Y.1564 Ethernet Service Activation Testing Methodology

By Paul Marshall | February, 2011

SUNRISE TELECOM®
Productivity Rising™

Overview

Y.1564, Ethernet Service Activation Test Methodology, is arriving from the ITU just in time to ensure that service providers can do a proper job of verifying the correct configuration and performance of Carrier Ethernet services at the time of service activation. Sunrise Telecom played a special role in this document, driving it to include the second generation test capabilities that expanded the bandwidth profile coverage. At the time of Y.1564 consent, only Sunrise Telecom's IntelliSAM™ could demonstrate this second-generation capability. This white paper will help anyone involved in the process better understand:

- How Carrier Ethernet works
- What each bandwidth profile parameter is
- How Y.1564 improves the reliability of the service activation process
- Why the older RFC 2544 standard needed to be replaced for service activation applications
- How recommendation Y.1564 verifies that parameters are set correctly
- How the recommendation verifies that all services work together simultaneously
- How and when to apply the procedures in appendix 1 to verify the CBS and EBS configuration.

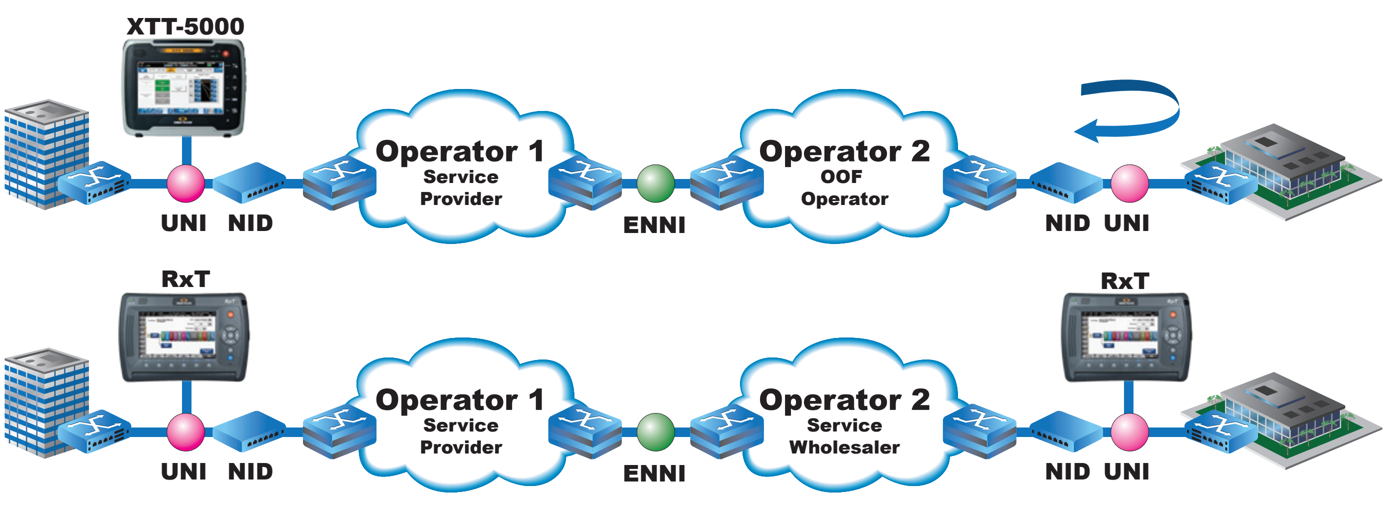

Figure 1 – Carrier Ethernet Service Activation Testing

IntelliSAM™

Carrier Ethernet Service

Carrier Ethernet Service defined by MEF, the Metro Ethernet Forum, carries a customer's Ethernet traffic across the Service Provider's network as described in the Service Types shown in Figure 2 below:

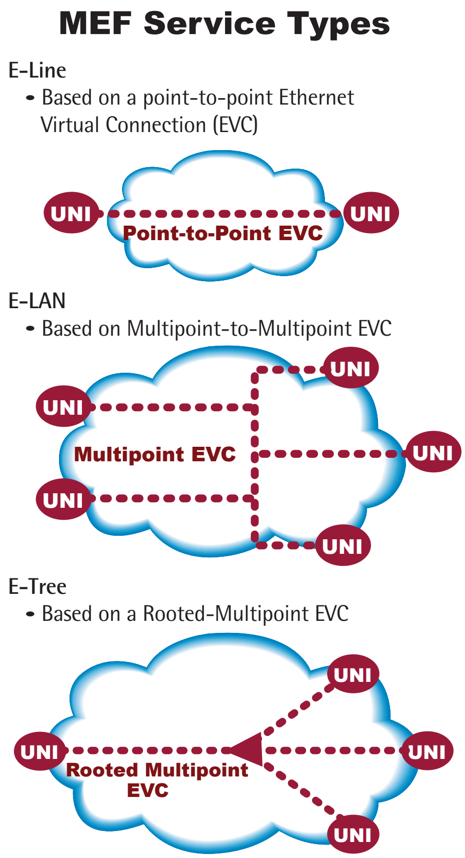

Figure 2 – The standard Carrier Ethernet service types

The E-line is the simplest configuration, forwarding all the frames between two locations. The E-LAN connects multiple locations together, all with equal connectivity as if they were on a standard LAN.

The E-Tree is a model that a service provider might use to distribute Ethernet service over a DOCSIS or GPON network, where all the subscribers (the leaves) need to be kept separate from each other for security purposes. They are only allowed to communicate with the service provider (the root).

Bandwidth Profile

With Carrier Ethernet, customers may transmit traffic into the network according to a set Bandwidth Profile. The profile includes the average rate at which they can transmit traffic, and specifies what kind of burst capacity is available to them. Bursting is important in Carrier Ethernet, because the applications using the service often have widely varying transmission characteristics, from voice and video streaming which are fairly rate-constant, to file transfer, which can burst at rates far outstripping the subscribed average service rate. If the customer transmits traffic into the network according to that profile, they are guaranteed a certain service performance level. The Bandwidth Profile consists of the following parameters:

CIR:

The CIR (Committed Information Rate) is the average maximum information rate that the customer is allowed to transmit traffic into the network at agreed service performance objectives. If the customer transmits faster than this for longer than allowed by the CBS, then their traffic might get dropped or suffer from longer delays or higher delay variation.

CBS:

The CBS (Committed Burst Size) defines how much freedom the customer is given to transmit his committed traffic in bursts at temporary rates that are above the CIR. Ethernet networks often have bursty traffic, and this is a major source of negotiation between customers and service providers. The customers want infinite burstiness so they are assured of not losing their frames, but the service providers don't want to buy a lot of excess transmission and switching equipment to safely handle all those bursts. The solution is to sell the burst capacity to customers by the byte, and the customer can pay for as much burst capability as needed, and the network provider will invest in the capacity needed to handle the burst performance that has been sold.

Traffic that conforms to the CIR/CBS rate and burst combination is said to be "green" traffic and is marked as such to be sure that it meets its performance objectives.

EIR:

The EIR (Excess Information Rate) is the average maximum information rate of additional traffic that the customer is allowed to transmit into the network over and above the CIR/CBS allowance. It is sold at a cheaper rate than the CIR, and there are no set performance objectives. Successful delivery of the EIR/EBS frames is subject to network congestion. Capacity is not reserved within the network to forward those frames.

EBS:

The EBS (Excess Burst Size) defines how much freedom the customer is given to transmit his excess traffic in bursts at temporary rates above the EIR plus CIR. If the customer sends traffic that exceeds both the burst limits provided by the CIR/CBS and the EIR/EBS, then the traffic is deemed "red" and may be dropped immediately upon entry to the service provider's network.

CM:

CM (Color Mode) determines whether the customer is allowed to tell the network which of the frames should be marked as green frames and which frames should be marked as yellow frames. "Color Aware" is used to describe the mode where the customer is marking each frame as green or yellow, and the network is supposed to take this marking into account at the bandwidth profiler and traffic policer. "Color Blind" is used to describe the mode where the bandwidth profiler - traffic policer ignores color markings on the customer's frames and instead independently profiles and marks the color of each arriving frame according to the purchased bandwidth profile and the frame arrival history.

Being able to mark the color can be important to the customer if he has different applications using the same service which have different sensitivity to performance problems. Voice traffic, for instance, needs low frame loss and low delay, and the customer usually wants to classify this traffic as "green". File transfer traffic utilizing a transport protocol such as TCP may be relatively insensitive to transmission problems, and may receive satisfactory service when classified as "yellow".

If the Color Mode is set to Color Aware, the customer's traffic must still conform to the CIR/CBS profile and the EIR/EBS profile to be classified as green or yellow; otherwise the network will change the customer's marking. If some of the customer's green-marked traffic doesn't conform to the CIR/CBS burst limit, it will be remarked as yellow by the network. Some of the customer's yellow-marked traffic and non-conforming green-marked traffic will be dropped at the user network interface by the network if it doesn't conform to the EIR/EBS burst limit.

CF:

For customers who have very bursty traffic, it may be possible to couple together unused green traffic capacity into the yellow traffic by setting the CF (Coupling Flag), which allows for higher throughput than would otherwise be achievable. Coupling is particularly applicable in Color Aware cases where the customer often sends the yellow class of traffic when not fully utilizing the green class. In this case, he's not using the green service he paid for, so unused tokens that would have overflowed an already-full green token bucket can be diverted to refill the yellow token bucket to allow a yellow traffic data rate as high as CIR + EIR.

The Bandwidth Profile Algorithm

MEF10.2, Ethernet Services Attributes, Phase 2, details how the bandwidth profiling algorithm works in Figure 13, The Bandwidth Profiling Algorithm. It is presented here for your review in our Figure 3.

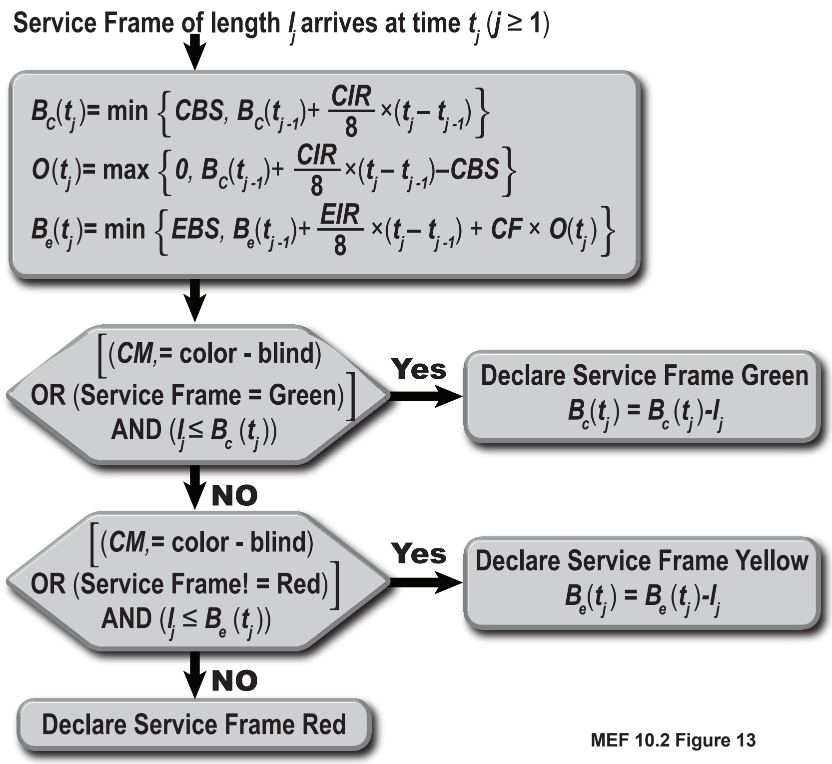

Figure 3 – the Bandwidth Profile Algorithm

This procedure defines exactly how frames are to be profiled when they enter the network. It is a beautifully constructed algorithm, simply connecting the primal forces at work in the forwarding of Ethernet frames.

On the one hand, it gives each user an average data rate at which they are allowed to transmit. Then, it gives them an amount of memory to work with, which allows them to transmit their data in bursts as computers often do, without dropping any frames. Next comes a second dimension which makes it even more powerful - It links the limited space available in the switch's memory to the rate at which frames enter and exit the switch.

If all the memory of the switch is allocated fairly among the burst sizes of all the users attached to the switch through their bandwidth profiles, and all the CIRs allocated to the customers do not exceed the outgoing transmission capacity of the switch, then no conformant frame will ever need to be dropped. Further, the total delay that all users will face will be a function of their aggregate utilization of their bandwidth profiles, and the service provider's capacity installed between the switches. It's such a simple algorithm and so appropriate.

The algorithm starts with a frame arriving at the bandwidth profiler - policer at time t_j. The bandwidth profiler - policer first refills the committed token bucket to a level of $B_C(t_j)$, according to how much time has passed since the last frame, and how many tokens were in the bucket at that time.

The committed token bucket refills at the rate of CIR. Note that there is a factor of 8 used to adjust the refill rate to equalize the bit rate to the number of bytes. The committed token bucket is considered full when it has a number of tokens equal to CBS. It is not allowed to have more tokens than this.

Next the excess token bucket is refilled to a level $B_E(t_j)$ based on the number of tokens in the bucket at the last frame, the EIR refill rate and the amount of time that has passed since the last frame. If the coupling flag is set to 1, then it also gets an optional amount of tokens $O(t_j)$ in the refill cycle equal to the number of tokens, if any, that were unable to be used by the committed token bucket during the refill because the committed token bucket filled to the CBS limit.

With both token buckets refilled, the frame is then checked against the number of tokens in the buckets to see if it can be marked green or yellow. In Color Blind mode, all frames are first checked to see if they can be marked green at the committed token bucket. If there are at least as many tokens in the committed token bucket as there are bytes in the service frame, then the frame is marked green, the committed token bucket is lowered by the number of bytes in the service frame, and the frame is forwarded into the network.

In Color Aware mode, only the frames that the customer marks green are checked against the number of tokens in the committed token bucket. If the frame can't be marked green because there are not enough tokens available in the committed token bucket, and if there is a yellow classification available, then the frame is checked to see if it can be marked yellow at the excess token bucket. If there are as many or more tokens in the excess token bucket than there are bytes in the service frame, then the frame is marked yellow, the excess token bucket's tokens are reduced by the number of bytes in the service frame, and the frame is forwarded into the network.

For color-aware service, frames that the customer marks yellow go directly to the excess token bucket for coloring, without stopping first at the committed bucket. If the frame can't be marked green or yellow, it is declared a red frame, and is immediately discarded at the bandwidth profiler – policer.

Most of us are not mathematicians, so the mathematical expressions used in the above Figure 3 may be a bit unfamiliar, and thus may not properly communicate the simple concepts underneath. See Figure 4 for a way to visualize the process using a leaking-bucket analogy.

In this analogy, the user releases their traffic to the network like bursts of water out of a faucet which is continuously being turned on and off, sometimes in big bursts, and sometimes in little streams. The water (traffic) is caught at the bandwidth profiler – policer in a bucket of size CBS, which keeps it from being lost as it is transmitted into the network.

Water that is released out of the faucet in any combination that doesn't overflow the bucket will meet its bandwidth profile. The bucket is emptied at the speed of the CIR.

A safety overflow bucket can be purchased by the customer that provides some backup protection via additional information rate and burst capacity. However the performance of frames going through the second bucket (EIR/EBS) is not guaranteed. Traffic that overflows both buckets is discarded.

The biggest limitation with the leaking bucket analogy is that unlike an ideal bandwidth profiler – policer, the leaking bucket doesn't release the water immediately, but instead makes it wait until the CIR can provide for release. In an ideal traffic policer, the frame is forwarded immediately after it is confirmed as green or yellow. There is no delay. See Figure 4 below:

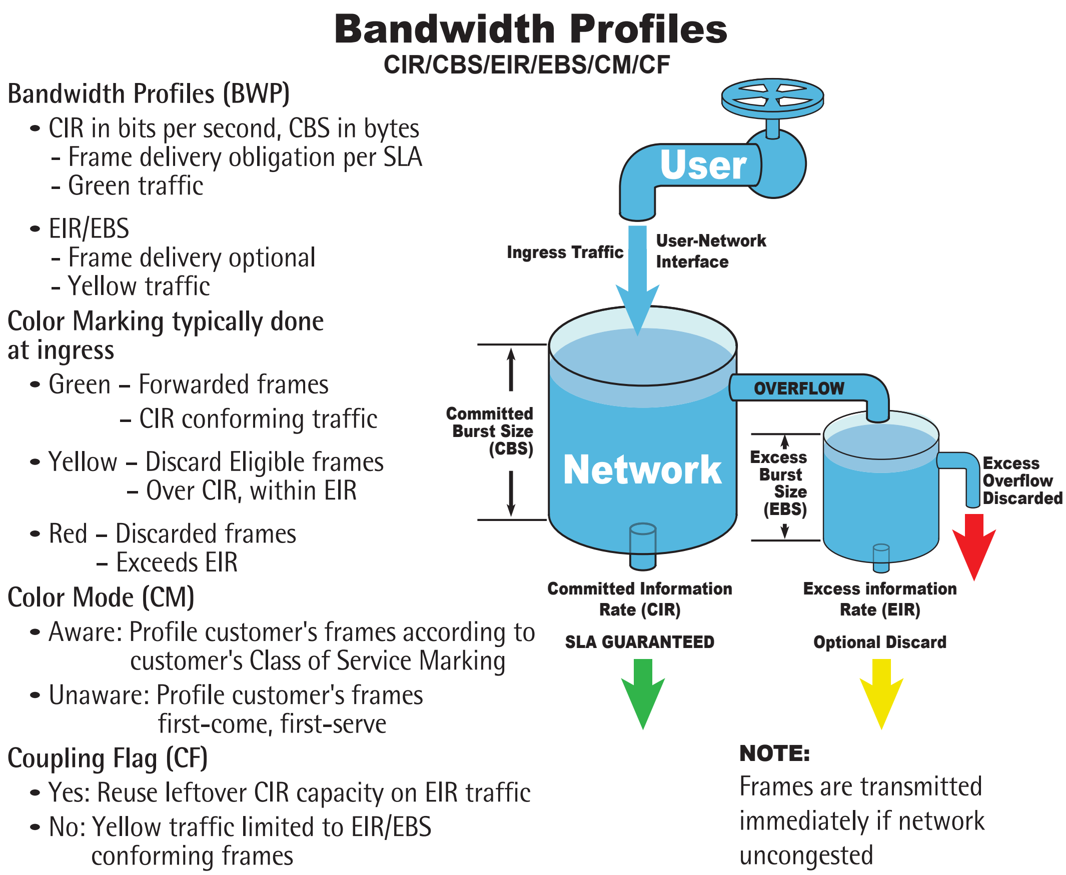

Figure 4 - A Leaking Bucket Analogy for CBS, EBS, CIR, and EIR

Traffic Shaping

One final word about the Bandwidth Profiler/Policer. Some networks add a third function of traffic shaping to this device. A traffic shaper works much like the leaking bucket portrayed in the Figure 3, providing a safety reservoir to collect bursty traffic, and metering it out into the network at the CIR.

In the shaper, the frames are held until they can be released without bursting at rates above the CIR. Instead of having a capacity of CBS, the shaper usually has much greater capacity.

There are two reasons to do this. One reason is to provide the network strict protection from overload, where all inputs to the network everywhere is limited to a maximum instantaneous information of the CIR, in which case the network backbone can be positively designed to never be overloaded. The second reason for the shaper is to protect the customer frames from being discarded if and when they exceed their CBS/CIR/EBS/EIR bandwidth profile. Rather than forward these red frames into a policer which will discard them, the shaper holds on to them until a time in which they can be forwarded in-profile so that the traffic policer will accept them.

Because the traffic shaper has a lot more memory available than the CBS/EBS bandwidth profile limits, traffic shaping can provide useful protection to customer traffic from being discarded, at the cost of adding additional delay.

Measurement Basis for the Bandwidth Profile

Figure 5 shows how the bandwidth profile CIR/CBS/EIR/EBS parameters are measured. The service frame is the fundamental entity being metered. It starts with the first bit of the destination MAC address and ends with the last bit of the frame check sequence. The size of the frame is measured in bytes, and it is the number of these bytes that is taken out of the committed token bucket or the excess token bucket when the frame is profiled.

The information rate is the rate in bits per second of the service frames. The CIR and the EIR are the rates that the committed and excess token buckets are refilling at, respectively, and are thus the maximum average rate that the service frames can achieve and be profiled as green or yellow frames.

Note that the EIR is an added capacity that with the CIR gives the total allowed information rate. So if a customer has purchased 10 Mb/s CIR and 20 Mb/s EIR, then the total information rate at his disposal is up to 30 Mb/s.

There is some variation between service providers and network equipment manufacturers in how they measure rates and frames. The standard way has been explained above. The variant way is to use Utilized Line Rate instead of Information Rate. With Utilized Line Rate, the parameters are still labeled as CIR, CBS, EIR, and EBS, but the measurement basis includes the extra bits for the minimum interframe gap, preamble and start of frame delimiter. So it is possible that the CIR from one service provider or manufacturer will be a little bit different than the CIR for another, by the difference between utilized line rate and information rate.

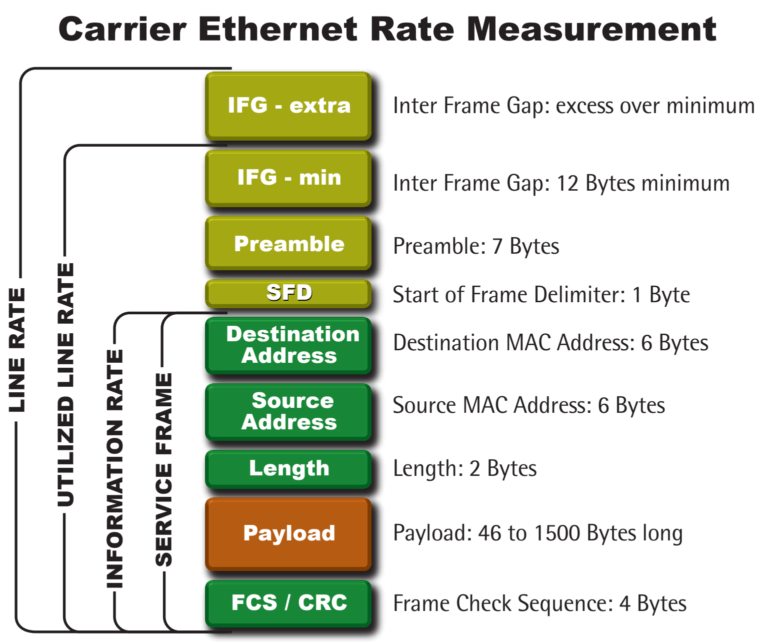

Figure 5 – Carrier Ethernet Rate Measurement

C-Tags and S-Tags

See Figures 6 and 7 for a look at tags and how they are added to Ethernet frames by the customer and service provider.

Another item which makes for small differences in the measured information rate at different points in the network is the addition of customer tags and service provider tags to the Ethernet frame.

The customer adds a C-Tag to the Ethernet frame to give it a VLAN ID number, and to possibly give it a priority code point setting, which shows what class of service that frame should receive. The C-Tag is 4 bytes long, with two bytes used for the Tag Protocol Identification set at 8100 in hexadecimal format, 3 bits used for the priority code point bits which mark the class of service to be used, and a Canonical Format Indicator bit which is not used.

The Bandwidth Profiler – Policer examines the customer's C-Tag and can "push" (add) a service provider S-Tag. The S-Tag accomplishes several functions for the service provider:

- It associates each customer frame to the Ethernet Virtual Connection that will forward that frame from one UNI to the other(s).
- It marks which class of service (Cos) the customer's traffic will receive by assigning the appropriate PCP (Priority Code Point) codes.
- It can use the PCP bits or the DEI (Drop Eligibility Indicator) to mark whether each frame is green or yellow.
- One or more customer VLAN IDs can be mapped to a single service provider VLAN ID through the use of separate S-Tags and C-Tags.

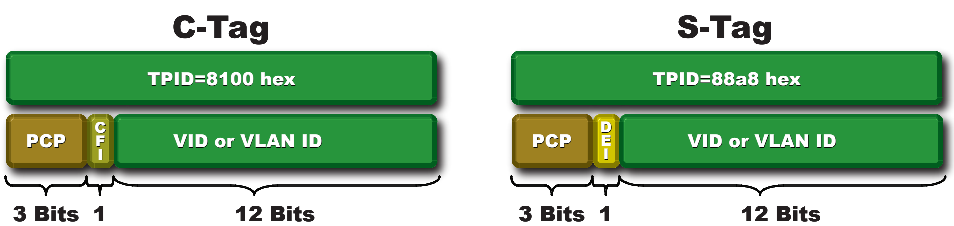

Figure 6 – S-Tag and C-Tag Format

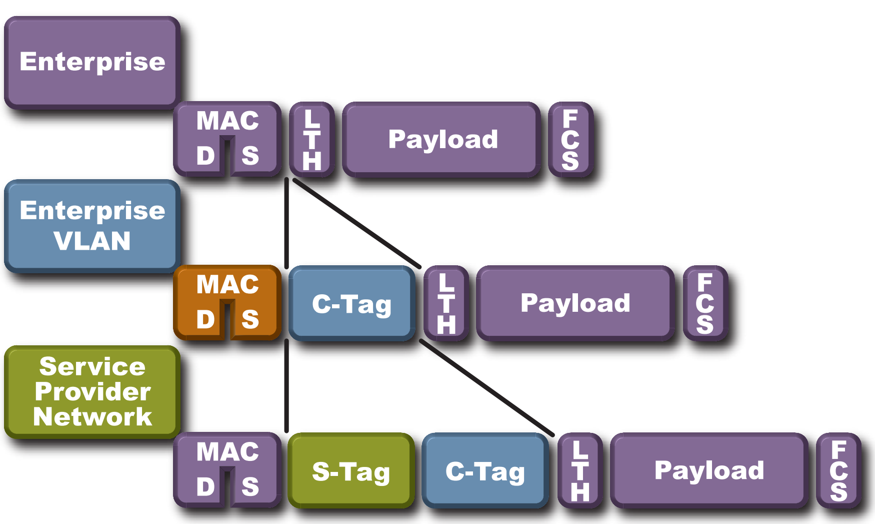

Figure 7 – Growth of the Ethernet Frame with C-Tags and S-Tags

Figure 8 shows how the measured bandwidth profile can depend on the measurement point. The customer will typically present the Ethernet frame at the user network interface (UNI) with no tag, or with just a C-Tag. Next the frame goes through the Bandwidth Profiler – Policer which adds an S-Tag to the customer's frame. In this case, the bandwidth profiling and policing function is performed by the service provider's Network Interface Device (NID) located at the customer premises.

The bandwidth profile that the customer traffic exhibits at the UNI will grow a little bit larger inside the service provider's network with the addition of the size of the S-Tag on each customer frame at the NID.

It is possible that the NID will actually be set with Bandwidth Profile Parameters CIR/CBS/EIR/EBS defined on the network side rather than the UNI side, and if so, then the testing that is done at the UNI will need to be done with all the CIR/CBS/EIR/EBS numbers adjusted downwards from the numbers set in the NID by the appropriate derating factor based on the frame size in question.

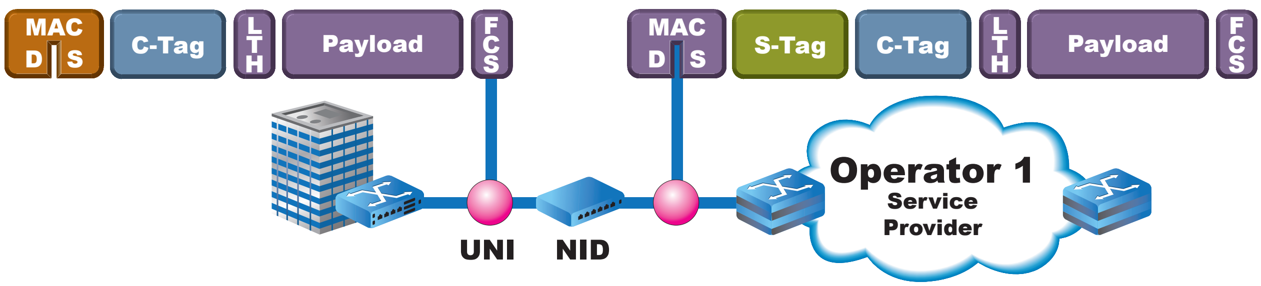

Figure 8 – The Measured Information Rate Depends on the Measurement Point

Class of Service Guarantees for Conforming Carrier Ethernet traffic

Customer traffic that conforms to the Bandwidth Profile (BWP) has a corresponding Class of Service (CoS) Performance Objective (CPO) according to the terms of the Service Level Agreement (SLA). There are many ways to name and define the Classes of Service, depending on how much the customer is willing to pay, how far the circuit goes, and how much priority the customer's traffic needs to receive relative to other traffic, to meet the requirements of the overlying customer application. In the end though, all the different Classes of Service are reduced to what level of performance is going to be met on the following performance parameters.

- **Frame Loss Ratio:**
  What is the ratio of frames that do not make it to the other side of the network compared to the total number of frames transmitted? Loss of frames causes problems for the customer, so a low Frame Loss Ratio (FLR) is desired. Some customers may be willing to pay very dearly to be assured of extremely low FLR.

  Errored frames are by rule dropped by the next network element that processes them, so frame errors usually appear as frames lost at the measurement point at the other side of the network. In applications where the local access link connected to the UNI is contributing a large proportion of the total errors, service providers may also want to count frame errors as frame loss, because the customer's equipment by definition will need to discard the errored frames when they are received and they will thus have the same effect as lost frames on the customer's equipment.

- **Frame Transfer Delay:**
  How long do the frames take to travel from one UNI to the other? Frame Transfer Delay (FTD) answers this question. One-way services like video streaming are well suited to this one-way measurement. There are some applications such as TCP file transfer, where the bandwidth delay product based on the round trip delay is the key performance factor.

- **Frame Delay Variation:**
  Frame Delay Variation (FDV) is particularly important for services that make use of playout buffers, such as music or video streaming. FDV must be matched to the buffer size, or buffer overrun or underrun may result, and the customer will experience impaired service.

  The FDV measurement reports the minimum, mean, and maximum measured variation between the frame transfer delays of a group of frames. The measurement is ideally presented only with positive values – the lowest possible variation is zero, representing at least two frames that travelled through the network with the same measured frame transfer delay, and where both of those delays were the lowest values measured. The maximum reported FDV is the difference in frame transfer delay between the slowest transfer and the fastest transfer.

  Like FTD, FDV is ideally measured from two ends of the circuit by measurement units that are fully synchronized with each other, with locally supplied synchronization sources to each unit such as GPS or local IEEE1588 master clock available locally at each end being the preferred synchronization sources. Where these clock sources are not available to synchronize both ends of a one-way delay measurement, then an inter-frame delay variation measurement may be used. The inter-frame delay variation measurement doesn't require clock synchronization.

- **Availability:**
  AVAIL (Availability) is measured as a count of unavailable seconds and a percentage availability. When the circuit is available, the customer can use it and performance measurements can be calculated. When the circuit is unavailable performance measurements become meaningless because the service can't even be used.

  Service Performance is often so good that even a relatively brief period of unavailability could grossly impact the overall values of the other performance measurements if they were counted during that unavailability. For this reason, FLR, FTD, and FDV are not calculated during unavailable time. Time becomes unavailable at the onset of n (for instance, 10) consecutive severely errored seconds, and becomes available again at the onset of n ((for instance, 10) consecutive non-severely errored seconds.

  The threshold for determining an unavailable second is based on the FLR, and may be different for different applications or Classes of Service, and may be set to 50%. A separate Class of Service Objective may be set for Availability.

  Availability results should be displayed to the technician during the Service Performance Test, and the measurements of FTD, FDV and FLR should be suspended during unavailable time.

Measurement Topology for Ethernet Service Activation Test

The measurement topology for accessing the circuit at the measurement points is shown in Figure 9.

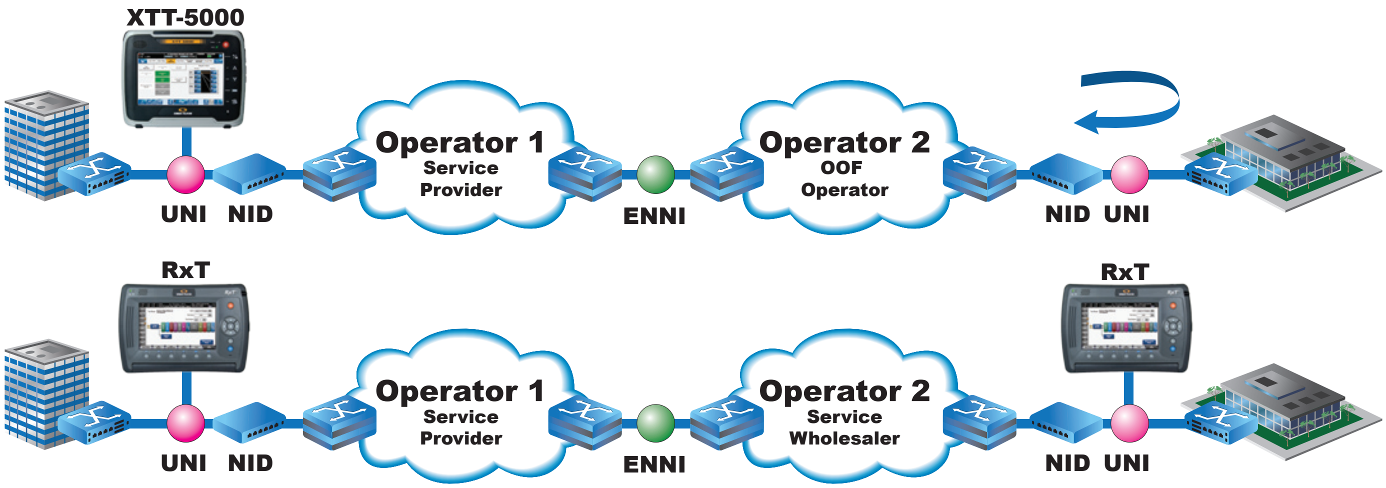

Figure 9 - Carrier Ethernet Service Measurement Topologies

Measurement topology is a key part of the measurement. The Y.1564 preferred topology is a two-point measurement which uses a test set plugged in to each UNI as the measurement points at which to make the measurements. Further, each test set is locally synchronized to an accurate network clock such as GPS or an IEEE1588 master clock.

Y.1564 favors one-way measurements, but notes that round trip measurements may be desirable from an operational point of view, and it will be up to the operator to verify that they are getting valid operational results if they use a loopback measurement topology. Where synchronizing clock sources are not available at each end of the line, it may be possible for two pieces of test equipment to get a quasi-synchronization that can provide one-way delay measurements. The limitation with this method, though, is that it is susceptible to error from asymmetric delay in each direction on the service. The test equipment has no way to remove this source of error without access to local clocks for synchronization.

Operators who don't have access to synchronizing clocks at each end, or who only have a technician at one end, may have situations where they need to investigate making a loopback measurement satisfy their needs using a NID or other equipment with known loopback performance.

Carrier Ethernet Service Activation Test Methodology

Now that we have fully reviewed how the Service Level Agreement is based on a bandwidth profile that allows the customer to transmit data to certain rates and bursts, and seen what kind of performance objective goes along with that SLA, let's turn our attention to what the service provider needs to achieve during service activation testing.

- First, the test procedures should find any configuration errors that might be preventing the service from performing as it should so that these errors can be corrected.
- Second, the procedures should verify that the network is operating with the required quality to meet the SLA.
- Finally, the procedures should take the minimum amount of time needed to do the job well.

We'll review the Y.1564 procedures here and see that they nicely satisfy these criteria. Before we review the procedures, we'll go over one more term used in Recommendation Y.1564 – SAC, or Service Acceptance Criteria.

At service activation, testing is performed to see whether the service is performing at or better than the Service Acceptance Criteria (SAC). The SAC will generally be chosen to give the Service Provider high confidence that a service accepted according to the SAC will meet its SLA. The SAC may vary from the exact SLA performance objectives due to the service provider partnering with other service providers to deliver the service, and the service may be activated one section at a time. Optionally, the service provider might want to allow some margin for the circuit performance to deteriorate over time, and still be better than the objectives or guarantees in the SLA. Whatever the rationale, the service provider sets the SAC.

The measurements for characterizing Carrier Ethernet Service performance are found in Y.1563, Ethernet Frame Transfer and Availability Performance.

The Ethernet Service Activation Test Methodology is broken into two parts, **the Service Configuration Test** which finds and corrects configuration problems, and **the Service Performance Test**, which verifies that performance meets the SAC and is stable over time.

First let's review **the Service Configuration Test** procedures:

CIR test – The CIR Configuration test is performed to verify that the proper data rate has been configured in the network. If the CIR is set properly, the test should pass. The stepped nature of the test makes it easy to differentiate a CIR configuration error from a general network performance problem. A single set of procedures adequately handles the color unaware case as well as the color aware case. A fast version of the test is provided, which goes directly to the CIR. A slower version of the test is illustrated, which is useful if the fast version doesn't pass.

EIR test – The EIR Configuration test has two procedures – one for color aware service and one for color unaware service. The separate procedures are necessary because the full set of performance information can only be determined from the color aware test, leading to a better understanding if there are any underlying problems in that case. Because the color unaware frames are not marked, it is not possible to see if there was a problem getting the green frames through the network when they were transmitted in the presence of yellow frames, so we are just limited to measuring the total information rate to determine if the test is successful.

Traffic Policing test – In the traffic policing test, frames are sent faster than CIR + EIR in order to see if the traffic policing function is working correctly. The goal is to see that the total transmission rate is limited to CIR + EIR, plus some small overage margin set by the service provider, according to how their traffic policers operate. Some services may not have a traffic policing function and don't need to have this procedure performed.

The Traffic Policing Test has two versions, one for non-color aware and the other for color aware. The separate procedures are necessary because there is more ability to make sure that all the green frames have been properly transmitted in the color aware case, just as in the EIR test. In comparison, for the non-color aware case it is not known how the bandwidth profiler profiled each of the frames, so it is not possible to say which frames could have been dropped or could have been delayed at levels worse than the SAC. So the non-color aware test may be limited to evaluating the received information rate and verifying that it is at the full CIR, less an allowance for SAC FLR. Figure 10 shows the interaction of the CIR test, EIR test, and traffic policing test:

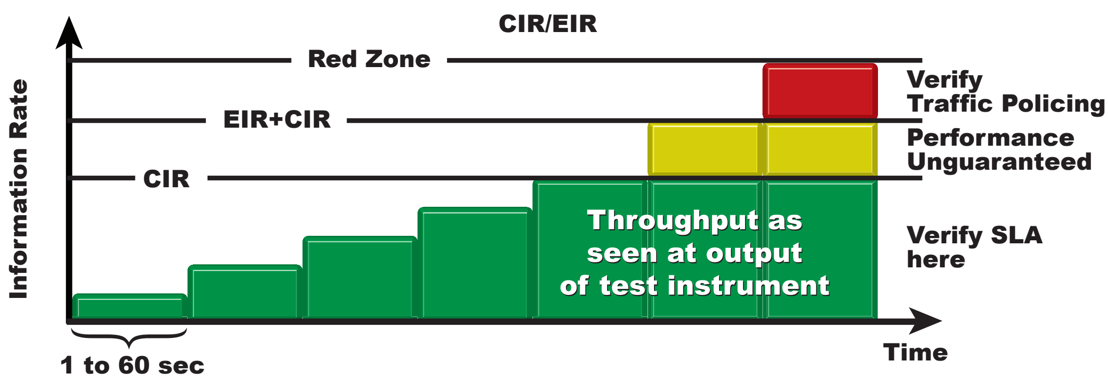

Figure 10 – The Y.1564 Rate Configuration Tests

CBS and EBS test status – The CBS and EBS tests are described as experimental in the standard. This is because at the time of consent there was limited field experience with the procedures, and limited understanding of the percentage of the embedded base of newer bandwidth profilers in the network which are conformant to MEF 10.2, Figure 13, The Bandwidth Profile Algorithm.

If a service provider has these kinds of bandwidth profilers deployed, then these procedures may be used as is to verify configuration of the CBS and EBS at the time of service activation. Further, these procedures may help service providers deal more effectively with complaints from end users who are having throughput problems with their service, as very often these problems are due to the bursty nature of their traffic. The constant-spacing-frame transmission procedures which have heretofore been used are not helpful for diagnosing burstiness problems. As frame transmission is inherently bursty, burst procedures should be used where they fit the traffic policers and bandwidth profilers deployed in the network.

CBS test – The CBS Configuration Test uses a single procedure for color aware and color unaware service. The test exercises the CBS token bucket in the bandwidth profiler and traffic policer to uncover a possible misconfiguration. It works the token bucket nearly through its whole range, leaving an adequate margin to ensure that no frame is marked yellow or is discarded.

The margins are set just a little bit loose, so that all brands of equipment that implement MEF 10.2 Figure 13 should pass. Within the procedure, there is room for about 1 msec of timing error in token bucket refilling, standard Ethernet 100 ppm clock accuracy, and large frame delay variation which can result in misordered frames occurring in the network between the test set and the traffic policer, which may not be co-located.

The test only uses a 10% burst duty cycle, which verifies the proper CBS configuration of the network equipment without putting undue stress on weaker bandwidth profiler – policers. The 10% duty cycle means that 10% of the frames are sent in burst mode, and the other 90% of the frames are sent with even spacing at constant information rate. Figure 11 shows a sample series of frames from the CBS test, along with a frame-by-frame depiction of the CBS token bucket.

The frame arrival sequence is shown in the upper half of the figure, and the token bucket level just before and after the marking of each frame is shown in the lower half of the figure. For simplicity's sake a relatively small CBS token bucket of 6000 bytes is used, and the frames are an even 1500 bytes long. The procedure starts out by stopping transmission to ensure that the CBS token bucket is full to overflowing. Then it starts the burst cycle, bursting at the maximum attainable rate until the token bucket is almost empty. It next pauses to let the token bucket refill about half way. Then it goes into constant rate transmission at the CIR to keep the token bucket half full. It then starts the cycle all over again, by filling the bucket to overflowing and then bursting. It continues this cycle for as long as the user sets. The service should meet its SAC during this test.

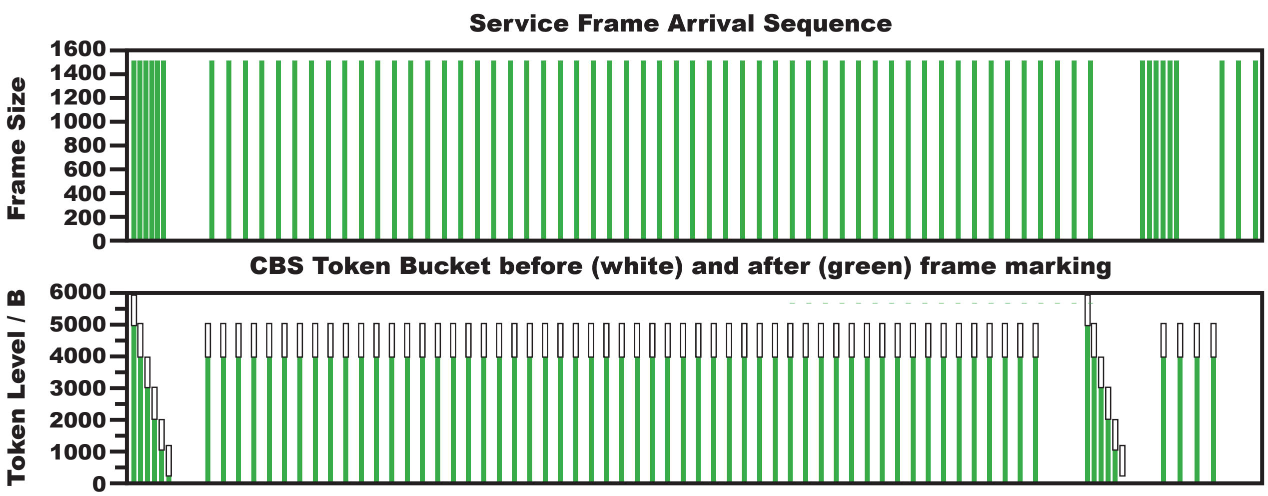

Figure 11 – CBS Test Service Frame Arrival Sequence and Token Bucket Status

EBS Configuration tests – Y.1564 Appendix 1 uses three EBS Configuration Tests, one for the special case where the CIR is zero, one for color aware, and one for non-color aware. When the CIR is zero, the procedure is essentially the same as the CBS Configuration test. In the color aware case, two separate test flows are set up, one for the green frames, and one for the yellow frames. They are both bursted at 10% duty cycle pretty much according to the CBS procedure. The one complication is that sometimes frames from the green flow and the yellow flow would be transmitted during the same period of time, and so guidelines are given to keep the process under control while allocating the transmitter between the two flows.

The color non-aware case tests the excess token bucket while keeping the committed token bucket nearly empty throughout the whole test. The committed token bucket is kept empty because the CBS token bucket has already been tested in the CBS test, and because the test achieves a higher overall information rate to exact the maximum performance out of the line at nearly CIR+EIR.

Users will note that the burst tests do not quite achieve the full information rate. This is because the tests are designed to go through any MEF 10.2 Fig 13 conformant traffic policer without dropping or misprofiling any frames, allowing for tolerances between different manufacturers' equipment, as well as allowing for maximum clock difference between test sets and traffic policers, plus some frame delay variation between the UNI and the traffic policer. At equipment type approval, service providers may want to use tighter margins to bench mark the accuracy of different policers against each other.

Service Performance Test

The Service Performance Test is performed after all the configuration errors, if any, have been identified and corrected. With a stable network configuration in place, a performance test is run, as shown in Figure 12. If there is more than one EVC on the line, they are all run together at CIR. The minimum recommended time is 15 minutes, and recommended times vary between 15 minutes and 24 hours, depending on network complexity, number of service providers involved, geographic span of service, reliability of network, and likelihood of trouble. The bandwidth profile is individually verified for each service, and may be associated with the entire port, to each of two or more EVCs, or to each of two or more Classes of Service inside an EVC. The multiple classes of service per EVC may be marked with S-Tag PCP bits or IP DSCP bits.

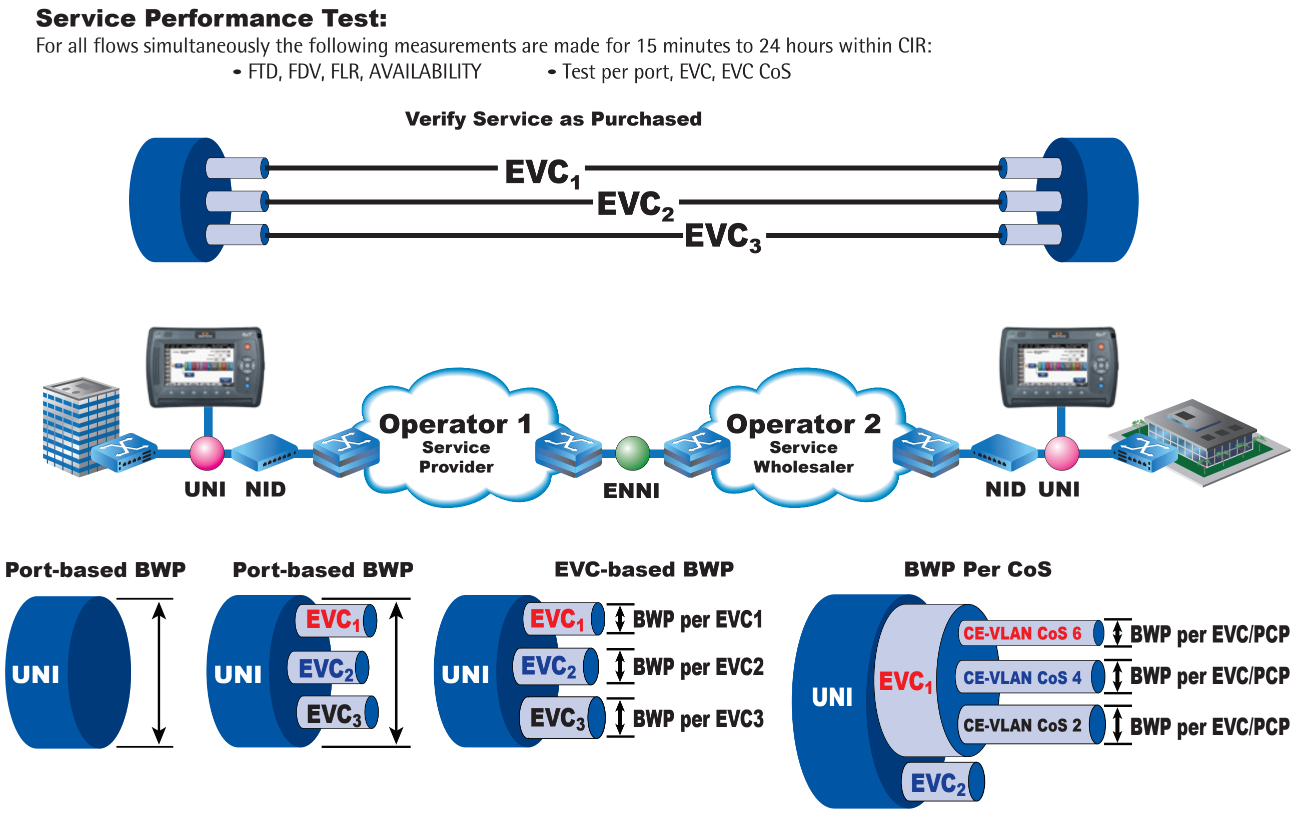

Figure 12 – Service Performance Test

Why Y.1564 is Better than RFC2544

Y.1564 is the first test procedure to be written specifically for verifying the configuration and performance of Carrier Ethernet services prior to turning the service over to the customer for their use. It is the first procedure to include verification of the Bandwidth Profile, and the first procedure to test to the measurements specified in today's service acceptance criteria.

Service Providers have been using RFC2544 at service activation for lack of any proper published standard. RFC2544, Benchmarking Methodology for Network Devices, is written for device benchmarking, and has procedures that test the absolute performance of network devices, allowing service providers to compare equipment made by different manufacturers to see which one they like best. It is not designed to test Carrier Ethernet services at the time of service activation.

The throughput test of RFC2544 is designed to see if the maximum throughput of the port can be obtained. Although this is a good test for benchmarking switches and routers, in general, it is not good to generate line-rate traffic on an in-service switch because it may harm service to other users. The RFC2544 latency test measures latency of one frame every two minutes, the limitation of test equipment a decade ago. Y.1564 includes no such limitations on FD measurement.

The frame loss rate is specified by RFC2544 is to be performed throughout the speed range of the device under test. That is not desirable now, as the line-rate frame transmission may overload in-service network equipment, and is inappropriate for customer services which are policed to rates far lower than the line rate. Similarly the burst test, which runs bursts at unlimited rates until frame loss is observed, has the potential to harm in-service networks which are not protected by a traffic policer.

Y.1564's burst test is the first in the world to be fitted to the actual CBS and EBS service parameters of the bandwidth profile which have been sold to the customer. Where a traffic policer is provided, the language of RFC 2544 creates an inefficient procedure, because it does not test to the size and operational characteristics of the CBS and EBS parameters in the bandwidth profile. In similar fashion RFC2544 provides no guidance for how to modify procedures for color-aware and color-blind service. RFC2544's system recovery test for recovering from device overload is unnecessary, and RFC 2544 network element reset test is totally inappropriate to be performed on an in-service network.

All this said, service providers and test equipment manufactures have generally modified their RFC2544 procedures for service activation so that they are more RFC2544 "inspired" than RFC2544 "compliant". And this is a problem, because people don't do it the same way, and no one's implementation has come close to what is really appropriate for Carrier Ethernet service configuration verification and performance verification. Thus, we feel that Y.1564 is an enormously valuable contribution to the industry, addressing a weak spot in the test and measurement toolbox available to service providers. Figure 13 summarizes the advantages of Y.1564 over RFC2544.

Advantages of Y.1564 over RFC 2544

- Verifies configuration and performance of CIR/CBS/EIR/EBS/CM
- Faster method
- Multiple streams tested, not just one
- Tests network rather than a single device
- Tests to SLA rather than to network failure
- No time wasted on excessive permutations
- Latency test tremendously improved and made on every frame
- Measured during normal operation when SLA should be met
- Adds frame delay variation
- Skips tests designed to crash the network
- Availability results calculated and presented to technician
- Guidance for using and resolving differences in information rate and utilized line rate

Figure 13 – Advantages of Y.1564 over RFC 2544

IntelliSAM™

Summary

Y.1564 Ethernet Service Activation Test Methodology provides the key procedures to test, troubleshoot, and accept today's sophisticated Carrier Ethernet services. With these procedures, service providers can unequivocally demonstrate proper service configuration and performance from UNI to UNI, giving total confidence that the SLA will be met according to the bandwidth profile purchased, and the promised class of service. These efficient procedures save technicians' time while making sure that all the problems are eliminated from service before it is handed over to the customer, enhancing the customer experience and reducing customer churn. They are a huge step forward from procedures previously used, and should be deployed in the network as rapidly as possible. Please contact your Sunrise Telecom representative today for the opportunity to try out these procedures using IntelliSAM™ in your network.

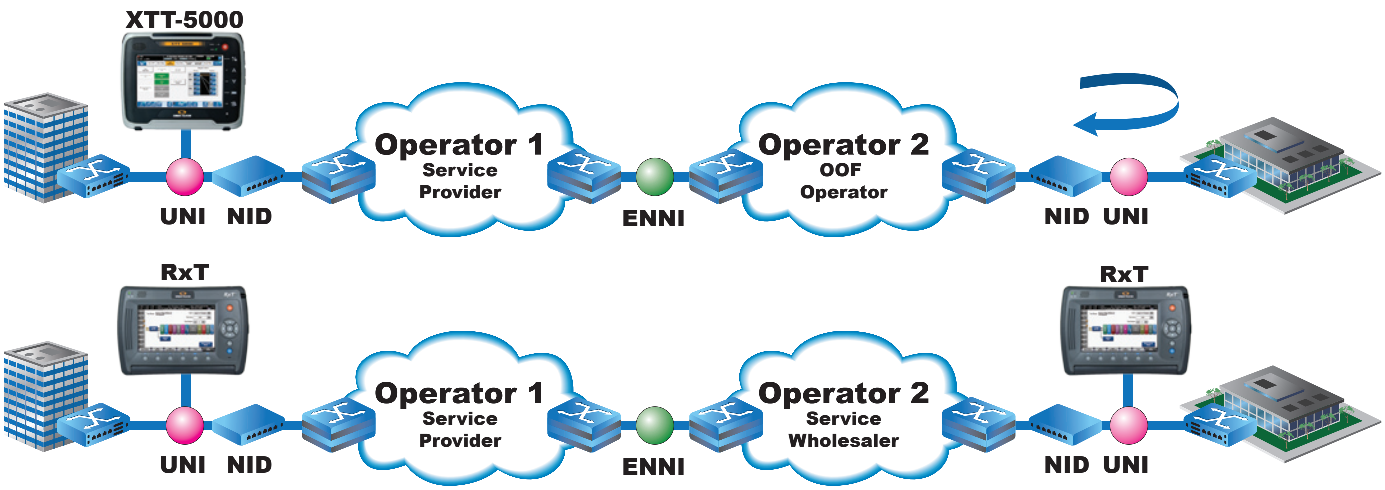

Figure 14 - Carrier Ethernet Service Measurement Topologies

Sunrise Telecom's Leadership Role in Driving to Second Generation Capabilities in Y.1564

Sunrise Telecom's numerous contributions to Y.1564 helped drive the document to its final form. In particular the company's original text contributions and theoretical explanations facilitated the evolution of the document from Y.156sam's first generation set of capabilities which included CIR/EIR/Traffic Policing functionality, to the core second-generation set of features found in the finished document, which expanded the bandwidth profile coverage to CIR/CBS/EIR/EBS/CM/Traffic Policing capability. As of the time of consent, Sunrise Telecom is believed to be the only company in the world with working bandwidth profile test coverage of this full range of CIR/CBS/EIR/EBS/CM/Traffic Policing, with its IntelliSAM™ test application.

About the Author

Paul Marshall is a co-founder of Sunrise Telecom, serves as its Chief Technology Officer, and serves on the company's board of directors. He represents the company at the ITU, MEF, SCTE, and HPNA forum. His is a Co-Chairman of the MEF Marketing Committee OAM Working Group, and serves on the board of the HPNA Forum. At the ITU he participated actively in the development of Y.1564 Ethernet Service Activation Test Methodology, serving as the focal point for Sunrise Telecom's contributions.

SUNRISE TELECOM®
Productivity Rising™

302 Enzo Drive San Jose, CA 95138 • Tel +1 408-363-8000 • Fax +1 408 363-8313

© 2011 Sunrise Telecom Incorporated. All rights reserved. Specifications subject to change without notice. All product and company names are trademarks of their respective corporations. Do not reproduce, redistribute, or repost without written permission from Sunrise Telecom. C_0200 REV. A00 Feburary 2011