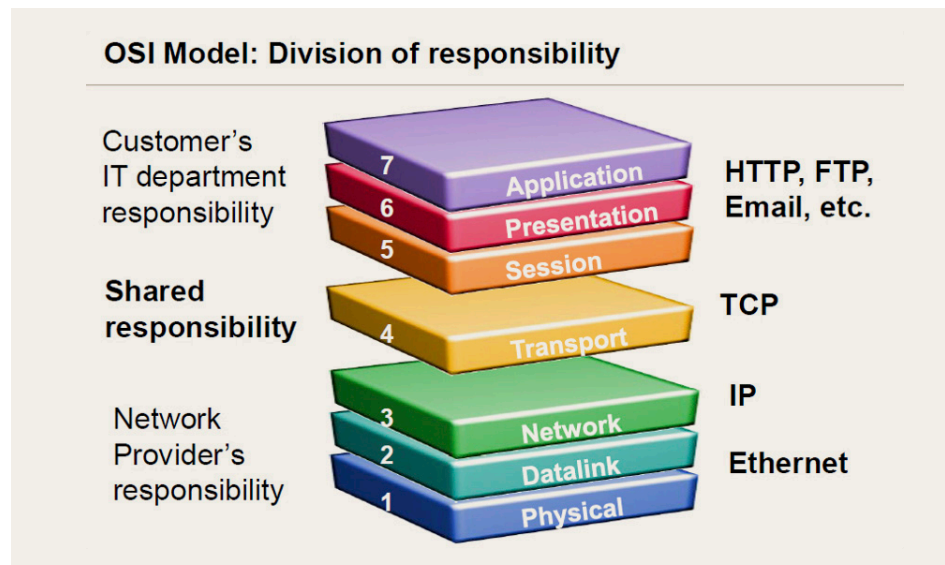


Figure 3 Shared Responsibility Transport Layer

By presenting this whitepaper to its users of Ethernet WAN services, Eurofiber believes that the user should be able to reduce throughput issues caused by configuration settings of their outgoing traffic.

Eurofiber uses ingress policing for their Ethernet services as depicted below. As mentioned earlier, policing will drop frames when the bitrate exceeds the CIR (ordered EVC bandwidth).

It is therefore important that the outgoing traffic (on UNI and NNI) is shaped by the user to avoid packet loss.

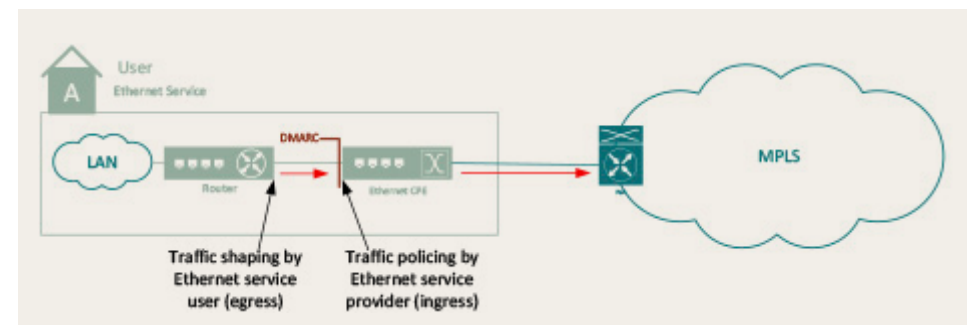


Figure 4 Traffic shaping and policing on an Ethernet WAN service with a service provider CPE

Traffic shaping is not needed when the CIR (EVC bandwidth) is equal to the port speed (i.e. 1 Gb/s EVC on a 1 Gb/s port) hence, the traffic is automatically limited by the port speed.

Conclusion

When using Ethernet WAN services for TCP communications it is important to shape the outgoing traffic to reduce packet loss and to optimize throughput. Shaping will perhaps increase the latency, but it will ensure a reliable and stable connection.



Curious about what we can do for your organization?

Eurofiber has been operating a high-quality digital open infrastructure since 2000. With our own fibre-optic network and datacenters we offer companies, governments and non-profit organizations a future-proof, smart and open infrastructure. Customers have the freedom to choose the services, applications and providers they need. In this way, they can fully exploit the innovation potential in digitalization. In addition to the extensive fiber optic network in the Netherlands, Belgium and France we also offer solutions for interconnectivity between almost all high quality carrier neutral datacenters in the Benelux. Eurofiber is thus laying the foundation for the digital society. The Dutch government has given Eurofiber the status of 'vital infrastructure'.

Go to www.eurofiber.nl



Eurofiber. Lifeline for the digital society

Safariweg 25-31, 3605 MA Maarssen
Postbus 7072, 3502 KB Utrecht
+31(0)30 242 8700, info@eurofiber.nl
www.eurofiber.nl

