



ETHERNET ACCESS TO PRIVATE IP USER GUIDE

Purpose and Objectives

This document demonstrates the maximum theoretical throughput when accounting for Ethernet overhead

- Fast Ethernet, Gigabit Ethernet, and even 10baseT have a certain amount of Ethernet overhead comprised of:
 - Interframe Gap (IFG), which is 12 bytes of equivalent-length silence
 - Preamble, defined as 7 bytes
 - Start Frame Delimiter (SFD), defined as 1 byte


Other Ethernet fields are defined here as packet overhead, e.g., MAC address, Length/Type, etc.

- Ethernet access speed to Verizon Private IP (our Layer 3 MPLS VPN service) includes, for any stated bit rate, the IFG, preamble, and SFD, so actual throughput calculations must be discounted accordingly
 - The overhead calculation is valid for any access speed or equipment
 - Overhead must be calculated for a specific Ethernet packet size



Ethernet Frame (untagged)

IEEE 802.3



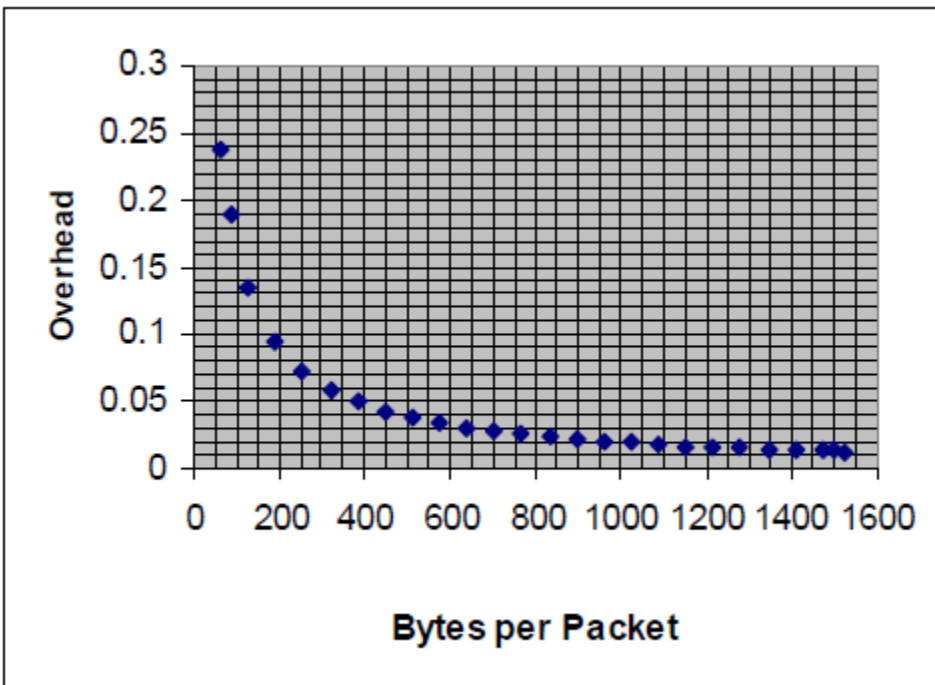
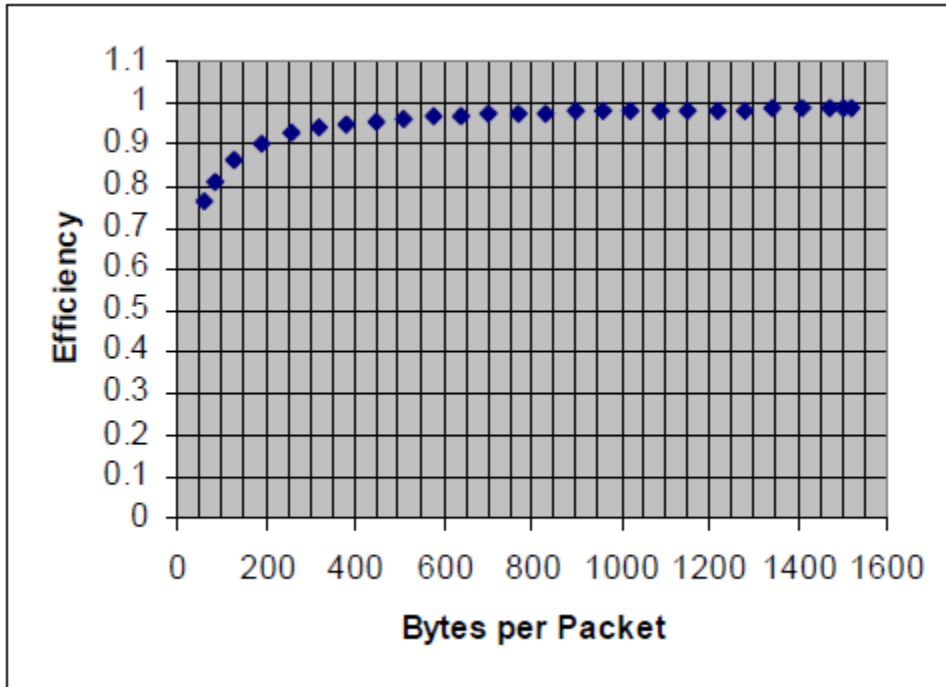
- For the purposes of our discussion, the IFG, preamble, and SFD are considered Ethernet overhead (20 bytes).
- The MAC addresses, length, and FCS comprise packet overhead (18 bytes), and if VLAN tagged, 22 bytes.

Overhead Calculation

The Ethernet overhead is calculated as follows:

- Choose the packet size in bytes
Packet size is the datagram from TCP, UDP, etc. plus the 18 or 22 bytes of packet overhead.
Example: 46 bytes of TCP + 18 bytes = 64 bytes.
- Compute these equations:
Efficiency = $\text{packet size} / (\text{packet size} + 20 \text{ bytes})$
Overhead = $1 - \text{Efficiency}$
- Example: Efficiency @64 byte packets = $64 / (64 + 20) = 0.76$
Overhead = $1 - 0.76 = 0.24$, or 24%
- Example: Efficiency @1500 byte packets = $1500 / (1500 + 20) = 0.987$
Overhead = $1 - 0.987 = 0.013$, or 1.3%
- When sending many small packets on a heavily utilized access circuit, the throughput will be decreased per these equations and can be quite noticeable

Efficiency and Ethernet Overhead Plotted



Impact on Maximum Throughput

- Given that Ethernet overhead grows as packets grow smaller, any maximum throughput calculations must be discounted by the efficiency, especially when small packets are used
- Example: Customer has “100 Mbps” access to PIP (the nominal bit rate). That is the maximum theoretical throughput when using 512 byte packets?
- Efficiency @512 byte packets = 96.2%, so the best the customer could ever achieve is $100 \text{ Mbps} \times 96.2\% \text{ efficiency} = 96.2 \text{ Mbps}$
- Many applications do not use a fixed packet size, so actual throughput can be difficult to estimate with the equations. An average packet size may be more useful.

How to Deal with Ethernet Overhead

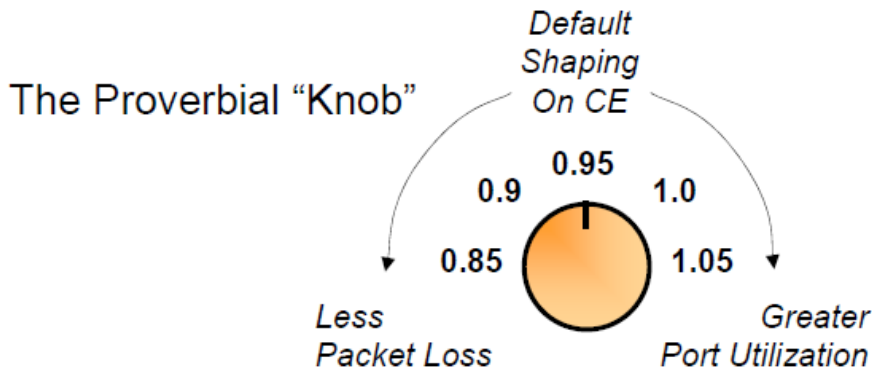
- Customers need to expect Ethernet overhead and plan accordingly
- The key is to shape the output of the CE router before sending the traffic stream to PIP. That is, shape in the CE in the CE>PE direction.
 - Aggressive shaping, i.e., close to or greater than the access rate, can increase throughput at the expense of packet loss
 - Conservative shaping, i.e., less than the access rate, can reduce packet loss at the expense of throughput
 - Only application-level testing can determine the acceptable level of packet loss. Can the application efficiently adapt to a circuit with packet loss? TCP handles this well, but other protocols may not.
 - Please be aware that “goodput”, or usable throughput, may decrease drastically depending upon the application’s tolerance of packet loss

Tuning CE Shaping

Verizon Business recommends a conservative shaping policy based upon an estimated average packet size of 384 bytes, but the customer is encouraged to review the packet size based upon actual traffic. So, set the output shaping initially in the CE (CE PE direction) to:

Shaping speed = 0.95 x Ethernet Access Speed

This setting will work for the majority of PIP customers, but it is a compromise between packet loss and port utilization. If the shaper shapes to 95% of the access speed, and if the average packet size assumption is correct, then packet loss should be tolerable and goodput should be *****HIGH*****. On the other hand, if packet loss is too high, shape at a lower speed. If the port utilization is consistently too low, increase the shaping rate.





Risk to QoS with Aggressive Shaping

QoS mechanisms such as CBWFQ and LLQ on Cisco routers do not engage to prioritize traffic if the egress queues on the CE are not congested

- The risk is that light loads, less than the shaping rate, MAY NOT exhibit the QoS behavior expected because QoS mechanisms do not engage
- During “unloaded” conditions, all traffic is FIFO. This is fine unless the shaping is so aggressive that the router fails to perceive the congestion that is actually occurring.
 - This might happen if the shaping rate is 110% of the port speed. The router will fail to sense that the port is congested at 100%.
- Customers have the ability to tune their own CE shaping speed if this phenomenon is observed
- If greater prioritization of traffic is desired, tune the shaping DOWN a few percentage points

Policing and Shaping on the PIP PE

- Ingress policing takes two primary forms in the PIP PE with regard to Ethernet access
 - When enforcing Expedited Forwarding (Gold) CAR, where the excess traffic is dropped
 - When using sub-rate Ethernet access, which may be applicable when a customer buys an access rate less than the physical port speed. For example, a 10 Mbps access is purchased on a Fast Ethernet port. The excess is dropped by the policer.
- Egress shaping is used to match the PE’s output to the access speed (PE - CE). In some cases, the PE’s shaping rate may have to be adjusted per the same rules used on the CE. For example, if packet loss is occurring because of a preponderance of small packets, the shaping rate may have to be adjusted downward.

Conclusion

- The maximum throughput of Ethernet must take into account the Ethernet overhead. Just because the access rate is nominally rated at “X” Mbps does not mean that “X” Mbps will pass through the network. Instead, multiply “X” by the efficiency for a given packet size.
- Problem: The method of shaping to Layer 1 speeds on Ethernet interfaces is causing non-QoS aware random drops (packet loss) between CE and PE during congestion conditions
- Solution: Adjust the PE and CE egress shaping on Ethernet interfaces depending on the type of traffic to compensate for Ethernet and GFP overhead.
 - Different encapsulation overheads between the CE and PE equipment
 - » Ethernet overhead: 20 bytes (IFG=12, preamble=7, and SFD=1)
 - » GFP overhead: 12 bytes (PLI=2, cHEC=2, Type=2, tHEC=2, pFCS=4)



- The average frame sizes defined for three types of customers are:
 - » 78 bytes (Voice centric)
 - » 140 bytes (Voice/Data combined)
 - » 404 bytes (Data centric)

– The following Ethernet shaping percentages will be applied on Layer 1 (L1) utilization to calculate the Layer 2 (L2)

Bandwidth:

- » 76% (Voice centric)
 - » 85% (Voice/Data combined)
 - » 94% (Data centric)
- Expected Results: During congestion, CE and PE routers will drop packets based on PIP QoS scheme and not on random drops.



Service Assurance User Guides Library

Documents can be found on the [Service Assurance User Guides](#) page.
The latest version of this document can be always found [here](#).

General Customer Training Information

Go to our [Customer Training Portal](#)* to enroll in training or to download other user and reference guides.
*Registration is required

Verizon Enterprise Center

The [Verizon Enterprise Center](#) portal is an easily accessible tool that supports you in dealing with Repair related technical issues via repair tickets, as well as with Invoice inquiries and Account Management requests, offering an alternative to emails and phone calls.

Getting started on Verizon Enterprise Center

Introduction to Verizon Enterprise Center and information on how to register can be found on the Guides & Tutorials page [here](#).



© 2023 Verizon. All Rights Reserved. The Verizon name and logo and all other names, logos, and slogans identifying Verizon's products and services are trademarks and service marks or registered trademarks and service marks of Verizon Trademark Services LLC or its affiliates in the United States and/or other countries. Microsoft and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks and service marks are the property of their respective owners.