IDEAL INDUSTRIES NETWORKS

IDEAL

Application Note

UniPRO MGig1
Carrier-grade Ethernet tester

UniPRO SEL1
Intelligent loopback device

Y.1564 Ethernet Service Activation Method for Testing

Contents

1. Why does the Carrier Ethernet Industry need a new test?

Today's Ethernet Wide Area links are far more complex than their LAN counterparts. Gone are the days when a leased-line or access circuit was a simple unmanaged bit-transport link.

The services WAN and Metro access links and leased lines now carry are increasingly sensitive to the quality of service they receive as they pass through the link and the network.

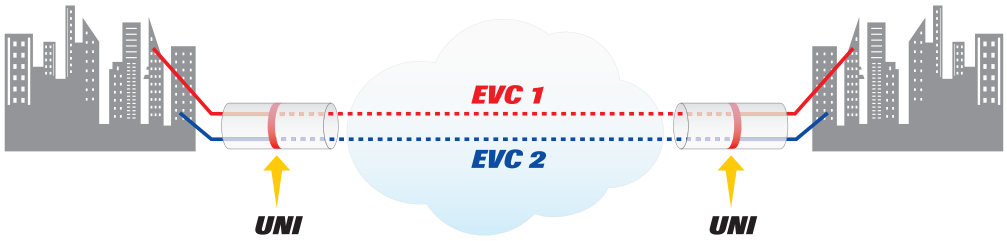

UNI - User Network Interface

Diagram 1: Ethernet virtual connection over carrier Ethernet

Today's Ethernet links frequently need the ability to deliver segregated services to multiple customers in a building (or a local area). This requires separate VLANs to be configured through the network for each customer.

Increasingly, each customer also has a requirement for multiple QoS (quality of service) levels to differentiate high priority traffic – for real-time and streaming services like voice and video - from less sensitive data. These often require separate VLANs to carry the differently classified services.

There are now many Ethernet links where VLANs within VLANs (often called Q-in-Q) have to be configured and tested. For example: to give each independent customer mentioned above its own services securely segregated from other customers. This dictates a 'customer VLAN' for each customer within the Ethernet link and then a further level of 'service VLANs' to enable the provision of different QoS (quality of service) levels for different customers.

Then there's MPLS (multi-protocol label switching), which is increasingly used within networks to send traffic rapidly to where it needs to be. MPLS has the great advantage that switches only need to read a few header bytes and not the whole header - so it's faster and more efficient than IP based routers. But from a service turn-up point of view it is yet another additional layer of complexity that the testing needs to cope with before a link can be checked as SLA compliant and stable.

But here's the growing problem: The industry's favourite testing scheme RFC2544 was never designed to test such services.

In fact it was designed as a lab-test to enable carriers to benchmark one piece of network equipment (such as switches, routers etc.) against another when making purchasing decisions.

It got adopted as the service turn-up acceptance test because, when Carrier Ethernet services first began, there was no other standardised test regime available.

But RFC2544 has a number of major drawbacks.

- It can only test one parameter of a service (stream) at a time. Which is hardly a realistic synthesis of the real-world of the multiple concurrent services (streams), VLANs and QoS requirements in most Access and Metro Ethernet links today.

- Because every parameter for each test has to be tested sequentially – RFC2544 testing can take days to complete. At considerable expense.

- As RFC2544 was designed to test isolated equipment, it tests many parameters by maxing-out to device overload or network failure.

Whilst this is harmless in lab tests, throwing maximum line-rate into a busy working network can overload network equipment and cause service problems to existing customers. This is far from ideal and unnecessarily time consuming.

- And, being a lab test, it was designed for single pieces of equipment in one physical location (the lab). Hence its unsuitability for far-end link testing.

2. Carrier Ethernet Link Topologies

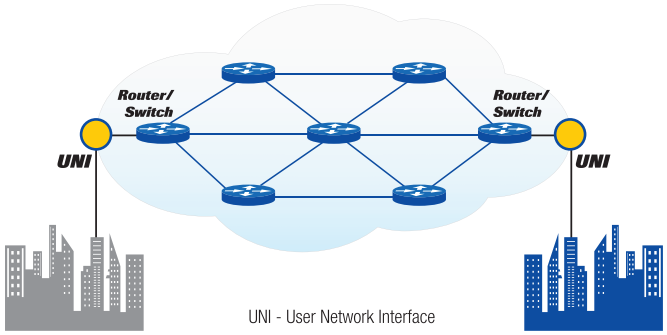

Diagram 2: Carrier network cloud

2.1 E-Line

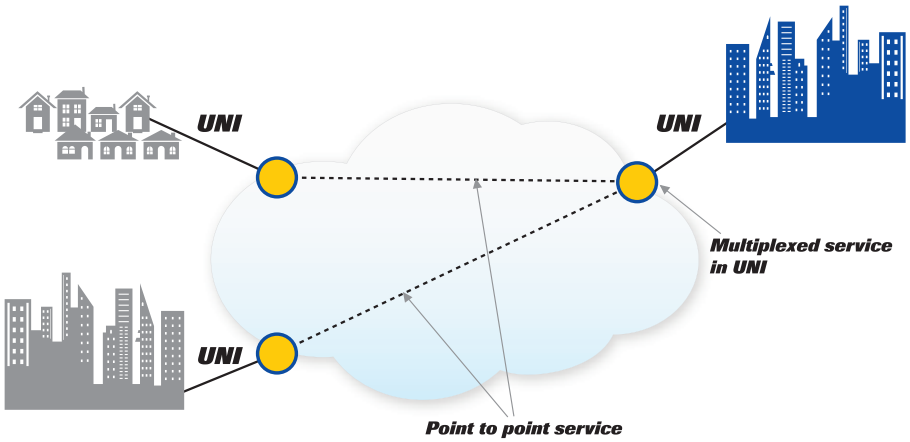

Diagram 3: E-Line services

- A point to point Ethernet Virtual Connection (EVC), Ethernet Private Line (EPL) or Ethernet Virtual Private Line (EVPL).
- Can run through one service provider's network or multiple networks
- Can be Same-City, National or International
- The simplest configuration. Forwards all frames between two locations

2.2. E-LAN

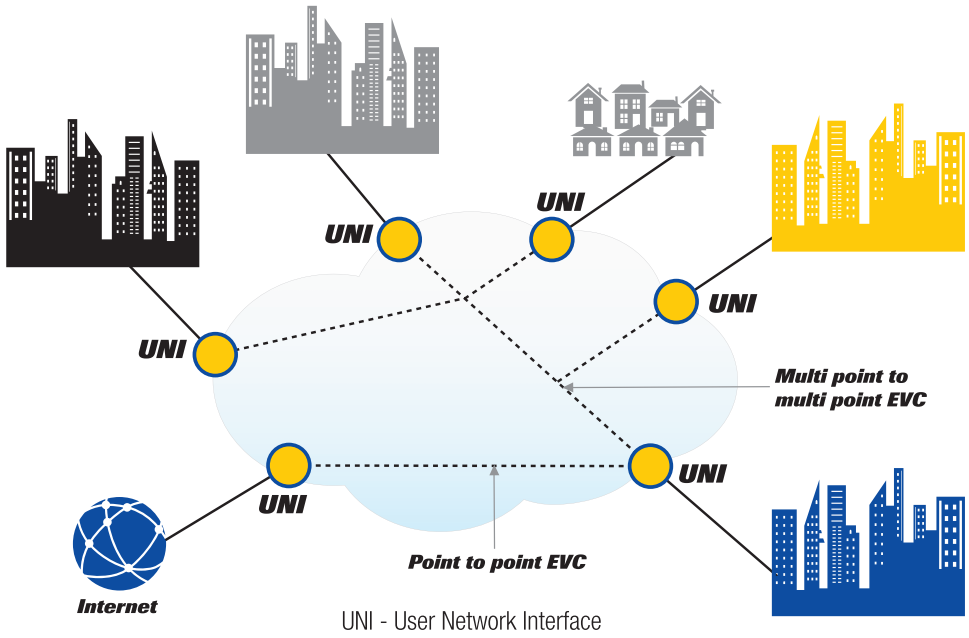

Diagram 4: E-LAN service

- A collection of multipoint-to-multipoint Ethernet Virtual Circuits, , Ethernet Private Lines (EPL) or Ethernet Virtual Private Lines (EVPL).
- Can run through one or several service provider networks
- Access (UNI) can be in same-city, national or international
- Connects multiple locations together with equal connectivity – as though they were on the same private standard LAN even though parts of the 'campus' may be many thousands of miles away.

2.3 E-Tree

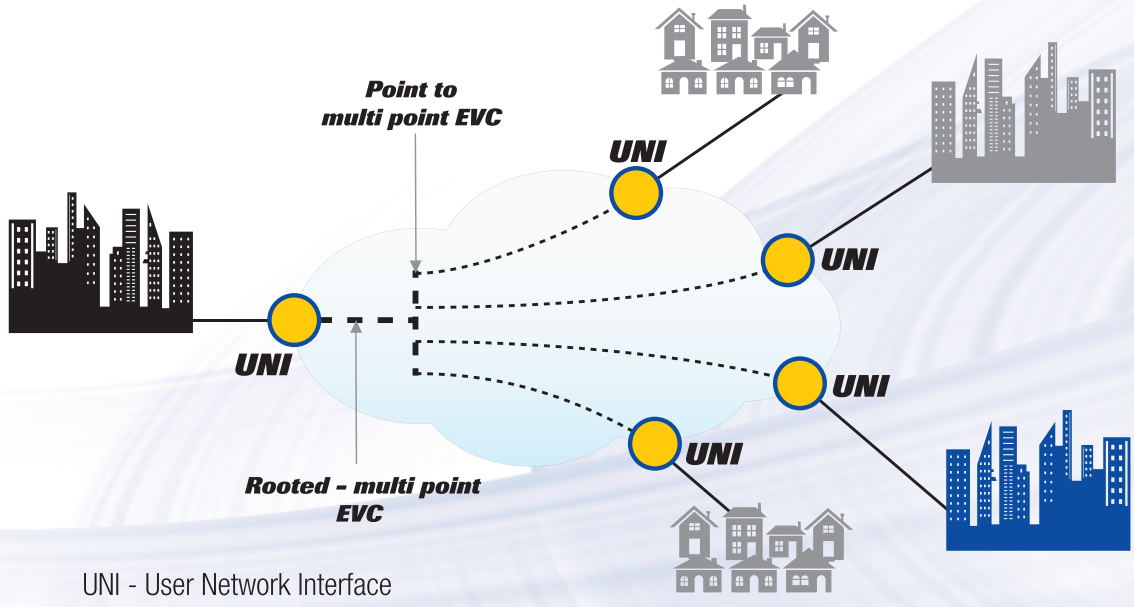

Diagram 5: E-Tree service

- Based on multiple point-to-multipoint Ethernet Virtual Circuits, Ethernet Private Lines (EPL) or Ethernet Virtual Private Lines (EVPL) with a 'root' and several 'leaves'

- Each leaf is only allowed to communicate with the root but not directly with each other except via the root

- A typical example is a company headquarters (root) connecting to many branch-offices (leaves in network terms) with all of the branches only communicating via the headquarters server or network.

- The network leaves can be same-city, national or international.

3. Why Bandwidth is Profiled

Ethernet traffic is by its nature 'bursty'. In its native state, when it has traffic to send, it will attempt to do so at or near full line rate.

Innately, Ethernet is a very 'selfish' protocol and individual interfaces have no interest in allowing other interfaces to share the available bandwidth.

But usually between each interface's 'bursts' there are periods of silence allowing other interfaces to transmit their 'bursts'.

It was because of the indiscriminate nature of Ethernet that QoS (quality of service) was introduced back in 1994.

	Bandwidth	Frame Size	Priority	Latency
Video (Conference)	Medium to high	Medium to large	Very high	Very low
Voice	Small	Small	High	Low
Data	Small to high	Variable	Low	Any

When it comes to WAN connections, generally there is far less bandwidth available than on the user's LAN and in the old days of fixed bandwidth WAN 'pipes', these bursts frequently led to the pipe being maxed out. However traffic would be evened-out because, in general, data is not highly time critical and could wait until capacity came available.

But when synchronous services like voice and video started to be encapsulated in asynchronous Ethernet – problems arose because as soon as network congestion was encountered the service would break-up (due to packets being delayed in switch and router buffers or being 'dropped' by the network and needing re-transmission).

3.1 Buried QoS information

Although QoS is used within the enterprise LAN, it is in Layer 2 which, by the time the LAN Ethernet packets have been encapsulated into Layer 4 frames for carrier Ethernet transmission, is effectively buried deep within those frames. The last thing we want is for switches and routers to have to disassemble large frames to access the low-level QoS or any other information – the processing, memory and latency overheads would be completely unacceptable.

Also, as all e-commerce traffic is heavily encrypted, even if the Layer 3 and Layer 4 routers/switches tried to access the QoS flags, they would find it impossible because it is within the encrypted data.

So a solution was needed that in some way codes the high priority frames in an easy to read location – inside the frame header.

3.2 The demise of fixed diameter pipes

The final piece of this puzzle is that most service providers no longer sell 'fixed diameter' pipes.

Instead, recognising the bursty nature of Ethernet, they install a larger 'pipe' on which they sell a guaranteed minimum bandwidth available at all times (called CIR – the committed information rate) PLUS an extra capacity (analogous to the fast lane on a motorway) which allows traffic to burst to a much higher speed. This additional bandwidth (called EIR – the excess information rate) is usually cheaper, but it is not guaranteed. In fact it is only available when the network is not busy and capacity is lying idle.

3.3 Why profile bandwidth?

3.3.1 Commercially, the service provider needs to ensure that the customer is getting the guaranteed minimum bandwidth (CIR) agreed on the SLA.

3.3.2 To ensure that each user's traffic is throttled to no more than the CIR+EIR bandwidth - since if customers were allowed to burst past this they would overload the network and bring about service degradation to all the other users of the network. Equally, if other customers were allowed to flood the network with traffic their own would suffer so it is in no-one's interest to allow traffic greater than the network is dimensioned for.

From the user's perspective the traffic also needs to be profiled so that high important, time-critical, or time sensitive traffic is treated as high priority from end-to-end. And - bearing in mind that the Layer 2 QoS flags (headers) which are used in the enterprise LAN are not readily accessible at Layer 3 and Layer 4 (unless specific action taken in the carrier network) - a new, simple, Layer 3 /Layer 4 marking scheme was needed.

3.4 How bandwidth is profiled?

The principle of bandwidth profiling requires user networks to profile their own traffic before sending it to the external service provider network. Otherwise, if large amounts of over-SLA traffic are simply sent, large proportions will be discarded by the service provider's profiling and policing functions leading to a very poor overall service from the link. The function of the own-network bandwidth profiler is similar to that of the service provider functions described here:

Bandwidth profiling gives each user of the service provider's network an average maximum data rate at which they are allowed to transmit. Allied to this it allocates them an amount of buffer memory at the network switch allowing them to transmit their data in bursts which can be temporarily stored in the buffer, if necessary, while waiting for network capacity to become free.

This is made extremely powerful because the profiler then links the limited amount of buffer memory on the service provider's switches to the rate at which frames enter and exit the switches.

In theory, and in practice, if the buffer memory of the switch is allocated in this way, and the totality of customers CIR's do not exceed the switch's outgoing transmission rate, than no good frame should ever need to be dropped.

The bandwidth profiling parameters are:

CIR – Committed Information Rate

The average maximum information rate at which the customer's SLA permits traffic to be sent to the service provider network.

CBS – Committed Burst Size

Customers ideally want infinite burst capability. But service providers cannot afford to dimension networks for that.

Many service providers allow customers to transmit bursts that are temporarily faster than the CIR - provided that the maximum number of bytes that are allowed to exceed the CIR in any one burst do not exceed CBS.

EIR – Excess Information Rate

This is the average maximum rate of additional traffic that the customer can send on top of CIR for the network to forward on a non-guaranteed 'best efforts' basis if capacity is free.

EBS – Excess Burst Size

If EIR capacity exists, customers can still burst above the maximum EIR rate. EBS defines the maximum number of bytes that can burst faster than CIR+EIR.

Colour Mode

We mentioned earlier the need for a simple classification method for high importance traffic (usually, but not limited to, video and voice) and carrier Ethernet now carries a simple colour-labelling system to do this.

Traffic is coded as green, yellow or red as it traverses the service provider network.

Users or service providers s have the ability to pre-code frames with green (high priority) or yellow (normal priority).

Where the customer or service provider is marking the frames as green or yellow – this is said to be 'colour aware' - the network should take this into account at the bandwidth profiler and traffic policer.

Some legacy networks are still 'colour blind' and the service provider network instead uses their own algorithms to set the colour for traffic from these networks.

Green would generally be reserved by for high priority user frames. As long as the green traffic remains below the CIR, it is guaranteed low latency and high priority.

All traffic above CIR and below EIR is labelled yellow and this uses any spare (non-green) capacity remaining in the CIR (or waits until there is capacity). It can also use the EIR when the network has spare capacity available.

Any traffic that is sent to the network and over a period of time exceeds the CIR and EIR averages is classified by the network as 'red' and discarded (lost).

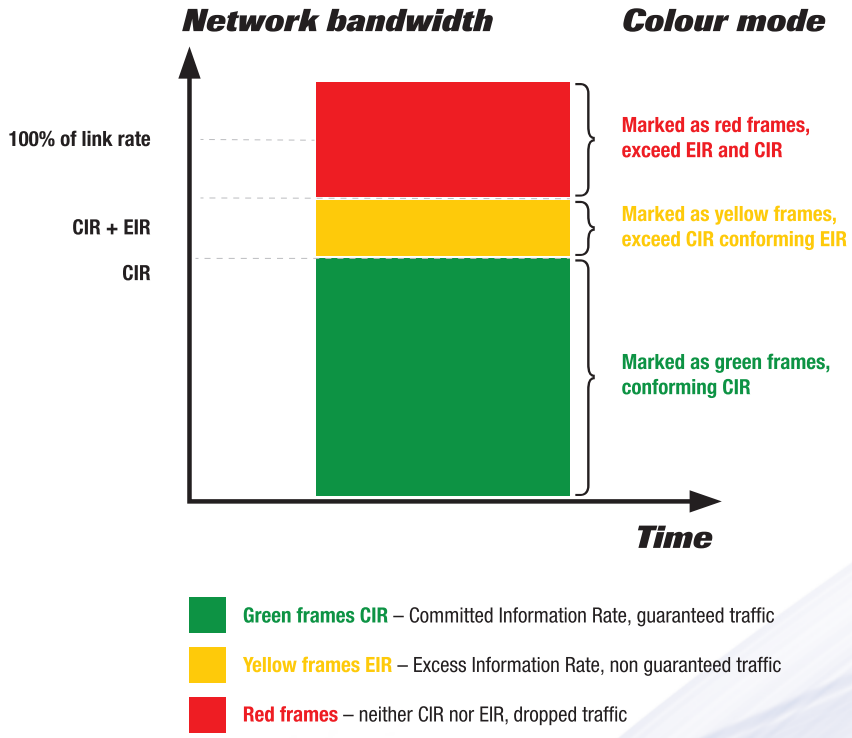

Diagram 6: CIR, EIR and colour mode

4. Buckets for The Traffic Police

Having decided that it is necessary to control the flow of Ethernet traffic into the service provider or wide area network, the question becomes how can this be done in a way that allows Ethernet traffic to 'burst' but at the same time constraining usage to the committed information rate (CIR) plus any agreed 'best-efforts' excess information rate (EIR) agreed in the SLA?

The answer is of course a software algorithm installed into all of the network elements such as edge switches and other NID (network interface devices) that are given the task of being 'traffic policers'.

When the traffic arrives at such a traffic policer, the first thing it does is to establish whether the incoming traffic colour mode (CM) flag is set or not.

If CM is set and the incoming frames are marked green or yellow by the user, the policer will re-assign frames as green or yellow marking. This re-marking depending partly on the colour they were marked by the user (other factors like the user's VLAN C-tags and CoS labels are taken into account if the user did not colour mark). – Importantly however, the service provider also takes into account the incoming information rate.

As a result not all incoming green frames necessarily retain their green marking with in the network. For example if the incoming level of green traffic is significant higher than the SLA permits, not all will continue their journey marked green.

4.1 How the assessment is made

Imagine two buckets. One Green, one Yellow. And into these two buckets are falling a regular supply of green tokens and yellow tokens.

Green tokens each represent one burst of CBS (committed burst size) and drop into the green bucket at the rate of one-eighth of the CIR each second.

Yellow tokens each represent one burst of EBS (excess burst size) and drop into the yellow bucket at the rate of one-eighth of the EIR each second.

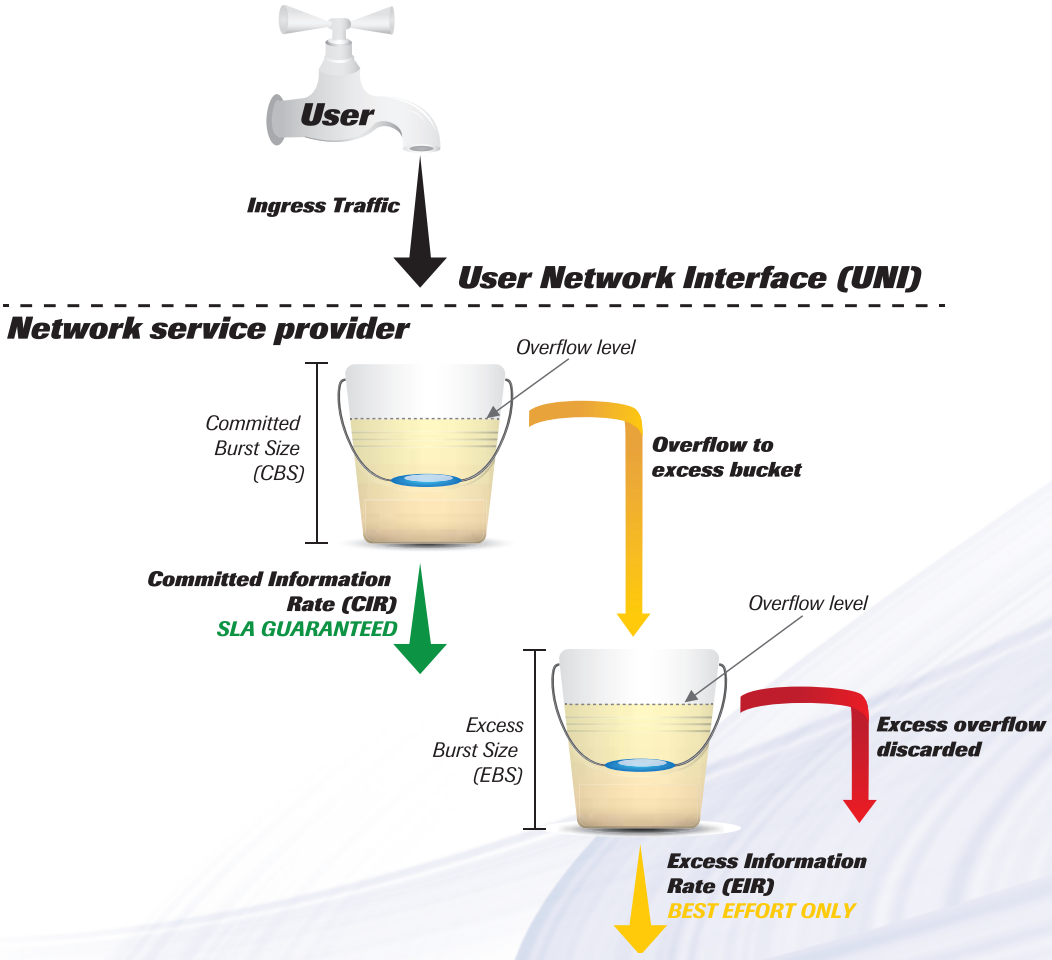

Diagram 7: Bucket theory

4.1.1 Non Colour-Marked Frames Arriving

If the user is not sending colour-aware traffic, all of the arriving traffic is initially handled in the same way as green frames. (As described in the next paragraph).

4.1.2 Green Frames Arriving

As each green frame (or non-colour marked frame) arrives, the policer checks to see if there is a spare green token in the green bucket. If there is, it takes away one token, re-marks the frame as green and passes the frame to the network for guaranteed transmission.

If the green bucket is empty of tokens (because the CIR is used up) the policer next looks to see if the yellow bucket has any tokens. If a token is available, the police takes one token away and the frame is re-classified as yellow before being passed to the network for 'best-efforts' transmission.

If neither the green nor yellow bucket have any tokens, the frame is marked red and discarded.

4.1.3 Yellow Frames Arriving

As each yellow frame arrives, the policer checks against the yellow bucket to see if any tokens are available. If a token is there, the policer takes it, re-marks the frame as yellow and sends it into the network for 'best-efforts' transmission.

If the SLA excludes the use of CF (the coupling flag) any frames arriving for which there is no token available are simply marked red and discarded.

On the other hand, if use of the CF is permitted, the polices will next look to see if any green tokens are available. If one is free, then the placer takes it and marks the frame yellow and passes it to the network.

If neither the yellow bucket nor the green bucket has any tokens available, the frame is marked red and discarded.

5. UniPRO MGig1 three stage testing

Y.1564 defines two steps of testing

1) The Service Configuration Test

2) The Service Performance Test

It also identifies two areas of problems that prevent the testing being commenced.

1) Troubleshoot Ethernet Services

2) Troubleshoot Service Configurations

Sadly most Ethernet Testers ignore these two 'problem areas' which together are often responsible for many hours of "trial and error" fault-finding before the link testing can even begin.

So we've split NetSAM Y.1564 testing into three logical steps – to give Engineers the tools at their fingertips to resolve those configuration issues speedily.

1) Pre-Y.1564 Troubleshoot Ethernet and Service configurations

2) Y.1564 Service Configuration Test

3) Y-1564 Service Performance Test

Let's look at those in detail:

5.1 Pre-Y.1564 - Troubleshoot Ethernet and Service Configurations

It's well known amongst Field Engineers the massive number of hours that are wasted on site simply because links frequently don't actually work when they arrive to test them.

Usually this is the result of a network equipment mis-configuration or a cable mis-patching somewhere along its route. But it is frequently difficult to diagnose whether the problem is due to a test-item parameter setup, a target setup or test path setup.

Frequently this leads to several hours of 'trial and error' manual network configuration checking and cable tracing at the central office to isolate the problem.

So on UniPRO MGig1, we have isolated target check, service path setup and test item parameters setup as three separated pre-testing steps because this can frequently save many hours on-site setup troubleshooting.

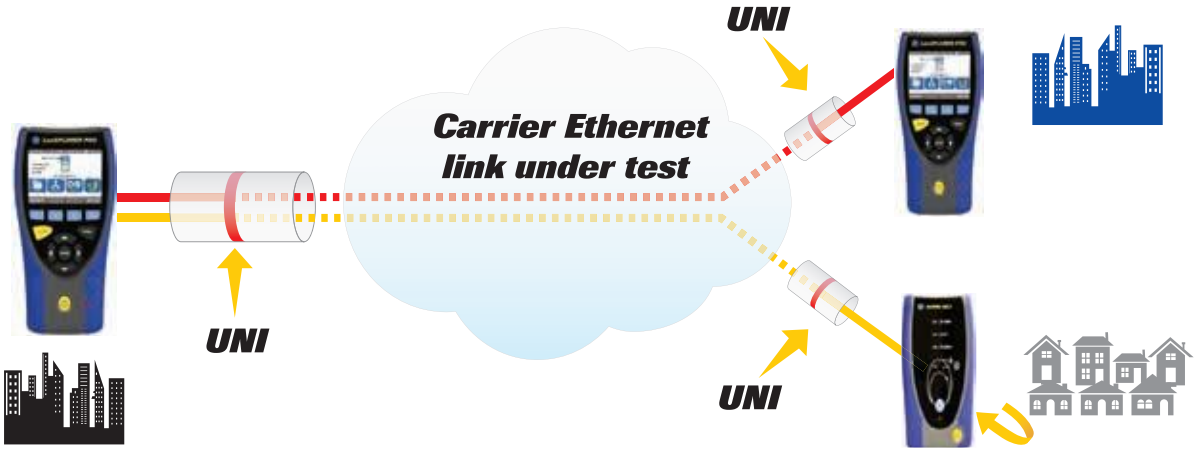

UNI - User Network Interface

Diagram 8: HQ to branch offices tests

5.2 Y.1564 Service Configuration Test

Once the Engineer has confirmed that all service targets are successfully acquired, UniPRO MGig1 commences the service configuration test. Each service or stream is tested in turn to confirm that it is correctly configured and that it properly forwards traffic at the CIR and is capable of the EIR. It also tests that overload traffic is correctly policed.

Diagram 9: UniPRO MGig1 target and service setup on search and configuration

One feature of the bursty nature of Ethernet traffic is that frequently traffic consists of different-sized packets. Therefore it is essential to be able to test the network with mixed frame sizes. Y.1564 defines a mechanism (EMIX) for testing this. EMIX allows representative frame size mixes to be used in a prescribed manner during each test.

The following screen shot shows the setup of EMIX in UniPRO MGig1 tester, highlighting the pre-determined mixed frame sizes A-H, and U being a user defined frame size.

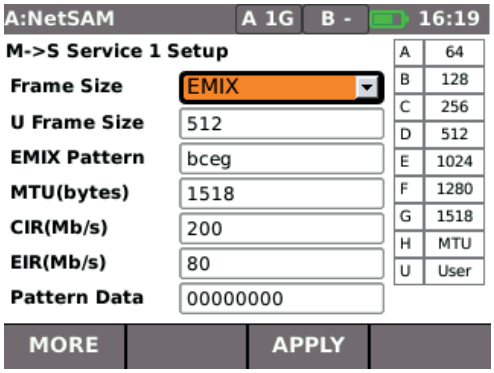

Diagram 10: Emix setup in Y.1564 (NetSAM)

5.2.1 CIR (committed information rate) test

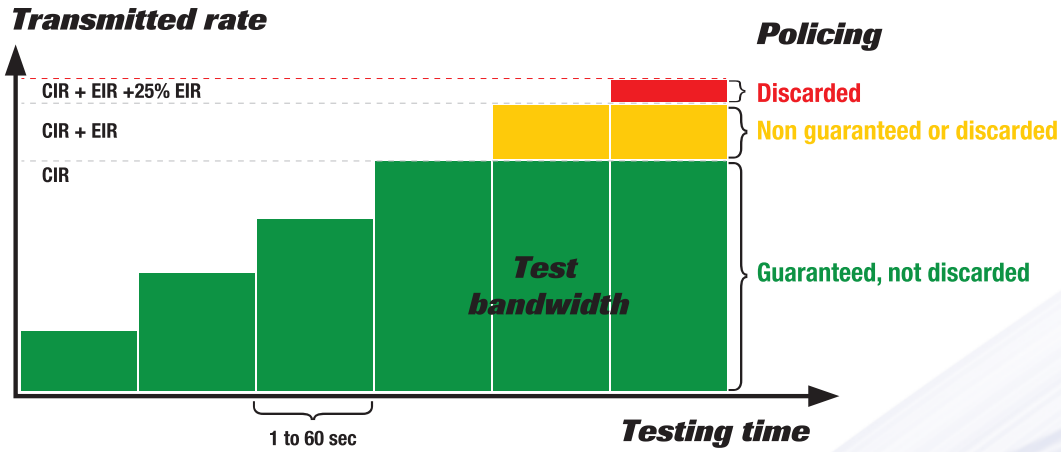

Diagram 11: CIR, EIR bandwidth test and policing

This test checks that the CIR information rate has been properly configured through the network delivered by service providers. Traffic can be presented at the full CIR (Y.1564 Test A1 on UniPRO MGig1) or it can be presented in a step increase format (Y.1564 test A2 on UniPRO MGig1). The traffic can be colour aware or non-colour aware as diagram 11 above shows.

5.2.2 EIR Test - Non Colour Aware

The network is tested to ensure that traffic is successfully carried at the combined CIR plus EIR rate as diagram 11 illustrated. (Y.1564 test B2 on UniPRO MGig1)

5.2.3 EIR Test - Colour Aware

In addition to the CIR+EIR throughput test, UniPRO MGig1 directly tests to ensure that green frames are carried without loss up to the full CIR in the presence of yellow frames. This is to ensure that green frames are correctly prioritised. (Y.1564 test B1 on UniPRO MGig1)

5.2.4 Traffic Policing Test

Traffic is sent above the rate of CIR+EIR (i.e. into the red zone). This is to ensure that the Traffic Policer is operating correctly and discards frames that exceed CIR+EIR (sometimes including 'M' a small safety margin defined by the service provider).

In the case of colour-aware traffic only, UniPRO MGig1 checks that in this overload scenario, yellow frames are discarded not green packets. (Y.1564 tests C1 & C2 on UniPRO MGig1). Diagram 11 gives an idea of the policing test principle up to the red zone where packets are dropped or discarded.

5.3 Y.1564 Service Performance Test

Once all configuration errors have been identified and rectified – and all service configuration traffic tests successfully passed, it is time to test the long-term performance of the Ethernet link.

As a result of configuring each individual service for the configuration tests UniPRO MGig1 now has all of the settings needed for it to be able to simulate 'real-traffic' into the link based on the correct services – up to eight services at the same time as diagram 12 illustrates.

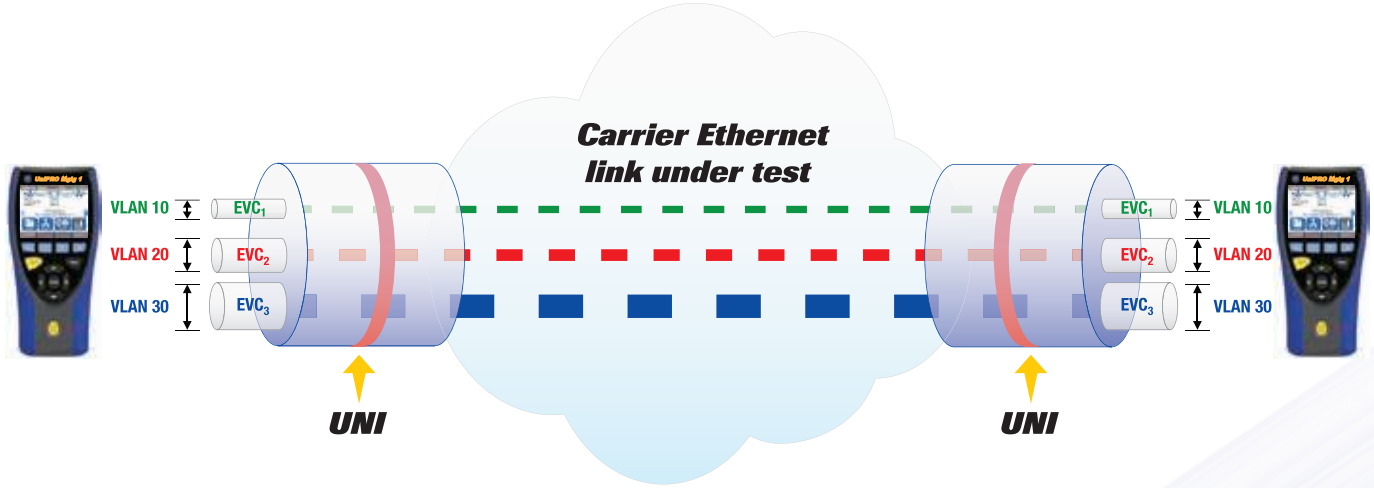

UNI - User Network Interface

Diagram 12: Multiple EVC within same physical interface with different VLANs

UniPRO MGig1 tests all of these services (streams) concurrently at their CIR for periods of 15 minutes (suggested minimum) up to 24 hours (suggested maximum) or beyond with simultaneous measurement of FTD (frame transfer delay), FDV (frame delay variation), IR (information rate), FLR (frame loss ratio) and AVAIL (availability) on each service as illustrated in diagram 13.

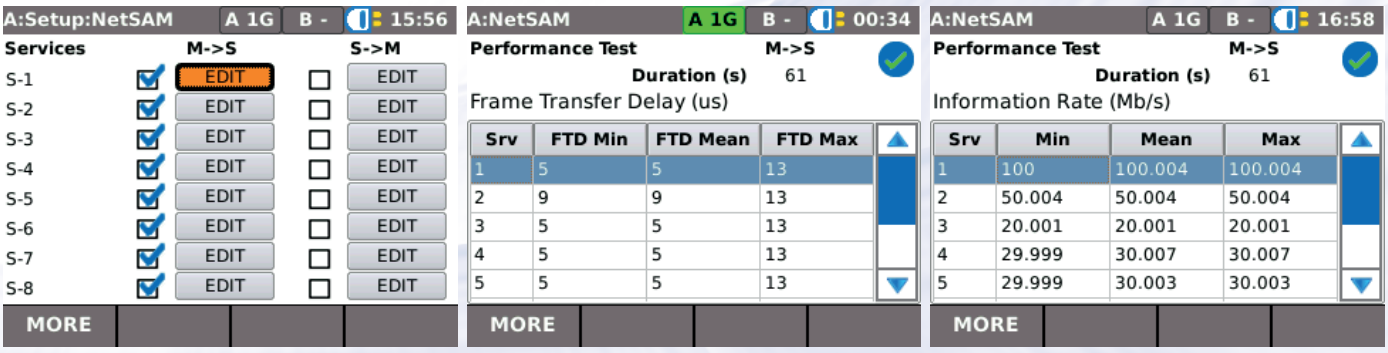

Diagram 13: UniPRO MGig1 Y.1564 (NetSAM) performance test setup and results

6. The Advantages of Y.1564 over RFC2544 in UniPRO MGig1

- Tests multiple services concurrently – not limited to single service test.
- Tests IR (information rate), FDV (frame delay variation or jitter), FTD (frame time delay) and FLRsac (frame loss ratio with reference to the service acceptance criteria)simultaneously rather than one by one
- Far more representative of 'real-world' usage and the demands of QoS on Ethernet Connection.
- Much faster testing – multiple service performance tests in parallel not series
- Verifies configuration of CIR/ EIR/ policing with Colour Mode
- Option of testing specifically to SLA not just to network failure (Versus RFC2544 which only has very time consuming maximum tolerance stress testing and so doesn't allow user to test directly against SLA).
- No time wasted testing permutation after permutation.
- Measured during normal network operation – when SLA should be expected to be met Adds measurement of Frame Delay Variation (jitter)
- Doesn't perform dangerous tests in RFC2544 designed to crash the device (network) under test.
- Tests the network not just a single device.
- Automated sequence test in UniPRO MGig1 NetSAM further improves the efficiency of Y.1564 Configuration Test and Y.1564 Performance Tests

7. Conclusion:

Y.1564 gives a faster, more realistic and more practical test methodology than RFC 2544 for Ethernet Service Activation.

IDEAL Industries Networks' implements Y.1564 testing in the UniPRO MGig1 Carrier-Grade Ethernet tester with its advanced NetSAM (Network Service Activation Method). This combines Y.1564 testing with IDEAL's unique methodology and tools for the speedy resolution of the highly time consuming 'trial and error' troubleshooting of Ethernet and Service Configuration problems which waste thousands of man-hours each year before testing can even commence.

8. Glossary of Y.1564 Related Abbreviations and Acronyms

ATM	Asynchronous Transfer Mode
AVAIL	Availability
CBS	Committed Burst Size
CE	Customer Edge
CF	Coupling Flag
CIR	Committed Information Rate
CM	Colour Mode
CoS	Class of Service
DST	Destination CE
EBS	Excess Burst Size
EIR	Excess Information Rate
EL	Exchange Link
EMIX	Ethernet Mix
ENNI	External Network-to-Network Interface
ETH	Ethernet MAC layer network
ETY	Ethernet physical layer network
EVC	Ethernet Virtual Connection
FCS	Frame Check Sequence
FDV	Frame Delay Variation
FL	Frame Loss
FLR	Frame Loss Ratio
FRE	Frame Reference Event
FTD	Frame Transfer Delay
GPS	Global Positioning System
IMIX	Internet Mix
IP	Internet Protocol
IR	Information Rate
LACP	Link Aggregation Control Protocol
LAN	Local Area Network
MAC	Medium Access Control
MP	Measurement Point
MPLS	MultiProtocol Label Switching
MTU	Maximum Transmission Unit
NID	Network Interface Device
NNI	Network to Network Interface
NS	Network Section
NSE	Network Section Ensemble
OAM	Operation, Administration and Maintenance
OTN	Optical Transport Networ
PDH	Plesiochronous Digital Hierarchy
PE	Provider Edge
QoS	Quality of Service
SAC	Service Acceptance Criteria
SDH	Synchronous Digital Hierarchy

ANY15640513V1.0-UK

IDEAL INDUSTRIES NETWORKS DIVISION
Unit 3, Europa Court, Europa Boulevard, Warrington, Cheshire, WA5 7TN. United Kingdom
Tel: +44 (0)1925 444 446 | **Email:** uksales@idealnwd.com