

302 Enzo Drive San Jose, CA 95138 Tel +1 408-363-8000 Fax +1 408-363-8313 www.sunrisetelecom.com

# The Field Engineer's Challenge - from LAN to Carrier Ethernet

### **Carrier Ethernet**

From its modest beginning in the early 1970s as a way to connect computers in a data center, Ethernet has evolved over the years into the dominant networking technology. Today it is used in a wide range of environments, from home networks connecting a few computers to large enterprise networks connecting technology, it is now increasingly used by telecommunications carriers as an alternative to traditional circuit switched networks over long distances. Carrier Ethernet refers to the extensions necessary in order to allow telecommunications carriers to deploy and provide Ethernet services to their customers.

The following are the main reasons behind the adoption of Ethernet in carrier networks:

**Ethernet is everywhere.** Retail as well as corporate customers are already using Ethernet. By adopting Ethernet in the backbone, the same technology can be used from end-to-end, thus eliminating the need for interface or frame conversion.

**Cost effectiveness.** Due to economies of scale, Ethernet components are inexpensive, thus lowering both the capital and operational expenses.

**Increased flexibility.** Unlike traditional point-to-point circuits that come in discrete speeds (T1/E1, T3/E3, OC-3/STM-1), Ethernet allows subscribers to add bandwidth more incrementally and more quickly. Typically, bandwidth upgrades only involve configuration changes and can take effect in days rather than weeks.

To promote the use of carrier Ethernet and to increase interoperability between vendors, the Metro Ethernet Forum (MEF), an industry consortium, produces technical specifications for Ethernet services in metropolitan areas as well as in the wide area.

In addition to extending Ethernet service to the customers, carriers and service providers are also building and converting their own core networks to Ethernet. For example, more and more mobile phone operators and DSL providers are using Ethernet as the back-haul technology instead of the traditional TDM or ATM network[1].

#### QoS

Over the years, Ethernet has evolved from its original best-effort, CSMA/CD based data communications technology to a QoS (Quality of Service) enabled service capable of supporting multi-media traffic in a converged network. As an example, MEF uses the concept of a bandwidth profile to describe the subscriber's traffic. It specifies the average rate of "committed" and "excess" traffic that a subscriber can

generate into the provider's network, and the associated performance objectives in terms of the delay, frame loss and availability. Thus each traffic profile consists of the following parameters:

- CIR (Committed Information Rate)
- CBS (Committed Burst Size)
- EIR (Excess Information Rate)
- EIS (Excess Information Size)

Service frames up to the committed rate are considered in-profile and are delivered as per the performance objectives. Those sent up to the excess rate are allowed into the provider's network but are considered out-of-profile and are delivered without any service performance objectives. Traffic sent above the excess rate is discarded. Thus subscribers who have different bandwidth profiles will be accorded a different level of service. [2]

## **Field Testing Requirements**

In an increasingly competitive environment, carriers and service providers must be able to provision circuits quickly and cost-effectively, while at the same time ensuring that the service meets or exceeds the Service Level Agreement (SLA). This requires engineers and technicians to be able to accurately and efficiently test circuits in the field. Field testing generally falls into a few categories – basic connectivity testing and service verification; performance testing and QoS testing.

To facilitate field testing, portability of the test equipment becomes an important consideration. Thus instead of chassis based units, hands-held devices are commonly used. In recent years, these units have become widely available and affordable. The SunLite GigE from Sunrise Telecom is one such example and will be used to illustrate the tests described in this white paper.

#### Basic Connectivity Testing and Service Verification

Whether it is a new customer or an existing customer ordering a new circuit, the field technician must ensure that the circuit is operational before it is turned-up. This type of testing usually consists of the following:

- Link check confirm the link is active and successfully negotiated a speed and duplex setting
- DHCP confirm that the CPE can obtain a valid IP address
- PING confirm that there is IP connectivity
- Traceroute confirm that IP traffic is taking the desired route

These tests are also commonly used to troubleshoot an Ethernet network. (*See sidebar on Ethernet Troubleshooting*) Additional tests such as loading a web page using HTTP or downloading a file using FTP are also often performed in order to verify that Internet connectivity is established.

#### **Performance Testing**

Basic connectivity testing is generally sufficient for best-effort service such as residential Internet access which has no implicit performance guarantees. For corporate customers who require services with specific performance objectives, it is common to employ the RFC 2544 tests.

RFC 2544 [3], published in 1999 by the IETF, defines the "Benchmarking Methodology for Network Interconnect Devices". It was originally designed to allow the standardized testing and benchmarking of a single interconnect device such as a router or a switch (known as the DUT or Device Under Test). This methodology has become the de facto standard performed routinely in QA labs and verification labs in order to quantify the performance of network devices. The following diagram shows the classical test setup.



Figure 1 - Standard RFC 2544 Setup for Device Testing

In this setup, a test system is responsible for injecting traffic into the DUT and then analyzing the frames coming out. By correlating the input and output, the test system can measure the transmission characteristics of the DUT.

The RFC 2544 benchmarking methodology defines a series of performance measurements. The terminology and metrics used in the test suite is defined in [4]. The following are the four benchmark tests that are most commonly implemented in an automated test suite<sup>1</sup> :

Throughput – This test measures the highest data rate at which none of the frames are dropped by the DUT. The following frame sizes in bytes should be used during the test – 64, 128, 256, 512, 1024, 1280 and 1518.

Latency – Once the throughput has been determined, this test measures the one-way delay through the DUT at the throughput rate.

Frame loss – This test measures the frame loss rate throughout the entire range of input data rates and frame sizes.

Back-to-back - This test measures the longest frames burst (frames transmitted with

<sup>&</sup>lt;sup>1</sup>In addition to these four, RFC 2544 also defines the system recovery test and the reset test which measure the speed at which a DUT recovers from an overload condition and a device reset. These are not used as much mainly because these metrics are typically not part of the carrier SLA. Furthermore, these two tests are also harder to implement in an automated test.

the minimum inter-frame gaps in between) that the DUT can handle without the loss of any frames.

In recent years, the RFC 2544 test suite has been adapted to measure the performance of systems and networks. Thus instead of measuring the ability of a switch or router to pass frames, we can now use the same methodology to quantify the transmission characteristics of a virtual circuit or a network. The following shows two examples. The first one shows an example of what MEF defines as the Ethernet Line (E-Line) service, which is essentially an emulation of a virtual circuit. The second example is an Ethernet LAN (E-LAN) service, which emulates a multipoint network.



Figure 2 - Testing an Ethernet Virtual Connection



Figure 3 - Testing an Emulated LAN Service

Adapting RFC 2544 from testing a DUT in the lab to a production network in the field requires a few changes:

- The test system used in the test lab connects ports that are physically close to each other. Typically, the same test system transmits traffic from one set of ports and receives traffic from another set of ports. In a network, the two measurement points are geographically separated. Thus two or more test units are used to measure the input and the output. In other words, one unit generates the traffic and the other unit receives the traffic.
- In the lab, the test system is stationary. In the field, the test unit must be portable and can be deployed quickly in various locations.
- RFCs 2544 and 1242 define latency in terms of the one-way delay. In a network where two or more test units are required to perform the test, round-trip delays are measured. Measuring one-way delay would require the test units to be time-synchronized.

The following are some sample results of a RFC2544 throughput test of a Gigabit switch:

TEST MODE : THROUGHPUT TEST

SPEED MODE : 1000 FULL	512 BYTES :
FLOWCONTROL : OFF	MAX RATE : 100%
FLOWCONTROL : OFF	FRAME LOSS : 0
PREAMBLE SIZE : 8	PASS/FAIL : PASS
DATA PATTERN : RANDOM	768 BYTES :
TX RATE(INIT) : 60%	MAX RATE : 100%
TX RATE(FINAL) : 100%	FRAME LOSS : 0
TX RATE(DELTA) : 10%	PASS/FAIL : PASS
TX PACKETS : 10000	1024 BYTES :
64 BYTES :	MAX RATE : 100%
MAX RATE : 100%	FRAME LOSS : 0
FRAME LOSS : 0	PASS/FAIL : PASS
PASS/FAIL : PASS	1280 BYTES :
128 BYTES :	MAX RATE : 100%
MAX RATE : 100%	FRAME LOSS : 0
FRAME LOSS : 0	PASS/FAIL : PASS
PASS/FAIL : PASS	1518 BYTES :
256 BYTES :	MAX RATE : 100%
MAX RATE : 100%	FRAME LOSS : 0
FRAME LOSS : 0	PASS/FAIL : PASS
PASS/FAIL : PASS	
	END

The automated test stepped through all the prescribed frame sizes testing each at the line rate. In this test, the switch was able to pass through wire rate traffic for all frame sizes.

## **Asymmetric Testing**

In field testing, another common test scenario is the broadband access network with asymmetric bandwidth. In ADSL, for example, it is important to test the two directions (upstream and downstream) separately. This is illustrated in the following diagram.



Figure 4 - Testing a DSL Access Network

The following shows the two test configurations for measuring the downstream and upstream performance.



Figure 5 - Measuring Asymmetric Throughput

Furthermore, by inserting the test units at various points of the upstream network, we can determine the throughput of the regional as well as access networks.

## Field Testing of QoS Enabled Network

Today's multi-service networks must support a variety of traffic such as data, voice and video. Corporate clients also have SLAs with stringent performance requirements. For example, some carriers offer a gold, silver and bronze service, each with a different set of service objectives. Thus increasingly, the network infrastructure must be able to differentiate and prioritize traffic. This is usually done using one or more of the following mechanisms. Physical Port. Traffic can be given different priorities simply based on the originating port. In a private network, the physical port may correspond to, say, the voice VLAN where all the IP phones are. In a carrier network, each port may correspond to a different client with a different SLA.

802.1p. Within a switched network, Ethernet frames can be tagged with an additional 802.1p header.<sup>2</sup> This provides QoS at the MAC level. There are three priority bits (referred to as Priority Code Point or PCP) which indicate the priority of the frame. This is sometimes referred to the Class of Service (CoS).

IP TOS or Diffserv. The network can also provide QoS at the IP layer using the TOS (Type of Service) or Diffserv (Differentiated Services). Unlike the 802.1p bits which are usually set by the end system, the TOS and Diffserv bits are usually assigned by the router based on either the IP address of the source or the type of traffic (TCP or UDP port numbers).

VLAN. In an aggregation network, all the subscribers within the same service class can be assigned to a unique VLAN. For example, in a DSL network, a DSLAM can assign all the Gold service subscribers to the same VLAN. The backbone network can then prioritize traffic based on the VLAN ID. This VLAN tag is usually referred to as the S-VID or the provider VLAN. In order to preserve the customer's own VLAN designation (C-VID), Q-in-Q is often used.

#### **Testing Prioritization**

Regardless of the prioritization scheme, one must then be able to confirm that indeed the network is prioritizing traffic in accordance with the service level objectives. Testing QoS requires the ability to inject streams of traffic that will be accorded different treatment by the network and then measure their respective performance. The following diagrams illustrate a few scenarios.



Figure 6 - Prioritization Based on Physical Port

In this scenario, prioritization is done based on the physical port. For example, all traffic originating from port 1 of the ingress switch might be given a higher priority ("Gold" service). On the receiving side, by monitoring the two traffic streams, we can confirm that when there is congestion (e.g., when both streams are transmitting at over 50% of the line rate), more packets from stream 2 will be dropped compared to stream 1.

<sup>&</sup>lt;sup>2</sup>The 802.1p header is an 802.1Q header with the VLAN ID set to 0.



Figure 7 - Prioritization Based on 802.1p

In this test setup, 3 streams of traffic are injected into the network each with a different PCP (0 being the lowest and 7 the highest priority). On the receiving side, the test set monitors each stream to look for impairments such as packet loss. In order to verify that traffic is being prioritized, there must be congestion on the output port, thus forcing the network to decide which packets to drop. This congestion can be simulated if necessary by injecting some additional background traffic from another port.<sup>3</sup>

The next two scenarios illustrate similar setups for verifying IP prioritization as well as prioritization based on VLAN.



Figure 8 - Prioritization Based on TOS or DiffServ

<sup>&</sup>lt;sup>3</sup>These few diagrams show a separate port and a separate test unit used to generate the background traffic. This is only necessary if the ingress and egress ports have the same speed and thus additional traffic is needed to artificially create a congestion condition.



Figure 9 - Prioritization Based on VLAN

The last example shows an example of Q-in-Q. Each test packet has two 802.1Q tags, the outer tag is the provider tag and the inner tag is the customer tag. In this example, a provider tag of 100 might indicate a "Gold" service and 101 might indicate a "Silver" service. Thus the test will determine if each provider tag indeed accords each class with a different service level.

Using the SunLite GigE, one can construct multiple simultaneous packet streams and perform a "BERT Throughput" test. In the example below, four streams are enabled that correspond to four different bandwidth profiles. Each packet is configured with 2 VLAN tags (Q-in-Q). The outer tag (S-VID) has a VLAN ID of 100 whereas the inner tag (C-VID) has a VLAN ID of 200.

To diagnose these problems quickly, the field personnel needs a number of troubleshooting tools that are now packaged in hand-held test units. The following are some of these common tools:

Cable tester – This function checks for any cable faults such as open and short circuits. Using TDR (Time Domain Reflectometry), the tester can also determine the length of the connected cable.

Port identifier – This function activates and deactivates the Ethernet port thus causing the port LED on the switch to turn on and off periodically. This is useful in identifying the switch port connected to a device.

Link status check - When the test unit is connected to a switch port, this function displays the link status such as the speed, duplex and flow control.

DHCP – This allows a test unit to confirm that an IP address can be acquired successfully from the DHCP server.

Device scan - This function scans a range of IP addresses looking for active devices.

PING – This standard function probes an IP address using ICMP. It returns the round-trip delay between the test unit and the target device.

Traceroute - This standard function traces the IP path to a target address identifying all the intermediate routers.

## Troubleshooting Ethernet

What to do when the network does not work? Here is where troubleshooting tools come in. In Ethernet, there are a number of common problems that can be identified quickly with the right tools. These include the following:

- Cabling error. The LAN cable is connected to the wrong port.
- Cable fault. There is a problem with the connector resulting in an open or short circuit.
- Speed or duplex mismatch. The NIC is configured for a speed/duplex setting that is incompatible with that of the switch port.
- DHCP error. Stations are unable to obtain an IP address from the DHCP server.
- Duplicate IP address. Multiple devices are configured statically with the same IP address.
- Routing problem. Stations are unable to reach destinations outside the local VLAN either because the default gateway is down or misconfigured.

I		Stream	ı Table	-q.	
	#	S	Structure		
	✓ 1	MAG	MAC-VLAN2-IP		
	<b>√</b> 2	MAC	MAC-VLAN2-IP		
	<b>√</b> 3	MAG	MAC-VLAN2-IP		
	✔ 4	MAG	MAC-VLAN2-IP		
	5		MAC		
	6		MAC		
	7	MAC			
		MAC			
				•	
	On/Off	Edit Menu Start			

Figure 10 - Enable four packet streams.

🔲 Stre	am Conf	figurati	ion #1 🚃 🔒
Frame Stri	ucture —— VL/	AN [	MPLS
IP IP	🗖 ΤርΙ	P/UDP [	PAYLOAD
Frame S	ize:	70	Bytes
[Traffic Sha	ape ———		
🛛 🔘 Const	ant 🔿	Ramp (	🔵 Burst 🛛
Mode: (	۵% С	)Kbps (	🔵 Mbps 👘
Rate:		10.00	%
On/Off	Edit	OK	Next

Figure 11 - Each packet has a MAC, VLAN and IP header. This stream will transmit at 10% line rate.

	VLAN		
✔ #1. TPID :	9100	CFI :	0
VLAN ID :	100	Pri. :	0
💌 #2. TPID :	8100	CFI :	0
VLAN ID:	200	Pri. :	0
🔲 #3. TPID :	0000	CFI :	0
VLAN ID:	1	Pri. :	0
Auto Fill OK			

Figure 12 - There is an outer tag and inner tag.

After configuring the streams, traffic was turned on. The following screens show the test results:

🖸 Statistics View-Stream 1 🚅 🔒			
Rx fps		139450	▲
Rx Frames	3	1277696	
Rx Utilizati	ion(%)	10.04 %	
Rx Line Ra	Rx Line Rate		ops 🖵
Rx Data R	Rx Data Rate		3
Tx fps		139450	
Tx Frames	Tx Frames		
Tx Utilization(%)		10.04 %	
Tx Line Rate		100.40 Mbps	
Tx Data Rate		78.0 Mbps 📃 💌	
Prev	Next	Menu	PgDn

Figure 13 - Every transmitted frame was echoed back by the responder at the other end - therefore no frame loss.

🔋 Sta	tistics Vi	ew-Strea	ım 1 🔜 🔋
Loss Fram	es	0	
RX Pause	Frames	0	
VLAN 1 Fr	ames	1277696	
VLAN 2 Fr	VLAN 2 Frames		
VLAN 3 Fr	ames	0	
Inter-Fram	e Delay Min.	1872 ns	
Inter-Fram	e Delay Max	. 2784 ns	
Inter-Fram	e Delay Avg.	2550 ns	
Inter-Fram	e Delay Var.	912 ns	
Pat Sync I	Error Frames	0	-
Prev	Next	Menu	PgDn

Figure 14 – This test also measures the frame delay statistics.

# Jitter or Packet Delay Variation

Real-time traffic such as IP telephony or IP video is sensitive not only to the amount of delay but also whether the amount of delay is relatively constant or varying. This measure is referred to as the packet jitter or packet delay variation<sup>4</sup>. To complicate matters, there are many ways of calculating delay variations and they can yield different results. The following are a few commonly used definitions:

- In RFC 3393, the one-way IPDV (IP delay variation) is defined as the difference between the delays experienced by two consecutive packets from the same stream.<sup>5</sup>
- In RFC 4689, jitter is defined as the absolute value of the IPDV.<sup>6</sup>
- MEF defines frame jitter as the difference between the maximum frame delay and the minimum frame delay within a group of samples.<sup>7</sup>
- In ITU Y.1540, the 2-point Packet Delay Variation is defined as the difference between the packet transfer delay of the packet and a defined reference IP packet transfer delay.<sup>8</sup>

## **Triple Play**

For real-time applications such as voice and video, network transmission attributes such as packet delay, jitter and packet loss can adversely affect the quality of the reception. For example, for IP telephony, the following are commonly recommended values in order to maintain acceptable voice quality:

- Packet delay: Less than 150 ms one-way (less than 100 ms one-way for tollquality voice).
- Jitter: Less than 50 ms
- Packet loss: Less than 1%

For triple play, one must be able to generate multiple data streams and verify that voice and video traffic is transmitted with the required performance. In this test, the different streams can represent data, voice and video. How the network treats each stream depends on the network configuration. In some networks, all IP phones are assigned to a separate voice VLAN. In others, the IP phones are programmed to tag all traffic with higher priority using 802.1p. Thus in designing a test, one must take into consideration the configuration of the network infrastructure in order to effectively simulate the traffic mix. The following diagram illustrates a typical setup.



Figure 15 - Testing Triple Play

<sup>&</sup>lt;sup>4</sup>Packet jitter or delay variation measures the change in packet delay. This is not to be confused with video jitter which happens when the picture "jumps". Video jitter can be the result of packet loss and excessive packet jitter.

## Future Requirements - Ethernet Service OAM

As carrier Ethernet becomes more widely deployed, an area which is evolving rapidly is Ethernet Service OAM (Operation, Administration and Maintenance). The objective is to provide carriers with a comprehensive set of tools to monitor the availability and performance of Ethernet services. There are a number of related efforts in progress:

- IEEE 802.3-2005 (previously 802.3ah) Ethernet Link OAM (Ethernet in the First Mile)
- IEEE 802.1ag Ethernet Service OAM
- ITU-T Y.1731 OAM Functions and Mechanisms for Ethernet Based Networks
- MEF 17 Service OAM Requirements and Framework

It is expected that as these standards mature, the prescribed tests (such as connectivity check, loopback, link trace and so on) will be implemented in the field test sets.

## Conclusions

Ethernet has become the ubiquitous networking technology. Carriers and service providers are under competitive pressure to provision service quickly and cost-effectively. One of the challenges of the field technician is to verify the integrity of the service which is becoming more complex. Basic testing uses simple IP tools such as DHCP, PING and traceroute to verify connectivity. Increasingly, Ethernet service must also be measured against the service level. RFC 2544 is the test suite that can be used to determine the throughput, latency, frame loss and back-to-back. In a QoS enabled network, traffic may be prioritized depending on the physical port, 802.1p, TOS or DiffServ, or VLAN tag. To verify proper QoS handling, the field engineer must be able to generate multiple packet streams with different characteristics and measure the performance of each stream separately. For example, in a network that uses Q-in-Q, the packets must be configured with at least two VLAN headers to allow the aggregation switches to prioritize traffic using the outer tag. Finally, to ensure that the network can handle real-time traffic properly, the packet delay variation (which is not part of the RFC 2544 test) must also be included.

# References

<sup>1</sup>TR-101, "Migration to Ethernet Based DSL Aggregation", Broadband Forum, April 2006.

<sup>2</sup> MEF 10.1, "Ethernet Services Attributes Phase 2", Metro Ethernet Forum, November 2006.

<sup>3</sup> Bradner, S. and McQuaid, J., "RFC 2544 – Benchmarking Methodology for Network Interconnect Devices", IETF, March 1999.

<sup>4</sup> Bradner, S., "RFC 1242 – Benchmarking Terminology for Network Interconnection Devices", IETF, July 1991.

<sup>5</sup> Demichelis, C. and Chimento, P., "RFC 3393 – IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", IETF, November 2002.

<sup>6</sup> Poretsky, S., Perser, J., Erramilli, S., and Khurana, S., "RFC 4689 – Terminology for Benchmarking Network-layer Traffic Control Mechanisms", IETF, October 2006.

<sup>7</sup> MEF 6.1, "Ethernet Services Definition – Phase 2", MEF, April 2008.

<sup>8</sup> ITU-T Y.1540 "Internet Protocol Data Communication Service – IP Packet Transfer and Availability Performance Parameters", December 2002.

For more information, please call: 1-800-701-5208 or 1-408-363-8000 or visit www.sunrisetelecom.com

Our e-mail: info@sunrisetelecom.com