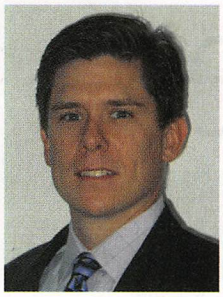

By Hunter Newby

(Internet Protocol + Ethernet Transport) x QoS - (The Public Internet / Net Neutrality) = Enterprise Peering

Although mathematical equations may seem to be complex, they can be better understood if their fundamentals are isolated. Enterprise Peering is an evolution and inevitable given the lessons of the past. All that is possible will be, particularly if time is combined with the formula above, because with time comes the repetition of history. In order to better understand the evolutionary network path we're all on, let's break down the equation.

Internet Protocol — not the Internet, but rather the protocol itself — is widely used and acknowledged as THE common language for machines and devices of all types to intercommunicate. Enterprise networks have been using IP successfully for many years. It has been so successful that enterprise network managers actually connected their remote offices using IP to support applications of all sorts, including e-mail and even Voice over Internet Protocol (VoIP). In the past, the transport links that were used for the wide area connections were clear channel TDM, ATM and frame relay, but more about that in a moment.

Ethernet transport has been on the networking scene in a meaningful way for about the past three years. Today, it is widely available in the metro, long haul, and even international long haul network segments. The legacy transport protocols of ATM and frame relay have given way to the King of the Enterprise Road, Ethernet. This is due in large part to the fact that Ethernet has successfully been in the local-area network (LAN) for 15+ years.

In the early days, frame relay succeeded X.25 due to its increased capacity and, since it was digital and packet-based, it could carry the data protocols of the LAN with relative ease. Frame relay was also squarely focused on layers 1 and 2 and did not route packets as X.25 did. This was a clear advantage for frame relay, as it lowered the overhead and increased performance dramatically. Ethernet focuses in layers 1 and 2 as well, but has a key element that frame relay lacks, and that is the ability to create Virtual Local Area Networks — VLANs.

Within the last few years, transport carriers have finally started to come around to the possibilities of Ethernet.

Within the last few years, transport carriers have finally started to come around to the possibilities of Ethernet. These services have been met with great acclaim from the IT and MIS directors looking to expand their Ethernet capability to the wide area and move away from legacy frame relay and ATM. Today, almost every transport carrier has an Ethernet service offering.

QoS is certainly one of the top reasons why have enterprises been building private (frame relay) networks for over 15 years. It is not just Quality of Service though, but also control. When the IT department established a frame relay connection, THEY set the parameters of throughput and availability, or in other words, the quality of the circuit. Those metrics are Committed Information Rate (CIR) and Extended Information Rate (EIR).

The CIR was the minimum guaranteed throughput from the carrier that the enterprise could count on. With that there was a Service Level Agreement and certain penalties for non-performance. The EIR was the size of the full port that the user could occupy if no one else was using that capacity across the carriers' network at the time. The delta was the "shared" portion of the network. There are many possible combinations of CIR and EIR, but one example from a bits perspective it looks like this (768k EIR - 64k CIR), or for a half a T1 (768 kilobits) there was one channel (64 kilobits) guaranteed. The difference is what came to be known as a "best efforts" class of service.

No single carrier network goes everywhere, but enterprises have offices in many places. So, how did a carrier sell a complete circuit to a customer? In the carrier world, that is known as a Type II, or off-net circuit. Type I is on-net for the carrier the whole way. Type II encompassed all types of circuits, but frame relay has its own special type of interconnection within that distinction. It is known as the NNI, or Network to Network Interface. This is how one carrier running frame relay over its network could directly interface its network with that of another carrier frame relay network. The NNI is based on

standards, it happens in a real physical place (usually a carrier hotel like 60 Hudson St.) and once it was complete it was seamless.

Certain carriers, such as WorldCom (UUNet), had very large frame relay networks with several large NNI's to other carriers such as Intermedia (Digex) and GTE/Genuity (BBN). These carriers' NNI's enabled their enterprise sales groups to sell end to end frame relay circuits with guaranteed QoS to the buyer since the carriers had established the NNI's at layer 2 all the way through their networks. This is exactly what is happening now with carriers establishing Ethernet connections end to end. I call them ENI's — Ethernet Network Interfaces; maybe that is a standard already, but I have not seen the term used.

The Internet, as we know, is a huge, shared network, based on routers running Internet Protocol that can be accessed in a variety of ways, frame relay being one of them. Does anyone remember the Burstable T1 Internet circuit? If not, it was the hottest selling "data" circuit of the late 90's. The Burstable circuit got its name from the delta between the CIR and EIR, the area in which the user could "burst" their traffic in a "best efforts" environment, of the transport carriers that supported it. (BBN's entire IP backbone was actually based on WorldCom frame relay circuits before it was sold to GTE.) That is the hierarchy of things.

Needless to say, the Internet created a new, non-enterprise class of buyer for the carriers, the Internet Service Provider, or ISP. They became the middlemen and facilitators for the enterprise to "get to" the Internet. The fact is that the Internet as a public network is decidedly different from frame relay, ATM, and other private, packet-based enterprise data networks carrying IP. The enterprise-managed, private Ethernet network phase we are all now entering is akin to enterprise private frame relay networks and is not new, unproven, or risky at all. It is actually a reversion back to what has already made sense.

The Net Neutrality debate actually throws fuel on the fire of the Internet versus private IP networks. The concept of a virtual private network, or VPN, has been around for a while. In the frame relay world it is what the carriers called enterprise networks that used "their" frame relay networks. Since the enterprise did not build the network, but rather leased ports, it was not really the user's network, but rather the carrier's. Also, it was not dedicated, as in point to point DS1/3 level circuits, but part of the carrier's frame relay "cloud." It is here, in this mentality and also that of POTS line, DSL, and even cable modem access to the Internet, where the carriers believe that ASPs, such as Google, are riding on THEIR pipes for FREE, the pipes that access the "public cloud," or Internet, that is, and not the cloud itself.

If your Access Provider (transport) is also your Internet Service Provider, then they have the ability to make this claim, but if you lease a data circuit (i.e., Ethernet), and connect that directly to another network (ENI), whether it be another carriers' transport network, or directly to an ASP or another enterprise, there is no ISP router in the middle and therefore no ability to apply a higher or lower standard of quality to the packets. The basis of QoS and the root of the Net Neutrality debate is the ability, or possibility of the access provider to dictate to the buyer what they can effectively use.

The truth is that the Access Provider-ISPs, like the RBOCs, cannot threaten to apply a higher or lower quality standard to any packets if they themselves lack the core NNI/ENIs and SLAs between each other. This is why the large last-mile access carriers are beginning to establish those very types of connections and agreements at the Ethernet level. No network goes everywhere all of the time and therefore no carrier can guarantee anything in a dynamic off-net scenario, especially over the "best efforts" Internet.

This equation, (Internet Protocol + Ethernet transport) x QoS - (The Public Internet / Net Neutrality) = Enterprise Peering, simply means that there is a better way for enterprise network operators to interconnect between their own sites, other enterprises, service providers, and ASP's directly, thus avoiding the Internet, but still using IP. This improves quality and reliability. In addition, this formula can equate to significant savings when avoidance of the PSTN and its associated costs are factored in, changing the equation to:

(VoIP + Ethernet transport) x QoS – (PSTN) + (The Public Internet / Net Neutrality) = Enterprise VoIP Peering. IT

Hunter Newby is chief strategy officer for telx. For more information, please visit the company online at www.telx.com.