THE IMPORTANCE OF TESTING TCP PERFORMANCE IN CARRIER ETHERNET NETWORKS

159

APPLICATION NOTE

Bruno Giguère, Member of Technical Staff, Transport & Datacom Business Unit, EXFO

As end-users are migrating from their legacy access architecture (Frame Relay, ATM or private line) to accommodate new applications like Voice-over-IP (VoIP), Customer Relationship Management (CRM), and Enterprise Resource Planning (ERP), service providers are now faced with the obligation to supply stringent service-level agreements (SLAs).

Historically, SLAs have been defined by these parameters, for both networks and applications:

Availability

- Uptime/downtime
- Mean-time-to-repair (MTTR)
- Protection switching

Performance

- Performance availability (throughput)
- Link burstability
- Service integrity (frame-loss rate)
- Transmission delay (latency)
- Frame-delay variation (packet jitter)

Although these parameters are great to characterize and define an SLA, they only cover the network performance up to the IP layer of a network (see Figure 1). With these parameters, service providers and end-users have an idea that the network is capable of transporting frames. However, they still will not know what level of performance they should expect from their mission-critical applications. How can service providers make sure that the end-user's most important application can make use of the full bandwidth of the newly installed Ethernet services? This application note will clarify these issues and discuss the role of Transport Control Protocol (TCP) in the transmission process.

Example of SLA parameters for business VPN services	
Latency	10 to 45-55 ms (one-way latency)
Jitter	5-10 ms
Packet loss	0.05%
Availability	99.98%
Class of service (CoS)	2 levels
Mean-time-to-repair (MTTR)	4 hours
Protection switching	< 50 ms

Table 1: *Sample values for different SLA parameters*

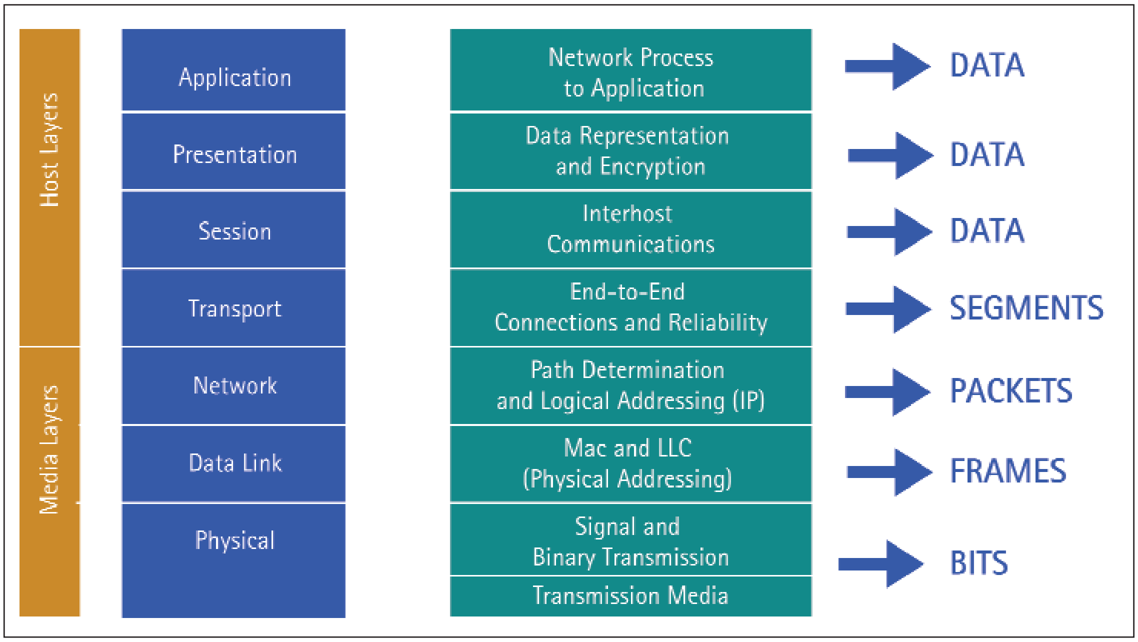

Figure 1. *OSI reference model and naming convention*

Communication-Network Protocols

To successfully exchange information, two networked devices need a series of protocols allowing the applications to communicate. This protocol suite, called TCP/IP, consists of a series of stacked protocols, layered one on top of the other. Each layer has a specific function and will provide services to the upper layers. The top layer – the Application Layer – uses the layers underneath it to communicate with another end-device. One of the most important layers in this process is the Transport Layer, as it is the entry point to the Host category of layers and it is responsible for the end-to-end connections; i.e., the Transport Layer ensures that the data segments are transfered from the network toward the application.

Applications use two types of transport protocols to communicate between all of their locations; User Datagram Protocol (UDP) and Transport Control Protocol (TCP). These protocols are part of the TCP/IP protocol stack that provides the complete architecture to exchange data between two networked devices. Depending on the application being run, the transport protocol will be different. If the application is used in real time and the loss of information is not critical, UDP will be used, as this protocol is simple, efficient and faster than TCP. For applications such as IPTV, VoIP or on-line gaming, it is the perfect protocol. On the down side, this protocol does not provide the reliability and ordering guarantees that TCP does. Information may arrive out of order or go missing without notice.

When an application needs a reliable and efficient delivery connection between two networked devices, TCP should be used. Examples of such applications are the Internet, e-mail, Customer Relationship Management (CRM), Enterprise Resource Planning (ERP) and file transfer.

TCP, in fact, ensures an entire process that consists of several chronological actions:

1. Establishes a connection between the two end-points
2. Manages the exchange of information, making sure that packets are delivered without errors and retransmitting them, if necessary
3. Reorders and removes duplicate segments received
4. Provides flow control *between the two end points*
5. Disconnects from the end-device once the exchange of information is completed

In addition, TCP also differentiates data from simultaneous applications (e.g., Web access and e-mail server) running on the same networked device.

As TCP is more sophisticated than UDP, it has multiple parameters that can be configured to optimize its utilization. Unfortunately, the default values used in different implementations can diminish the performance of the transmission across a network and can create situations that cause service providers and end-users to debate on the capability of a network to transmit TCP traffic.

The other layers underneath the Transport Layer are the Network, Data Link and Physical layers, which are also grouped together and called the Media Layers (see Figure 1). Although they are essential to deliver information, they are outside the scope of this application note.

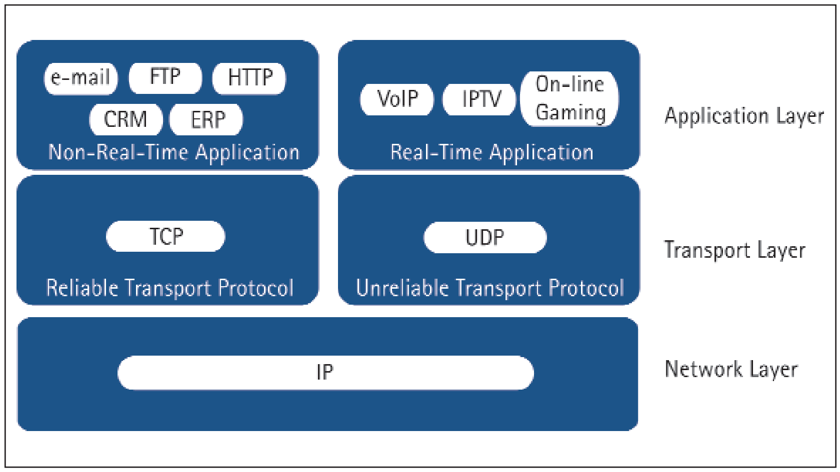

Figure 2. *Applications Running Over TCP and IP Protocols*

Non-Real-Time Applications on Multimegabit Networks

As service providers are deploying new high-bandwidth access technology to replace their legacy networks, end-users expect their current mission-critical application to perform better. Going from a 1.5, 2 or 45 Mbit/s access to a full 100 Mbit/s should enable them to fill this new pipe with the maximum traffic available. Unfortunately, this is not always the case.

As described earlier, applications running over TCP require the traffic to get transported to the other end without error and in order. Because of this, TCP will add overhead and delay in the transmission, as it needs to make sure that the transmitted information is received correctly at the other end.

Without going into all the details of the TCP protocol, there are certain parameters that are configured in the protocol that will influence the capability of a device to transfer information efficiently across the network. These parameters are: the size of the transmission window; the size of segment being transmitted; and retransmission time-out. Parameters external to TCP will also affect its performance. Round-trip delay and frame loss are the most important players in the operation of a TCP link.

There are also other factors that can affect the capability of an application to transmit data across a network. Factors like the application being used, the type of TCP/IP stack and the performance of the computers/servers running these applications will have an impact on the performance of the applications.

From a purely theoretical perspective, the maximum possible throughput from TCP can be defined by the following equation (also known as the bandwidth-delay product):

$$\text{Capacity (bits)} = \text{bandwidth (bits/sec)} \times \text{round-trip time (seconds)}$$

For example, if the round-trip time is 40 ms, end-users might find some interesting limitations in their current TCP implementation, depending on the circuit rates (see Table 2).

Bandwidth-delay product for different circuits based on a 40 ms round-trip time			
Circuit rate	Payload rate (Mbit/s)	Capacity (in bits)	Capacity (in bytes)
DS1 (1.5M)	1.536	61 440	7680
E1 (2M)	1.984	79 360	9920
DS3 (45M)	44.21	1 768 400	221 050
100BASE-T	100	4 000 000	500 000
OC-3/STM-1 (155M)	149.76	5 990 400	748 800
OC-12/STM-4 (622M)	599.04	23 961 600	2 995 200
1000BASE-T	1000	40 000 000	5 000 000
OC-48/STM-16 (2.5G)	2396.16	95 846 400	11 980 800
OC-192/STM-64 (10G)	9584.64	383 385 600	47 923 200
10GBASE-SW (WAN)	9584.64	383 385 600	47 923 200
10GBASE-SR (LAN)	10 000	400 000 000	50 000 000

Table 2: *Bandwidth-delay product for different circuit capacity*

The column of interest is the ***Capacity (in bytes)***. This theoretical value provides the maximum number of bytes in the system at any time so the circuit is filled at the maximum and that TCP can resend any dropped or errored segment. In a standard TCP implementation, the maximum allowable TCP window is 65 535 bytes; this means that at a rate of 45 Mbit/s and more, with a round-trip time of 40 ms, a server running normal TCP cannot fill the circuit at 100%.

So unless the TCP implementation used can extend its window size to more than 65 535 bytes, the end-user will never be able to transmit a single stream of TCP data at more than 13.1 Mbit/s for a round-trip time of 40 ms. As mentioned before, this is the theoretical value; unfortunately, the network might drop frames along the way, making it unrealistic to achieve such a throughput.

Advantage of Testing TCP Performance

As shown above, TCP performance across a network is dependent on multiple parameters, so what are the options for a service provider? Historically, service providers have used a test methodology based on RFC 2544 (Benchmarking Methodology for Network Interconnect Devices). This methodology provides a great way to assess the performance of a network if the applications running on it are UDP-based. The throughput, frame loss, burstability and latency tests provide a thorough snapshot of the quality of the network and are at the base of all current SLAs. That being said, if the applications are running on TCP, this methodology will only provide a general idea of how good the network is, but it cannot assess the quality of service an end-user will experience.

An end-user will always measure his TCP performance according to an end-to-end scheme. They will either base their tests on the bandwidth statistics provided by the computers/servers that are running their applications or use software to emulate TCP traffic. This last process will lead them to the conclusion that the service provider is at fault because their measurement shows that the maximum throughput they get is nowhere near the bandwidth available as it is supposed to be provisioned. These software tools are running on computers and operating systems. Unfortunately, all operating systems are not created equally. Some have their TCP/IP stack locked and use the basic windowing scheme as defined for TCP, which is 65 535 bytes. These software tools are also as good as the computers they are running on. Lack of performance from the computer will reflect in lack of performance in the measurement, and therefore will provide a false view of the performance of the network.

Some test methodologies also use multiple TCP sessions to fill the bandwidth because of the TCP window limitation. Although this methodology will show that it is possible to fill a service provider's circuit with TCP traffic, it will not demonstrate that a single application can achieve this. There are also issues with having multiple TCP sessions running at the same time. Each of them is trying to send maximum traffic into the test circuit. Since they are sharing the same bandwidth, they will sometimes send traffic at maximum allowed rate, but sometimes they will be in congestion mode and leave the bandwidth to another session. A measurement could be derived from this methodology, but since this is an attempt at multiple averaged measurements, they might not always be repeatable.

As end-users will test their network with a wide variety of tools, service providers will need a way to prove beyond any doubt that their network behaves as designed, or end users will deem their SLA as unmet.

By having a tool that can send TCP traffic across the network from each demarcation point in the network and should that tool be based on a neutral TCP/IP implementation, service providers would have the perfect methodology to actually prove that their network is not at fault. After that, they can provide the parameters measured to the end-users so they can try to replicate the test results from their test methodology.

Advantage of EXFO's TCP Testing Methodology

Once a service provider recognized the necessity of testing TCP performance in their networks, it becomes important to find a test tool that has the required feature set. Like any test equipment or methodology, the chosen instrument must have the ease of use, accuracy and repeatability.

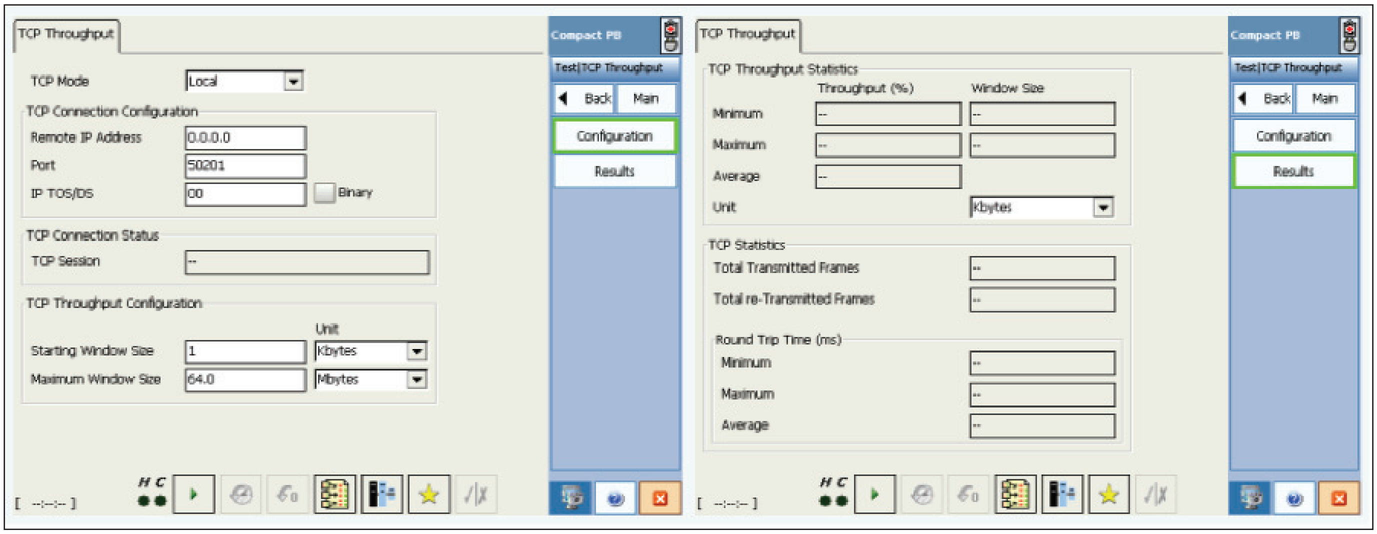

Figure 3. *FTB-8510/8510B TCP Throughput user interface as found in FTB-200 platform*

The TCP Throughput feature of the Packet Blazer FTB-8510/8510B Ethernet module will help service providers establish that the Ethernet services being delivered can provide the quality of service that an end-user expects from a TCP application perspective. With the ease of use found in the graphical user interface of the FTB-200 or FTB-400 platform, technicians and network professionals alike will test TCP performance in no time.

Since the TCP Throughput feature is implemented in the hardware, the Packet Blazer will always provide an accurate measurement because it does not rely on the implementation of any communication stack found in PC operating systems or servers.

The methodology used by EXFO to deliver TCP Throughput measurements is based on the TCP window scale option as described in RFC 1323; a single stream can then be used to provide the TCP throughput measurement. Therefore, it fills a circuit at full bandwidth with TCP traffic when the round-trip time or the transmission bandwidth is too large for standard TCP implementation.

This method also provides additional ease of use versus software-based solutions as the professional performing the test does not need to figure out how many sessions are required for the configuration and which TCP port to use on each. From a results perspective, the user does not need to average multiple test results to verify if the circuit is capable of transporting TCP application traffic. Furthermore, having only one TCP test session provide repeatability. If the network conditions (frame loss, round-trip time, etc.) are the same, the TCP throughput test should provide the same results.

Conclusion

As demonstrated previously, the TCP protocol is used by all non-real time application to deliver mission-critical information end-to-end in a network. Because the TCP protocol needs to acknowledge that the information got transmitted without any errors, it has built-in functionality that can limit its capability to perform in high-latency or high-bandwidth networks.

As all vital applications are not created equal and run on a wide range of computers/servers, the TCP/IP implementation and configuration will vary a lot between each end-customer. Since most end-customers measure the performance of a network with servers' statistics or software tools running on platforms, how can a service provider prove that their network is running according to specification?

The present test methodology, based on RFC 2544, will provide part of the test solution, as it covers the capability of a network to transfer information between two end-points in a real-time mode. Unfortunately, it does not take TCP's reality into account.

The only solution is to test TCP performance with a tool that can prove without any doubt that with the proper TCP configuration and because of the parameters of the network (round-trip time and frame loss), it is possible to get the maximum TCP traffic across the network.

EXFO Corporate Headquarters > 400 Godin Avenue, Quebec City (Quebec) G1M 2K2 CANADA | Tel.: 1 418 683-0211 | Fax: 1 418 683-2170 | info@EXFO.com

Toll-free: 1 800 663-3936 (USA and Canada) | www.EXFO.com

EXFO America	3701 Plano Parkway, Suite 160	Plano, TX 75075 USA	Tel.: 1 800 663-3936	Fax: 1 972 836-0164
EXFO Europe	Omega Enterprise Park, Electron Way	Chandlers Ford, Hampshire S053 4SE ENGLAND	Tel.: +44 2380 246810	Fax: +44 2380 246801
EXFO Asia	151 Chin Swee Road, #03-29 Manhattan House	SINGAPORE 169876	Tel.: +65 6333 8241	Fax: +65 6333 8242
EXFO China	No.88 Fuhua, First Road Central Tower, Room 801, Futian District	Shenzhen 518048, CHINA	Tel.: +86 (755) 8203 2300	Fax: +86 (755) 8203 2306
	Beijing New Century Hotel Office Tower, Room 1754-1755 No. 6 Southern Capital Gym Road	Beijing 100044 P. R. CHINA	Tel.: +86 (10) 6849 2738	Fax: +86 (10) 6849 2662

EXFO EXPERTISE REACHING OUT

ISO 9001

Printed in Canada 07/03