The Importance of Testing TCP Performance in Carrier Ethernet Networks

By Bruno Giguere, Advisor, CTO Office, EXFO

Cloud computing, virtualization, globalization – these terms regularly appear in communications from enterprise customers to service providers, but what do they actually mean to service providers? The answer is usually broadband access. However, customers are now seeking high-performance, Ethernet-based services to connect to the outside world. So what is the impact on service providers? They are obligated to deliver services with stringent service level agreements (SLA).

Historically SLAs, for both services and applications, have been defined by the following parameters:

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)

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
MTTR	4 hours
Protection switching	< 50 ms

Table 1: Sample values for different SLA parameters

Although these parameters characterize and define an SLA, they only cover network performance up to the IP layer (see Figure 1). They may allow service providers and end-users to validate that the network is capable of transporting frames. However, they still won't 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 makes full use of the bandwidth of the newly installed Carrier Ethernet services? This application note will clarify these issues and discuss the role of Transport Control Protocol (TCP) in the transmission process.

COMMUNICATION NETWORK PROTOCOLS

To successfully communicate application information, network devices must use a series of protocols. TCP/IP is a protocol suite consisting of a series of stacked protocols, i.e., layered one on top of the other. Each layer has a specific function and provides services to the layers above it. The top layer, the Application layer, uses the layers underneath it to communicate with other devices. One of the most important layers in this protocol is the Transport layer, because it serves as the entry point for the Host category of layers. It is also responsible for end-to-end connections; specifically, the Transport layer ensures that the data segments are transported from the network toward the application.

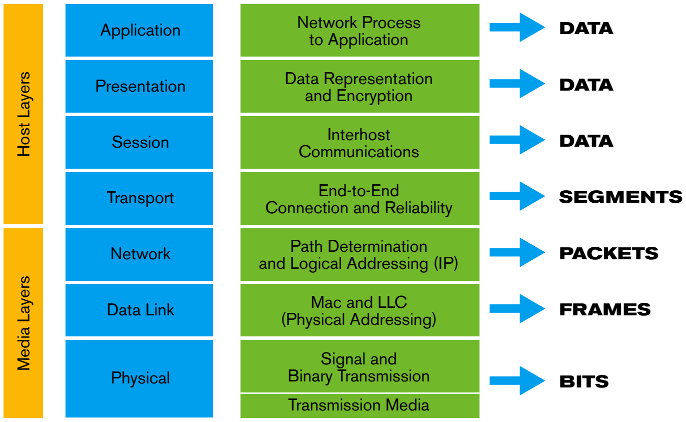

Figure 1. OSI reference model and nomenclature

Two types of transport protocols are used to communicate between all end-user 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 through which data between two networked devices is exchanged. Depending on the application being run, the transport protocol will be different.

For real-time applications where loss of information is not critical, UDP will be used because this protocol is simple, efficient and faster than TCP. For applications such as IPTV, VoIP or on-line gaming, it is the perfect protocol. The only downside is that this protocol does not provide the same reliability and ordering guarantees that TCP does. Information may arrive out of order or go missing without notice.

When applications need a reliable and efficient delivery connection, TCP must be used. Examples of such applications are the Internet, e-mail, Customer Relationship Management (CRM), Enterprise Resource Planning (ERP) and file transfers.

TCP oversees an entire process that consists of several chronological actions:

1. Establish a connection between two end-points.
2. Manage the exchange of information, making sure that packets are delivered without error, and retransmitting them, if necessary.
3. Reorder and remove all duplicate segments received.
4. Provide flow control.
5. Disconnect from the device once the exchange of information is complete.

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

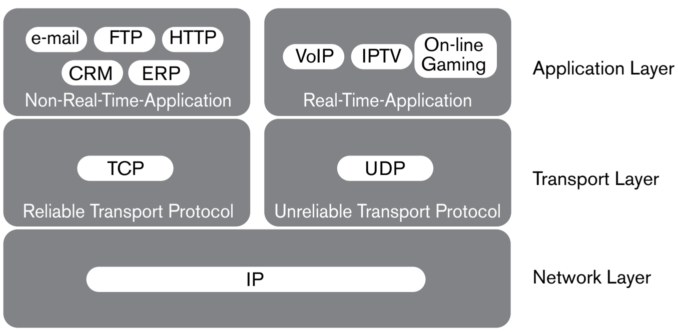

Figure 2. Applications running over TCP and IP protocols

Since TCP is more sophisticated than UDP, it can be configured for optimal use by setting multiple parameters. Unfortunately, the default settings 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 capacity of a network to transmit TCP traffic.

Under the Transport layer are the Network, Data Link and Physical layers. Grouped together, they are referred to as the Media Layers (see Figure 1). Although they are essential to information delivery, they are outside the scope of this application note.

NON-REAL-TIME APPLICATIONS ON MULTIMEGABIT NETWORKS

While service providers deploy new high-bandwidth access technology to connect to the outside world, end-users expect their current mission-critical applications 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 stated earlier, TCP will transport application data from one end to the other without error and in perfect order, but it will add overhead to do so and there will be a delay in the transmission.

Without going into great detail, certain TCP protocol parameters can be configured to influence a device's capacity to transfer information efficiently across the network. These parameters are: transmission window size, segment size and retransmission timeout. However, parameters that are external to TCP will also affect its performance, namely roundtrip delay and frame loss. These are the most important factors in TCP link operation.

Other factors will also affect an application's ability to transmit data across a network. These factors include the application being used, the type of TCP/IP stack and the performance of the computers/servers running the applications.

From a purely theoretical perspective, the following equation represents the maximum possible throughput of TCP (also known as the bandwidth-delay product):

Capacity (bits) = bandwidth (bits/sec) x roundtrip time (seconds)

For example, if the roundtrip time is 40 ms, end-users might discover limitations in their current TCP implementation, as per the circuit rates (see Table 2).

Bandwidth-delay product for different circuits based on a 40 ms roundtrip time			
Circuit rate	**Payload rate (Mbit/s)**	**Capacity (in bits)**	**Capacity (in bytes)**
DS1 (1.5M)	1.536	61440	7680
E1 (2M)	1.984	79360	9920
DS3 (45M)	44.21	1768400	221050
100BASE-T	100	4000000	500000
OC-3/STM-1 (155M)	149.76	5990400	748800
OC-12/STM-4 (622M)	599.04	23961600	2995200
1000BASE-T	1000	40000000	5000000
OC-48/STM-16 (2.5G)	2396.16	95846400	11980800
OC-192/STM-64 (10G)	9584.64	383385600	47923200
10GBASE-SW (WAN)	9584.64	383385600	47923200
10GBASE-SR (LAN)	10000	400000000	50000000

Table 2: Bandwidth-delay product in terms of circuit capacity

The **Capacity (in bytes)** column says it all. This is the theoretical maximum number of bytes in a system, at any given time, such that the circuit is filled to the brim and TCP can resend any dropped or erroneous segments. In a standard TCP implementation, the maximum allowable TCP window is 65 535 bytes. This means that at a minimum rate of 45 Mbit/s and a roundtrip time of 40 ms, a server running normal TCP cannot entirely fill the circuit.

So unless the TCP implementation used can extend its window size beyond 65 535 bytes, the end-user will never be able to transmit a single stream of TCP data (TCP segment) at more than 13.1 Mbit/s for a roundtrip time of 40 ms. As stated above, 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, the TCP performance across a network depends on multiple parameters. So what options does a service provider have? The current service testing trend is based on ITU-T Y.1564. This methodology will perfectly assess network performance, as long as the applications running on it are UDP-based. By assessing the configuration and performance at service activation, Y.1564 provides a comprehensive snapshot of network quality, which is the basis of all 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; it will not assess the end-user's quality of service.

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 the applications, or use software tools that run on computers to emulate TCP traffic. The latter will lead them to the conclusion that the service provider is at fault because their measurement shows that their maximum throughput is nowhere near the bandwidth available. Unfortunately, not all operating systems were created equal. Some have a locked TCP/IP stack and use the basic windowing scheme as defined for TCP, which is 65 535 bytes. These software tools are only as good as the computers they run on. Poor computer performance will reflect poor performance in the measurement, and therefore will provide a false view of network performance.

Given the TCP window limitation, some test methodologies also use multiple TCP sessions to fill the bandwidth. Although this 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. Moreover, simultaneously running multiple TCP sessions will create issues because they will all be trying to send maximum traffic through the test circuit. Since they share the same bandwidth, some will send the traffic at the maximum allowed rate, while others will be in congestion mode and will leave the bandwidth to another session. A measurement could be derived from this methodology, but since this implies multiple averaged measurements, it might not always be repeatable.

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

To prove that their network is not at fault, service providers need a tool that is based on a neutral TCP/IP implementation and can send TCP traffic across the network from each demarcation point. This will allow them to provide the parameters measured to the end-users who can then replicate the test results using their test methodology.

ADVANTAGE OF EXFO'S TCP TESTING METHODOLOGY

Given the necessity of testing TCP performance in their networks, service providers must find a test tool that has the required feature set. Like any test equipment or methodology, the chosen instrument must be easy-to-use, accurate and allow repeatability.

The TCP Throughput feature will help service providers prove that the Carrier Ethernet services being delivered provide the quality of service that an end-user should expect from a TCP application perspective. Thanks to an intuitive graphical user interface, technicians and network professionals alike will test TCP performance in no time.

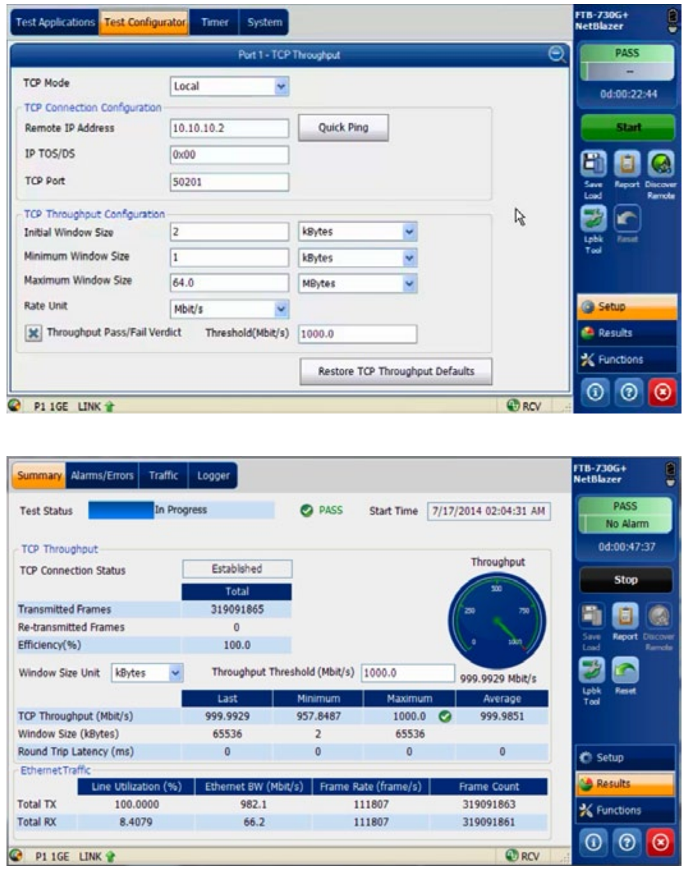

Figure 3. TCP Throughput user interface

Since the TCP Throughput feature is designed in hardware, the EXFO test instrument will always provide an accurate measurement. The reason for this is that it does not rely on the implementation of any communication stack found in PC operating systems or servers.

The methodology used by EXFO is based on the TCP window-scale option as described in RFC 1323, in which a single stream, as documented in RFC 6349, can be used to provide the TCP throughput measurement. It fills a circuit at full bandwidth with TCP traffic when the roundtrip time or the transmission bandwidth is too large for standard TCP implementation. This greatly simplifies the test because it eliminates the need to calculate the number of sessions required for the configuration and which TCP port to use on each. From a results perspective, the user does not have to average multiple test results to validate that the circuit is capable of transporting TCP application traffic. Furthermore, having only one TCP test session should ensure repeatability. If the network conditions (frame loss, roundtrip time, etc.) are the same, the TCP throughput test should generate the same results.

CONCLUSION

To summarize, the TCP protocol is used by most non-real-time applications to deliver mission-critical information from one end to the other of a network. Since the TCP protocol must validate that the information was transmitted without any errors, its built-in functionality will limit its capacity in high-latency or high-bandwidth networks.

Since applications come in all shapes and sizes and run on a wide range of computers/servers, the TCP/IP implementation and configuration will vary from one end-user to another. This can cause issues for service providers because most end-users will base the performance of a network on a server's statistics or software tools running on platforms.

The Ethernet service activation test methodology, which is based on Y.1564, provides part of the solution because it covers a network's capacity to transfer information between two end-points in real time. However, it does not take all of TCP into account. This means that the best solution is to test TCP performance with a tool that can prove that given the proper TCP configuration and network parameters, it is possible to get the maximum TCP traffic across the network.

EXFO Headquarters > Tel.: +1 418 683-0211 | Toll-free: +1 800 663-3936 (USA and Canada) | Fax: +1 418 683-2170 | info@EXFO.com | **www.EXFO.com**

EXFO serves over 2000 customers in more than 100 countries. To find your local office contact details, please go to **www.EXFO.com/contact**.

APNOTE159.2AN

Printed in Canada 14/09