Understanding Bandwidth Throughput Webinar
Frequently Asked Questions

General Administration

Is the Webinar with audio available for download? Is it possible to obtain a copy of the slides from this presentation?

The recorded webinar can be viewed:
http://www.metroethernetforum.org/PPT_Documents/OAM/flash/movie/Understanding-Carrier-Ethernet-Throughput-FINAL.wmv

Slides are available:
http://metroethernetforum.org/PDF_Documents/OAM/Understanding-Carrier-Ethernet-Throughput-Webinar-FINAL.pdf

For more information, please visit:
http://www.metroethernetforum.org/oamwg

Carrier Ethernet

Does CIR (Committed Information Rate) include the bandwidth used for preamble and inter-frame gap?

No, CIR is defined as the average bit rate of Ethernet service frames. Ethernet service frames do not include the preamble, start of frame delimiter, or inter-frame gap.

Is there a maximum number of customer VLANs that can be mapped to the single EVC?

Clearly, the 12-bit VLAN ID limits the number of VLANs to 4096. Otherwise, there is no standard limiting the number of customer VLANs per EVC. Note however, there may a practical limit depending on what your networking equipment is capable of supporting.

Policing is done as rate, i.e. bits/time. What time period is used and is it standard across all router/switch vendors?

The simple answer is that the time period used is mostly controlled by the Committed Burst Size (CBS) and the Line Rate. For example, for a service with a 50 kB CBS and 50 Mbps Committed Information Rate (CIR) operating at a 1 Gbps Line Rate (User network Interface) UNI, the maximum time period for a line rate burst that a policer will allow is about 0.4 to 0.5 mS. Read on for a more detailed answer.

Current network standards call for policing to be done to confirm that the offered traffic does not exceed the combination of the CIR (Committed Information Rate) which is expressed in bits per second, and the CBS (Committed Burst Size) which is expressed in bytes. When the service is turned on, the bandwidth profiler used to police the traffic sets up a token bucket, with as many tokens in it as there are bytes in the CBS. Each Service Frame that comes in to the network first must go through the bandwidth profiler where it is checked against the available tokens in the bucket. If enough tokens are available, the frame is forwarded and the bucket is emptied by that many bytes worth of tokens. When the next service frame comes in, the bucket is first refilled with as many tokens as are allowed by the amount of time that has passed and the subscribed Committed Information Rate, with a limitation that the bucket cannot be refilled to a level greater than the CBS. Then the Service Frame is checked again against the available tokens, and if there are enough in the bucket, the Service Frame is forwarded. If there aren't enough tokens, the Service Frame will likely be discarded by the policer. Not all carriers may offer fully-configurable committed burst sizes, and some implementations may be equipment-specific or carrier-specific.

For "Yellow" traffic, are the original DSCP / 802.1p markings kept or are they marked down (for example, mark AF31 traffic down to BE traffic)?

Note that IEEE 802.1p has been replaced by IEEE 802.1Q, which specifies VLAN tagging. On the one hand, a Carrier Ethernet service specifies whether the VLAN priority is to be preserved. When priority is to be preserved, 802.1Q VLAN priority bits will not be changed. In either case, DSCP also will not be changed. On the other hand, customer traffic can be encapsulated in a service provider tagged (S-Tagged) VLAN frame. While the class of service of an S-Tagged frame is determined initially by customer tagged (C-Tagged) VLAN frame, the priority of an S-Tagged frame can be changed. Depending on the service provider, a yellow frame can have its 802.1Q bits remarked from C-Tag to S-Tag, have the "drop eligible indicator" bit set, or pass the frame through with S-Tag class of service bits unchanged from C-Tag. The Service Provider is likely to mark the S-Tag differently than the customer's tag when the customer's drop eligibility indication does not match the green/yellow determination made by the Service Provider's Bandwidth Profiler.

Does a packet marked green by the first operator network also get marked green in the second operator network?

In a properly designed E-Line service, yes. Carriers are standardizing their network interfaces and service configurations to be able to reliably provide this sort of functionality across networks. However, the MEF External Network to Network Interface Technical Specification that carriers use to connect their networks was just ratified in January of this year, and there may be different implementations of this feature around the world in the hardware already deployed in the network, and those differing implementations may bear on this answer. You should direct this question to the potential team of providers that will cooperate to connect your particular access locations.

What are your thoughts on Jumbo Frame and the ability to pass those frames through multiple operator networks?

Jumbo frames or Maximum Transmission Unit (MTU) size in excess of 1518 are ideally used in gigabit Ethernet or higher-speed networks. The reason is that large frames in low-speed networks result in more frame delay (latency) and more frame delay variation (jitter) which can cause a problem when a mixture of high and low priority traffic share a connection. Support for jumbo frames, is a parameter that is agreed upon by network operators working together to deliver an Ethernet service, and in practice, jumbo frame support is common for higher bit rate services. There is a potential throughput increase available to some applications such as TCP if jumbo frames are used in the wide area network on high speed backbones.

Where can I find the specifications about Ethernet Ring Topology in MEF documentation?

MEF has not defined Ethernet Ring topology. Ethernet Ring topology is defined in ITU-T G.8032.

Factors Affecting Throughput

Is Windows scaling an application dependent or OS dependent?

TCP windows scaling can be both application and OS dependent. The OS establishes a TCP Window size that will be used by default. However, an application may also set the TCP Window size each time it initiates a TCP connection.

Exactly where can a customer adjust their TCP window size – i.e., on their router, server or individual end-user work stations?

TCP window size can be set on individual workstations or on a server. TCP window size is set at the end-points of a TCP session. TCP window size is therefore generally not set on router. TCP Window size can be tuned by a system administrator, but it is not usually set by individuals. Users wishing to employ TCP Window scaling in their Microsoft Windows based computers can find and download information found on the internet which gives instructions for editing the system registry to increase the TCP Window size. This question also suggests a broader issue of performance tuning. TCP window size is only one factor

which can be used to tune network performance. Network performance may be improved more effectively by looking at other factors such as traffic shaping and reducing end-to-end latency.

Proving Service Performance

What is the role of Y.1731 in the process of measuring an EVC?

Y.1731 provides tools for the in-service monitoring of Ethernet Virtual Circuits (EVCs), and performing corresponding Fault Management functions when the performance monitoring turns up a problem. Y.1731 specifies Continuity Check Messages which can be sent every second or faster to measure Frame Delay, Frame Delay Variation, and Frame Loss while the service is in-service. The Link Trace message traces the path that an EVC takes for fault management purposes. The loopback and loopback response messages allow further on-demand testing to take place in response to reported problems. Y.1731 loopback or test frames can potentially be used to measure throughput, however the standard is not specific on how to do so, and not all network equipment may be able to respond to Y.1731 loopback frames at full CIR or line rate.

Are RFC 2544 and/or Y.156sam intrusive tests?

Yes, both RFC 2544 and Y.156sam are designed to be run in the absence of "normal" customer Ethernet service traffic.

If I understand this correctly, Y.156sam is not a replacement to RFC2544 as we still need to test the limit of the Network. Can you please repeat the difference between RFC2544 and Y.156sam?

RFC 2544 provides a methodology for measuring throughput, latency, frame loss and back-to-back (burstability) for the purpose of verifying the design of network equipment. It can also be extended to the application of verifying the acceptable transmission of traffic over the network at the time of service activation. However, RFC 2544 has some well-known shortcomings. RFC 2544 is not Ethernet service aware, thus tests are run as a single flow at a time. Also, RFC 2544 tests are performed sequentially, so each parameter is measured in a different timeframe. RFC2544 doesn't test latency to the depth that Y.156SAM does. Finally, important Ethernet Service attributes such as Frame Delay Variation are not part of the methodology. Y.156SAM in comparison, allows the testing of all relevant parameters at the same time for faster performance of the activation test, and specifically tests to the CIR to verify SLA performance, and tests to the EIR limit and just beyond to verify policing behavior. Y.156SAM is expected to be completed by the ITU in the early 2011 time frame.

Will RFC 2544 be updated to include frame sizes greater than the standard 1518?

It is unlikely that RFC 2544 will be updated. However, most vendors who have implemented RFC 2544, provide their own extensions to address shortcomings of the standard. Many vendors providing RFC 2544 testing test frames sizes larger 1518.

Is there a set time for testing for each frame size in the rfc2544? Or maybe is there a standard time for each test in the RFC2544?

RFC 2544 recommends at least a 60 second test time for each frame size in the final determination stage. During the binary search algorithm, shorter times are allowed. In practice users often adjust the final determination time to be longer or shorter to provide faster or slower overall test times in line with their test objectives.

Does frame loss affect the integrity of RFC 2544 results?

One of the purposes of performing an RFC 2544 test is to determine whether frame losses occur. RFC2544 makes the simplifying assumption that no frame losses occur until the traffic gets so high that congestion occurs and many frames start to be lost. A person strictly performing an RFC 2544 procedure may have difficulty determining where the frame loss occurs from congestion and where it occurs from background network performance issues, and this may reduce the integrity of the RFC2544 test.

Does Y.156sam measure multicast or broadcast traffic or is it "-cast" agnostic?

Y.156sam is measuring frame rates. Y.156sam does not specifically require the use of broadcast or multicast frames.

What is the difference between RTT and RTD?

Round trip time (RTT) and Round Trip Delay (RTD) can be used interchangeably. Generally, the terms refer to the time it takes from when a data packet is sent to the time the time it the acknowledgement of that data to be received.