White Paper

File Transfer Protocol (FTP) Throughput Testing

by Rachel Weiss

Introduction

In today's complex networks it is often difficult to correlate different measurements that are reported by network equipment, test equipment, or an end customer. This is because telecommunication networks are comprised of several different layers of services, each with its own protocols and measurements. The measurements are designed to verify that network performance meets the requirements specified for particular protocols and in each customer's service level agreement (SLA). The purpose of this paper is to explain how actual and estimated File Transfer Protocol (FTP) throughput measurements are derived, how they relate to other common network measurements, and how specific factors affect FTP throughput.

Network view

Figure 1 below depicts a typical service provider's Ethernet/IP network. In this example, the customer hand off is a switched layer 2 Ethernet link or a routed layer 3 IP link. The service provider's network connects multiple customer sites (Premise A and Premise B). A common customer practice for verifying FTP throughput is to transfer a file (using FTP) from one location to another, and then measure the time it takes to complete the transfer. The customer then compares the measured FTP throughput with their SLA throughput measurement.

The inherent problem with this approach is that FTP throughput is measured in the upper layers of the network stack, while the throughput requirements stated in most SLAs are typically a layer 1, 2, or 3 measurement. Throughput measured at layer 7 will always be less than that measured at layer 1, 2, or 3 for reasons which will be explored throughout this paper.

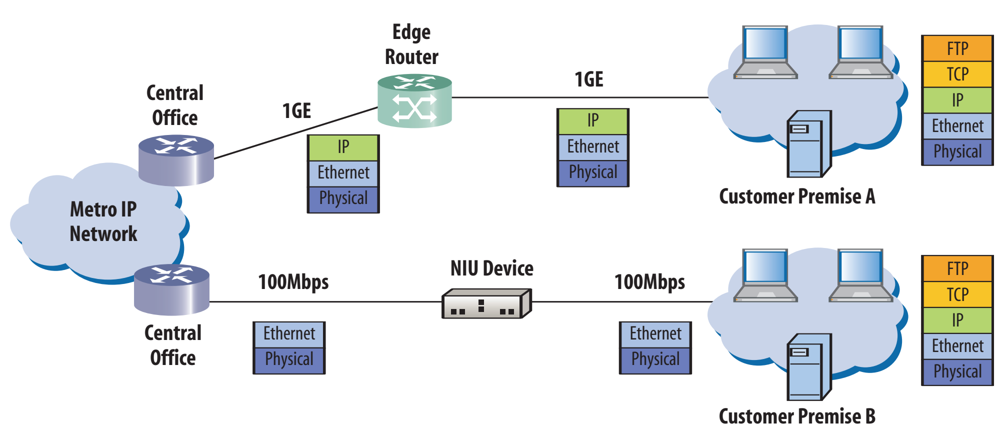

Figure 1: Metro IP Network

FTP, TCP, and the OSI model

FTP is a widely used application that enables a user to transfer files between two computers over a TCP/IP network. The user runs an FTP client application on one computer and the other computer runs an FTP server program. The server is identified by a text name or an IP address. The remote user logs into the FTP server using a login name and password, which the server then authenticates.

FTP runs on the upper layers of the OSI model (see Figure 2) and uses the Transport Control Protocol (TCP) to transport the transferred data. TCP is a connection-oriented protocol residing at layer 4 of the OSI Model. It provides extensive error control and flow control to ensure that data is delivered successfully.

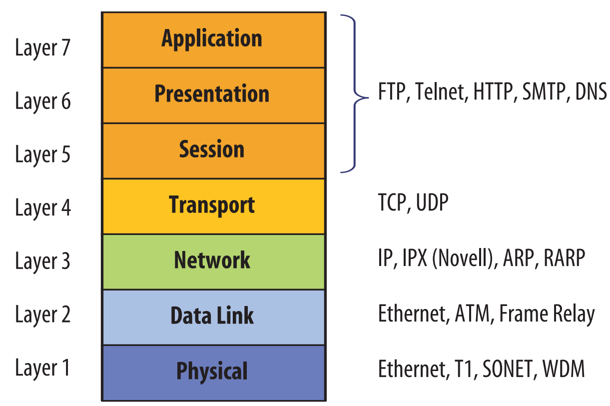

Figure 2: OSI Model

When an FTP client sends a request to transfer a file, a TCP connection or session is initiated through a three-way handshake between the client and server, and synchronization is established. TCP then sends segments of data, each carrying a unique sequence number. Every time the receiving device receives a segment, it must send an acknowledgement back to the sending device. TCP communication is thus a system of transmissions and acknowledgements. The sending device monitors the status of the connection to ensure that each segment is received without errors, and retransmits data if necessary.

TCP also provides flow control by negotiating a window size which dictates how much data a sending device can transmit before receiving an acknowledgement from the receiving device. Before data is transmitted, both devices advertise, and then agree upon the window size that will be used during the session. This prevents situations where the sending device transmits data more quickly than the receiving device can process it, which would result in a loss of data.

Measuring actual FTP throughput

To measure actual FTP throughput, a file of a known size is transferred from a test set (either manually, or using a script) and the time it takes to complete the transfer is measured. FTP throughput is measured in bits per second; therefore, if it took 10 seconds to transfer a 20 Mbyte file, the throughput would be 2 Mbytes per second. This is then multiplied by 8, because each byte contains 8 bits. The resulting actual measurement is 16 Mbps of FTP throughput.

$$FTP\_Throughput = \frac{20\,MByte}{10\sec} \times \frac{8bits}{1Byte}$$

$$FTP\_Throughput = 16\,Mbps$$

It is very important to note that throughput can be measured at each layer of the network stack. Because FTP is an application running at layer 7 of the OSI model, FTP throughput is a layer 7 measurement. This is not the same throughput measurement commonly referred to in network SLAs, which usually refer to layer 1, 2, or 3 throughput. Layer 1 throughput is the maximum rate of traffic that can traverse the link independent of the actual traffic type, and it is the simplest measurement of throughput. For example, if there is a 100 Mbps pipe between two locations, the maximum potential layer 1 throughput is 100 Mbps.

It is not valid to equate FTP throughput with throughput at the lower layers because each layer of the network stack introduces overhead and more complex protocols that will affect the throughput measurement, some of which are out of the service provider's control. Making this assumption would be analogous to comparing apples to oranges. Therefore, directly comparing SLA throughput to FTP throughput is not a correct way to verify the service provider's SLA.

Key factors for calculating or estimating FTP throughput

When measuring the actual FTP throughput on a link or estimating the potential (theoretical) FTP throughput that may be supported on a link, four key factors must be considered: SLA or link speed, frame size, round trip time (RTT), and window size. Based on the values of these factors at any given point in time on a live network, FTP throughput can vary widely. These four factors are used to derive the theoretical calculated FTP throughput measurement in the JDSU FTP script. The following section of this paper will discuss each of these factors in more detail and how this theoretical calculation is derived.

Factor 1 - Link Speed

The first factor is the link speed as contracted in the SLA (for example, a 30 Mbps link). The contracted speed represents the maximum FTP throughput possible on that link. A faster link speed provides higher potential FTP throughput; a lower link speed reduces potential FTP throughput.

Since most networks are comprised of multiple links, it is also important to bear in mind that FTP throughput is limited by the slowest link in the network path. For example, the customer hand-off might be 30 Mbps, but the link to a server from which the file will be transferred could possibly be 10 Mbps. In this case, the maximum potential FTP throughput would be 10 Mbps (not 30 Mbps).

Link speed is just one of the factors used to derive FTP throughput. The following sections discuss additional factors.

Factor 2 - Frame Size

Frame size impacts FTP throughput because each layer of the network stack adds overhead information to the actual user data. Various forms of encapsulation, such as VLAN, also add overhead to the data. The addition of the overhead at each layer lowers the actual data throughput.

For example, in Ethernet untagged frame lengths can vary between 64 and 1518 bytes, and the overhead represents a total of 20 bytes. For a 64 byte frame, 20 bytes of overhead results in an actual data rate on the link of about 75%. The percentage of bytes used for overhead decreases in larger frames so they provide a more efficient data rate. A frame of 1518 bytes still uses 20 bytes for overhead so the actual data rate is close to 99%.

Figure 3 depicts the overhead added at each layer of the stack. It also illustrates how FTP messages are assembled and then processed at each layer.

This entire process can be described using the analogy of nesting Russian dolls. These wooden dolls come in a set, each decreasing in size and placed one insider another. Likewise, assembling an FTP message involves placing the smallest doll (the FTP user data) into a larger doll (a layer 4 TCP message with a TCP header), and then placing this doll into a larger doll (an IP datagram with an IP header), and so on. This process continues down to Layer 1, where the Ethernet frame is converted to ones and zeroes.

When the FTP message is processed on the far end, the far end device opens the largest doll, takes out the next smallest doll, and so on.

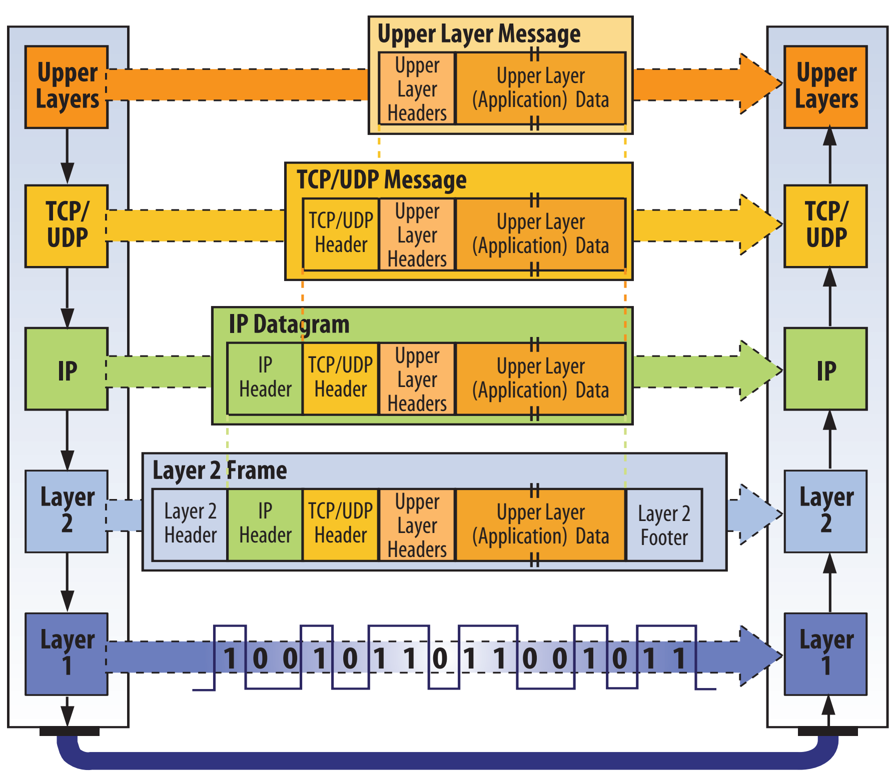

Figure 3: Overhead and Data at Each Layer

It should now be clear that each layer's overhead bits must be considered when calculating FTP throughput. When transporting IPv4 traffic using TCP over Ethernet, the overhead accounts for a total of 78 bytes (see Table 1).

Layer	Overhead	Bytes
1	Preamble, SFD, Interframe Gap	20
2	Ethernet MAC header & CRC	18
3	IP Header	20
4	TCP Header	20

Table 1: Overhead Bytes

To calculate FTP throughput, the number of user data bits carried in a frame (represented by L_{data}) is divided by the total number of bits including the overhead (represented by L_{max}). This results in a scaling factor, which is then multiplied by the link speed to determine the highest practical amount of FTP data throughput that could possibly be achieved on the link.

$$FTP\_Throughput = \frac{L_{data}}{L_{\max}} \times LinkSpeed$$

Calculating the highest practical FTP data throughput for a link speed of 1 Gbps carrying 1518 byte frames yields:

$$FTP\_Throughput = \frac{L_{data}}{L_{\max}} \times LinkSpeed$$
$$FTP\_Throughput = \frac{1460}{1518} \times 1Gbps$$
$$FTP\_Throughput = 962Mbps$$

Note that this is significantly lower than the link speed of 1 Gbps specified in the SLA. If the same calculation is used for a l Gbps link carrying 64 byte frames, the highest practical FTP data throughput is only 94 Mbps! Therefore, by simply varying the frame size, FTP data throughput can range from 94 Mbps to 962 Mbps on a 1 Gbps link.

But, in addition to link speed and frame size, two more factors must be considered: Round Trip Time and the TCP Window Size.

Factor 3 - Round Trip Time

Round trip time is the third factor that must be considered when calculating FTP throughput. Round trip time has a negative affect on FTP throughput. Typically, the longer the round trip time the lower the FTP throughput.

The maximum acceptable round trip time for an Ethernet link is usually specified in the customer's SLA. It is defined as the time that it takes for the data to travel to the far end of the link, and then back again to the transmitter. In live networks, round trip time is affected by the distance between the transmitting node and the far end node, the routing and switching times of equipment on the link, and any traffic congestion that exists on the link.

This factor must be considered because the TCP protocol requires that each received segment of data be acknowledged. For every segment that is transmitted, the receiver must send back an acknowledgement to ensure error free transmission. Therefore, the time it takes to send the data to the far end and the time that it takes to send the acknowledgement back to the transmitter must be considered.

In contrast, layer 1, layer 2, and layer 3 throughput (typically specified in the customer's SLA) is calculated based on the time it takes to transmit in one direction. This is because the TCP protocol does not apply at layer 1, 2, or 3. Essentially, both directions must be taken into account when measuring FTP throughput because FTP resides on the upper layers and uses the TCP protocol.

The following section discusses how round trip time fits into the FTP throughput calculation, along with a related factor, the window size.

Factor 4 - Window Size

The window size parameter used in the TCP protocol defines how much data a transmitter can send at one time. In a practical sense, this addresses the limitations of the receiving device at the endpoint of the file transfer, such as the processor and memory speeds. In another sense, the window size represents the receiver's link capacity, since it dictates how much data can be transmitted over the link and successfully received.

Link Capacity: Ideal Window Size

In order to understand how window size impacts FTP throughput, the idea of the physical capacity of a link must be explored. Physical capacity can be defined as the maximum number of bits that can traverse the link at a particular point in time. This can be thought of as the size of the pipe.

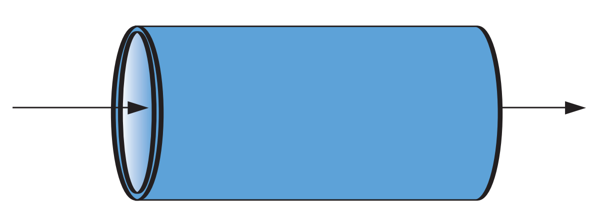

Figure 4: Link Capacity

The physical capacity (C) of an Ethernet link is determined by multiplying the link speed (LinkSpeed) by the measured round trip time in seconds (RTT):

$$C = LinkSpeed \times RTT$$

The ideal window size would be equal to the physical capacity of an Ethernet link. In other words, the best scenario is that in which TCP uses the full physical capacity of the link as its window size. The capacity of a link is a derived concept that will be used in the following section to build the FTP throughput equation.

Practical Window Size

Window size can be a significant limiting factor when calculating FTP throughput, because regardless of how large the physical capacity of the link may be, if the TCP window size is smaller than the physical capacity, the window size must be used to calculate the FTP throughput since it represents the true capacity of the link.

For example, using a link speed of 1 Gbps and a round trip time of 20 milliseconds (ms), the physical capacity of the link would be equal to:

$$C = LinkSpeed \times RTT$$

$$C = 1000000000 \times .020$$

$$C = 20{,}000{,}000 bits$$

This calculation results in a capacity of 20,000,000 bits. However, in the TCP protocol the largest window size that can exist in the Microsoft Windows 2000 operating system is 65,535 bytes, or 524,280 bits. In this example the value used for capacity in the FTP throughput equation must be the window size.

Using the concept of capacity and window size, the FTP throughput equation provided in Section 5.2 can now be modified to account for the round trip time. To do this, the window size is divided by the physical capacity to determine another scaling factor for the equation:

$$FTP\_Throughput = \frac{L_{data}}{L_{\max}} \times LinkSpeed \times \frac{W}{C}$$

$$FTP\_Throughput = \frac{L_{data}}{L_{\max}} \times LinkSpeed \times \frac{W}{LinkSpeed \times RTT}$$

$$FTP\_Throughput = \frac{L_{data}}{L_{\max}} \times \frac{W}{RTT}$$

Assuming the same 1 Gbps link, with RTT of 20 ms, a frame size of 1518 bytes, and a window size of 524,280 bits, the result is as follows:

$$FTP\_Throughput = \frac{L_{data}}{L_{\max}} \times \frac{W}{RTT}$$
$$FTP\_Throughput = \frac{1460}{1518} \times \frac{524280}{.020}$$
$$FTP\_Throughput = 25.2Mbps$$

When the window size and round trip time are taken into consideration, this yields a practical maximum FTP throughput result of 25.2 Mbps. This is significantly less than the layer 1, 2, or 3 throughput that will be measured on this 1 Gbps link!

If the window size happens to be greater than the physical capacity of the link, it can be ignored and the physical capacity can be used in the calculation. Actual window sizes vary in different operating systems.

Conclusion

The actual FTP throughput measurement and the theoretical calculation are complex concepts because the protocols residing at each layer of the network stack must be taken into consideration. As illustrated in the preceding sections, FTP throughput is typically measured significantly lower than the contracted SLA's layer 1, 2, or 3 throughput. This is due to a number of factors previously discussed, and a lower FTP throughput measurement does not necessarily indicate a problem with the link or that the link is not performing per the customer's SLA.

The four factors described in this paper can vary significantly in live networks, yielding different results under various circumstances. By varying frame sizes and window sizes, a full range of potential FTP throughput values can be calculated. One solution, the JDSU FTP script, provides this information in a table that can be a very useful tool for fully characterizing FTP throughput over a link. The table can be used to determine how some of these parameters might be changed to improve FTP performance on an Ethernet or IP link.

Finally, FTP throughput should not be used alone to evaluate the performance of a contracted link with a service provider. There are many other essential measurements required to fully characterize the link, such as layer 1, 2, or 3 throughput, round trip time, and frame loss.

Test & Measurement Regional Sales

NORTH AMERICA	LATIN AMERICA	ASIA PACIFIC	EMEA	WEBSITE:
TEL: 1 866 228 3762	TEL: +55 11 5503 3800	TEL: +852 2892 0990	TEL: +49 7121 86 2222	www.jdsu.com/test
FAX: +1 301 353 9216	FAX: +55 11 5505 1598	FAX: +852 2892 0770	FAX: +49 7121 86 1222	