JDSU

White Paper

The Need for Application-Centric Turn-Up Testing

By Barry Constantine

Introduction

As communication providers strive to improve their Ethernet service turn-up processes, the fundamental goal is to provide customers with expected service quality levels. Closely coupled to this goal is the ability to verify that the network is performing acceptably when customers complain that their applications are slow due to the network.

At the same time, service providers have begun offering Layer 2/Layer 3 (L2/L3) service level agreements (SLAs) and Ethernet-based turn-up testing has evolved from more traditional bit error rate test (BERT)-based turn-up (Physical Layer only testing) to L2 (Ethernet)- and L3-(Internet Protocol [IP])-based turn-up testing. In this application note, L2- and L3-based turn-up testing are referred to as "network-centric turn-up."

While network-centric turn-up verifies the L2/L3 components of the carrier SLA, it does not guarantee customer satisfaction. Specifically, customer applications (Layer 4 [L4] and above) can perform poorly even over carrier networks that have been tested for proper L2/L3 performance, or acceptable throughput, packet loss, delay, and jitter testing.

To further ensure that the customer applications perform satisfactorily, service providers must also conduct application-centric turn-up testing to qualify their networks. Application-centric turn-up testing will not only help eliminate improper carrier network configuration issues (such as misconfigured quality of service [QoS] settings and Transmission Control Protocol [TCP] tail drop due to router queue settings), it will enable the carrier to "defend" the network's performance and also provide solid customer recommendations for possible enterprise equipment issues.

Figure 1 highlights an enhanced workflow for turning up L2/L3-based IP services, emphasizing application/service level verification.

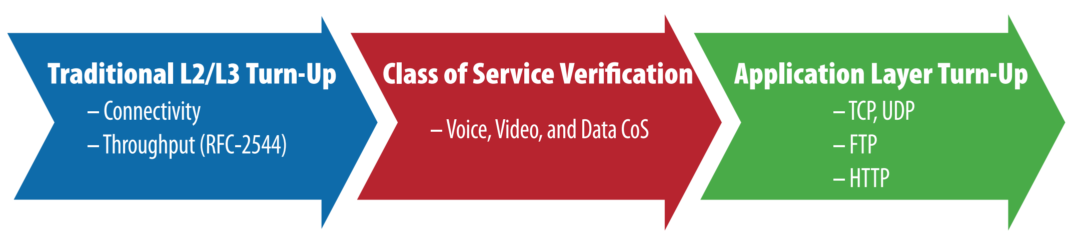

Figure 1: Workflow for Application-Centric Turn-Up

This application note presents the following topics:

– Overview of network-centric turn-up (today's traditional L2/L3 turn-up process)
– Review of the enhanced turn-up process (adding an application-centric test suite)
– Introduction to application-centric turn-up (using the JDSU T-BERD/MTS-8000 or T-BERD/MTS-6000 Metro Ethernet services testers)

Today's Network-Centric Service Turn-Up Process

In today's Ethernet service turn-up process the communications provider must verify the ability of the network to meet various SLAs for L2 (Ethernet) and L3 (IP) performance. Generally these SLAs specify throughput, loss, delay, and jitter limits under varying network loads. Figure 2 will be used throughout this section because it represents the general process for network-centric turn-up.

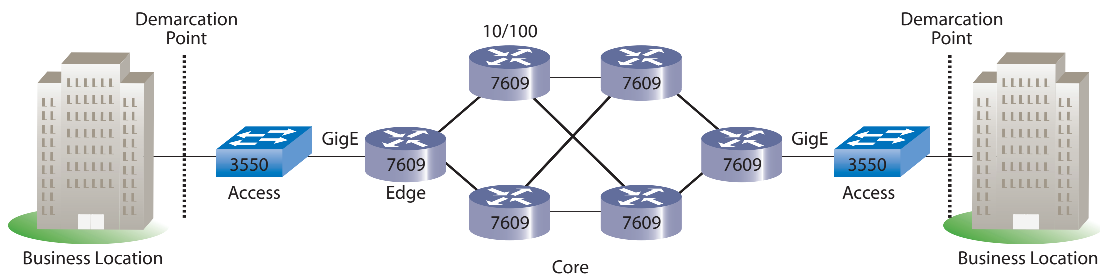

Figure 2: Representative Network Turn-Up Scenario

The following sections summarize the basic turn-up test flow and portray a general, rather than specific, workflow.

Basic Connectivity and Throughput Testing

A common first step in turn-up testing is to verify connectivity in both switched L2 networks and routed L3 networks. Throughput testing cannot be conducted until end-to-end network connectivity is established.

In a L2, switched Ethernet network, it is impossible to verify connectivity using Internet control message protocol (ICMP) pings (since the ICMP is a L3 protocol). The test equipment must generate and analyze special L2 test frames to verify L2 connectivity. Figure 3 portrays a typical end-to-end test scenario using test equipment to generate test traffic. The connectivity of L3-based networks can generally be tested using ICMP pings (covered in the next section) or with an end-to-end test scenario, also shown in Figure 3.

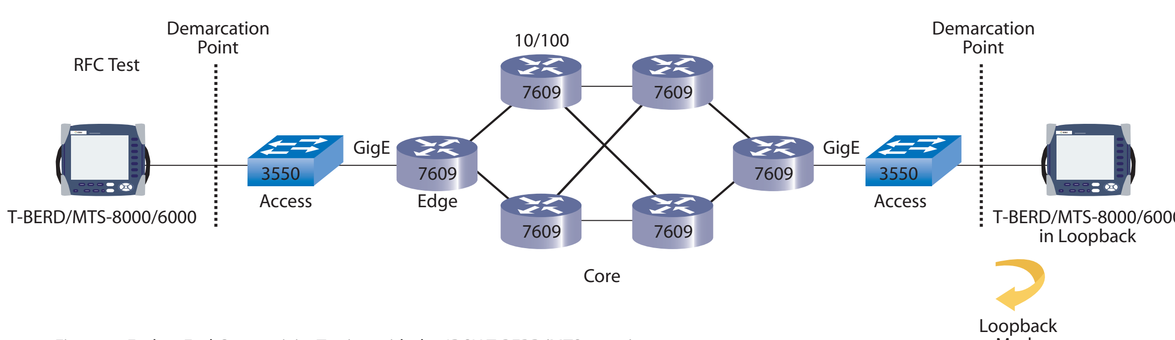

Figure 3: End-to-End Connectivity Testing with the JDSU T-BERD/MTS-8000/6000

Proper destination addressing, such as mandatory access control (MAC) addresses and virtual local area network (VLAN) tags, must be configured to ensure that the local test device can send L2 frames successfully to the remote test device. The remote test device transmits a frame that must also be capable of L2 address source and destination "swapping." In other words, the local test device reaches the remote test device which then must place the address of the local test device in the destination address of the frame, so that test frame can travel back to the local test device.

For L3-(IP)-based networks, connectivity testing can generally be conducted using an ICMP ping capability. While ping testing is a good way to establish basic connectivity, it does not provide a good measure of the end-to-end SLA. ICMP can also be blocked by routers or given low priority under congestion conditions. So it is advisable to conduct L3 throughput testing in the manner previously described for L2.

To add automation and repeatability to the throughput testing process, RFC-2544-based network testing is extremely useful. RFC-2544 is an industry standard benchmarking methodology for network elements such as routers and switches, which is now the de-facto standard for installation testing of Ethernet/IP networks. The RFC-2544 process encompasses throughput testing at various frame or packet sizes to verify that the network topology (or subsection) can sustain prescribed traffic requirements. The RFC test conducts these throughput tests and verifies that the required SLA thresholds for the network under test can be met.

Class of Service Turn-Up Verification

After basic connectivity and throughput testing are successfully completed, proper Class of Service (CoS) settings throughout the network must be verified. Traffic prioritization mechanisms, including VLAN tag, VLAN priority, type of service (TOS), and diff serve code points (DSCPs), set the different traffic priorities and allow providers to create differentiated classes of service, which enables coexistence of mission-critical applications such as voice and video with best-effort traffic.

Because CoS testing should emulate the profile of Triple-Play traffic such as voice, video, and data, it is important to verify that the intended CoS for voice and video service is being provided, even when bursty data traffic exceeds normal thresholds. Although, it can be very cumbersome to configure the various IP test streams in terms of representative frame lengths, bit rates, and CoS settings to properly emulate Triple-Play customer traffic.

JDSU has simplified this process by providing a Triple-Play Turn-up test application that greatly simplifies the test and verification of a network's capability to carry Triple-Play traffic. Figure 4 represents the Triple-Play services that ride on various physical network links.

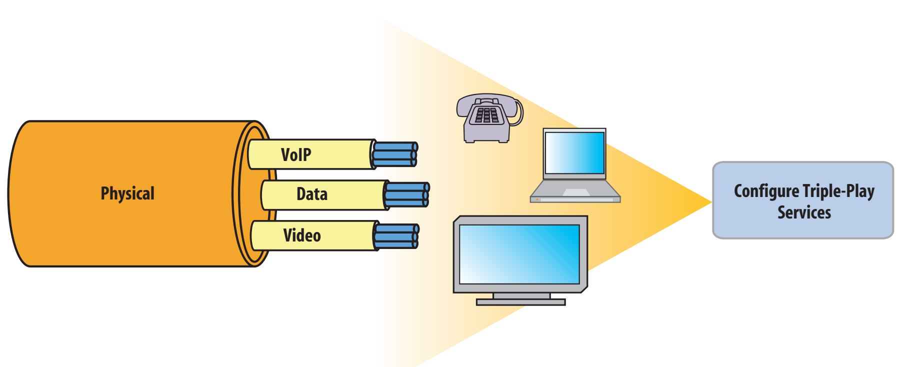

Figure 4: Conceptual View of Triple-Play Services

The JDSU T-BERD/MTS-8000 and T-BERD/MTS-6000 Triple-Play configuration wizard simplifies the setup of these services and eliminates the need for the user to be a voice, video, and data expert. For example, understanding the differences in bandwidth or compression of Standard (SDTV) or High (HDTV) Definition Television is not required. Users simply configure the number of representative SDTV and HDTV channels desired (along with voice calls and data traffic). Figure 5 shows the easy-to-use Triple-Play configuration screen of the T-BERD/MTS-8000 and T-BERD/MTS-6000.

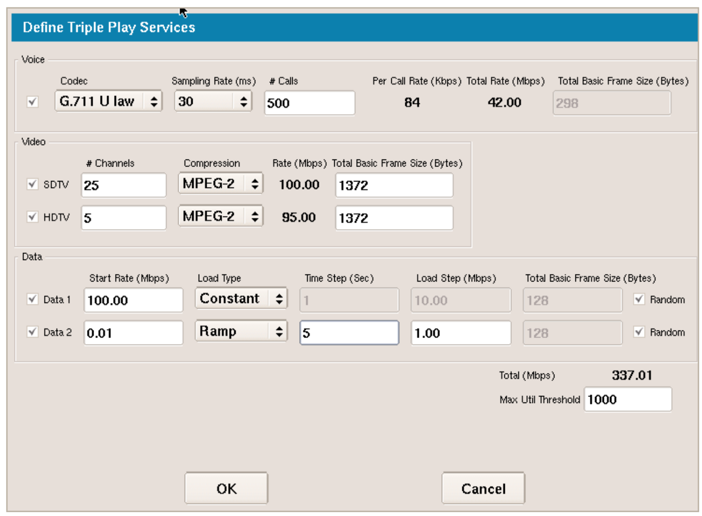

Figure 5: Configuring Triple-Play Profiles

Once the user configures the Triple-Play profiles screen, the unit presents an easy to understand summary screen including a network "pipe" diagram similar to the one shown in Figure 6.

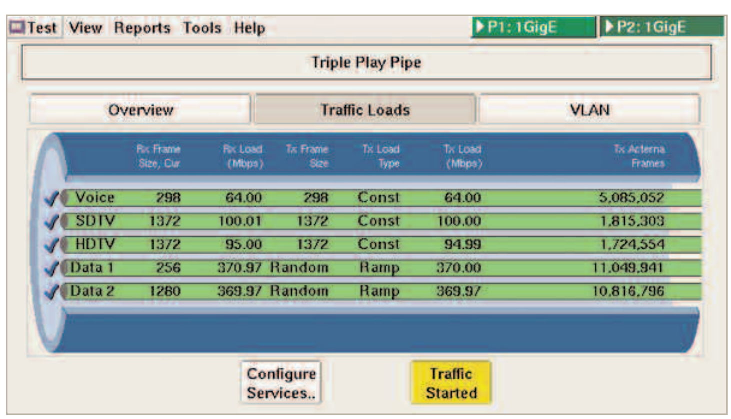

Figure 6: Triple-Play Summary Network Pipe Screen

Once the Triple-Play configuration is complete, the test is run between two JDSU T-BERD/MTS-8000 and/or T-BERD/MTS-6000 testers in an end-to-end manner (at opposite ends of the network under test). Again, the main purpose of CoS in a network is to ensure that higher priority traffic such as voice and video maintain the required SLA thresholds (throughput, packet loss, delay, and jitter) in the midst of dynamically varying, lower-priority traffic such as data, for example, Hypertext Transfer Protocol—Worldwide Web Protocol (HTTP) and File Transfer Protocol (FTP).

Interpreting these results is greatly simplified by viewing graphs for the various network performance metrics such as throughput, frame loss, delay, and jitter for each of the emulated Triple-Play service streams. The T-BERD/MTS-8000 and T-BERD/MTS-6000 both provide green- and red-light indicators for the overall physical link status. Figures 7 and 8 display typical Triple-Play result graphs provided by both JDSU T-BERD/MTS test platforms.

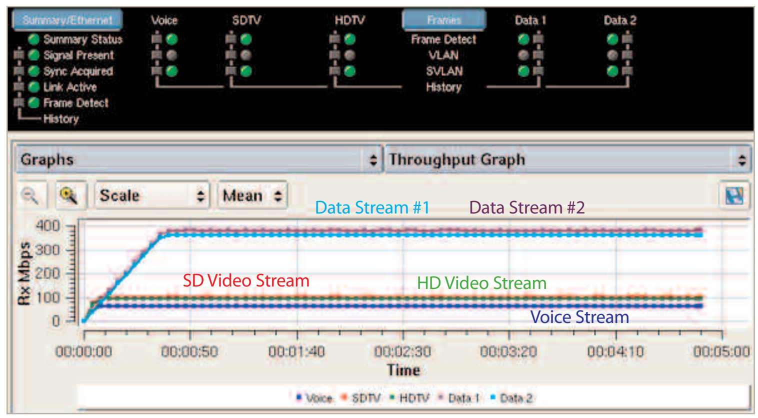

Figure 7: Throughput of Data Service Ramping Up

In Figure 7 note that as the Data Streams 1 and 2 ramp up to full link capacity (1 Gb/s in this example), the voice and video streams receive proper network priority and are not affected.

If the HD video content increases beyond the link capacity of the network, and the network has been configured properly from a CoS perspective, the lower priority data service would be discarded rather than the video or voice services, as Figure 8 illustrates.

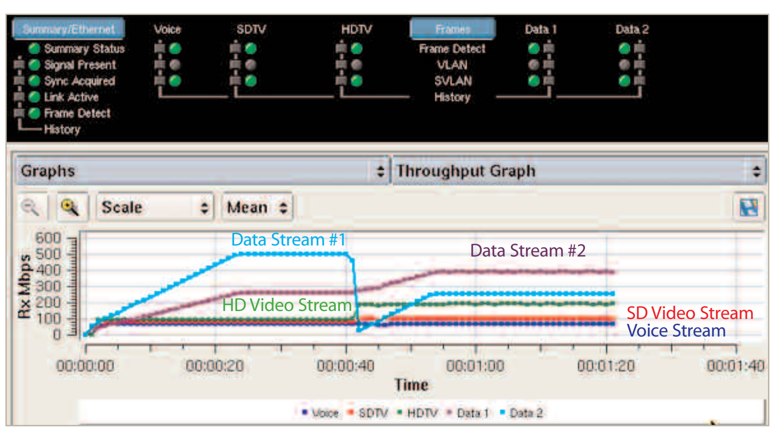

Figure 8: HD Video Increases, Data Stream Experiences Frame Loss

The Triple-Play correlation capability of the JDSU T-BERD/MTS test platforms provide easy-to-understand graphs. These graphs quickly highlight the outcome of the test, which would otherwise be a complex task of interpreting the statistics in grids/tables.

Enhanced Network Turn-Up with Application-Centric Turn-Up Process

The traditional network-level turn-up process highlighted in the previous section is insufficient as the sole means for verifying network functionality prior to customer usage. Application-level testing, such as L4 TCP throughput, FTP, and HTTP, is necessary for verifying the true user experience.

Conducting a turn-up test that emulates real customer traffic enables service providers to confidently activate a new customer and resolve post turn-up finger-pointing should problems arise. A classic example of this problem occurs when the customer has misconfiguration issues associated with TCP settings. For example, on a 10 Mb/s SLA (over a 100 Mb/s Ethernet Metro network) the network turn-up process (L3) verifies proper SLA by running RFC-2544 testing. However, upon service turn-up, the provider immediately gets a complaint from the customer concerning slow network performance. The finger pointing begins and only ends when a Tier 3 field support engineer travels to the customer site and validates the network SLA by running L4 tests (usually with a laptop computer).

In this example, finger pointing could be eliminated when the network is tested from an application perspective. With L4-enabled test equipment, a Tier 1 technician can easily run L4 throughput tests and deliver a service turn-up report, which would identify the specific TCP settings used to achieve the promised SLA.

TCP throughput testing is one example of application-centric turn-up. The following sections will specifically demonstrate how to use the JDSU T-BERD/MTS testers to verify:

- **TCP Window Size:** Automated verification of the optimal TCP Window size for the customer network configuration using the JDSU T-BERD/MTS-8000 or T-BERD/MTS-6000, which will "walk the window" (vary its size)
- **Router Queue Settings:** Tests for expected TCP performance in the midst of up to 10 background IP streams (User Datagram Protocol [UDP] or TCP)
- **FTP Throughput:** Automated verification of satisfactory throughput with FTP downloads and uploads
- **HTTP Throughput:** Automated downloads from either a single or multiple Web sites to provide key performance statistic

TCP Throughput Testing

The TCP Window is one of the most important settings in terms of an application's performance over a network. This setting varies greatly between operating systems and can also be affected by L4 proxies and any device that regenerates TCP connections.

In reality, two windows are configurable in TCP; the Receive Window and the Congestion, or Send Window. In most cases, it is the setting for the Send Window that is not optimized for network latency. The Send Window must be tuned properly with the latency of the network to achieve full throughput. The establishment of the proper Send Window size can be a complicated and tedious task, but the JDSU Walk the Window script automates the process.

Upon activating the Walk the Window script, the user simply configures the starting and ending Window sizes, the number of Windows to test (between the starting and ending Window size), and the duration of each trial test (as demonstrated in Figure 9).

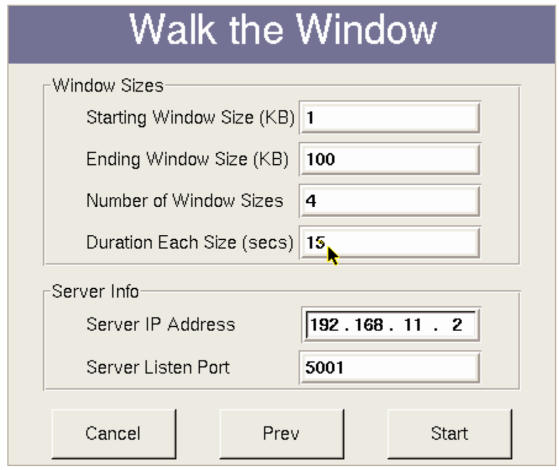

Figure 9: TCP Walk the Window Script Configuration

The test is then run between the two end-to-end JDSU T-BERD/MTS-8000 or T-BERD/MTS-6000 testers. After completion, the user receives an easy-to-understand test report that clearly highlights the Window size versus throughput performance (Figure 10). This example shows that the optimal TCP Window size is 127 Kbytes, which achieves maximum L4 throughput of ~7 Mb/s.

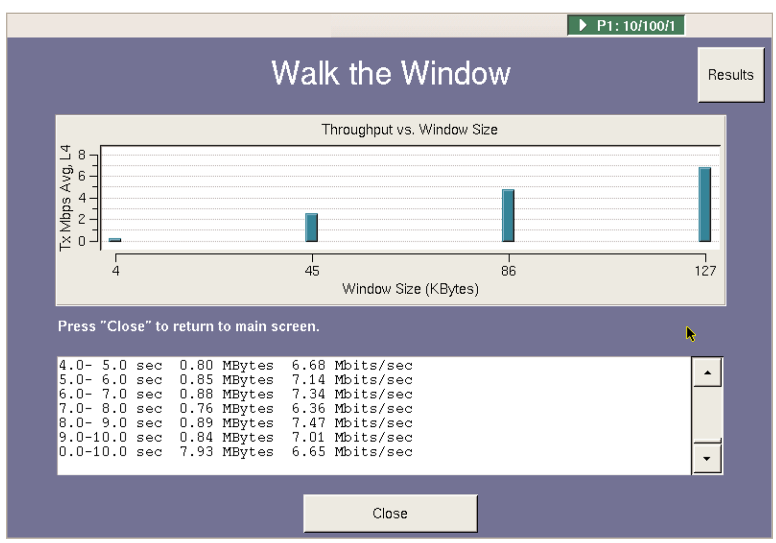

Figure 10: TCP Walk the Window Script Results

Once the optimal Window size is established, the next step is to conduct TCP throughput testing with a mixture of background traffic, such as UDP, to verify customer throughput under realistic network conditions. Router queue settings are a common cause of customer throughput issues and cannot be diagnosed unless the network is under load. TCP Tail Drop issues are surprisingly widespread in networks and many times go undetected until a customer exceeds a particular traffic threshold.

TCP Tail Drop Explained

In properly configured routers, the buffer mechanism is set to periodically drop packets when the router becomes congested. Routers use various "Early Discard" algorithms to alleviate congestive situations. An improperly configured router allows buffer fills and will drop the "tail" of the buffer, which causes a sequential block of packets to drop and adversely affect the built-in retransmission mechanisms of TCP. The net result is that the multitude of TCP hosts such as servers and PCs begin to retransmit in synchronicity, causing an endless loop of router-congestive overload and discards. The user experiences an incredible slowdown in TCP sessions such as HTTP and, in many cases, causes the infinite "spinning hourglass" on Windows workstations.

The JDSU T-BERD/MTS test platforms provide for full background traffic loading and representative TCP foreground testing at lines rates up to 10 Gb/s. Current users find this feature very easy to use because the TCP foreground traffic simply appears as another stream in the multiple streams configuration (Figure 11).

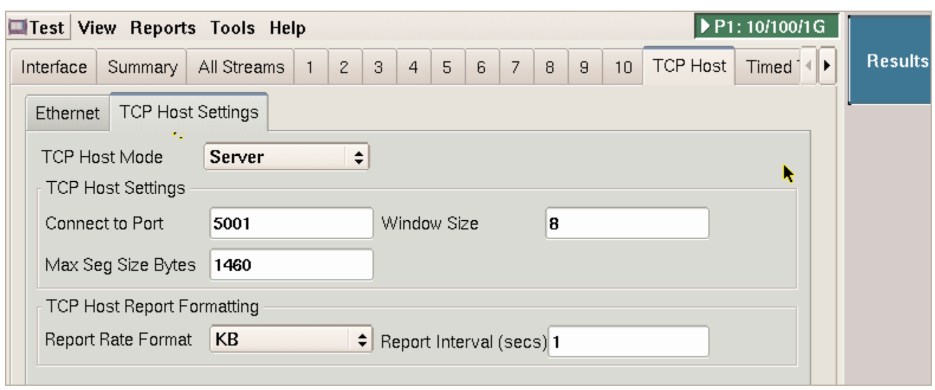

Figure 11: TCP Stream and Background Streams Configuration

The setup of the TCP stream (TCP Host Settings tab) is simple. Users can set the Window size to the optimal size automatically determined in the Walk the Window script from the previous section.

Each background stream can be configured as L3 or L4 traffic (UDP or TCP), with various encapsulations, such as VLAN; 802.1q Tunnel Tags (Cisco), or Q-in-Q; and DSCP. Each background stream can also be allocated specific throughput, frame sizes, and traffic profiles (constant or ramp).

After configuring the TCP foreground and IP background streams, users can review the network load profile with an easy-to-understand "network pipe" diagram. The diagram summarizes the relative bandwidth allocated to each stream, traffic load type, and other pertinent network settings (Figure 12).

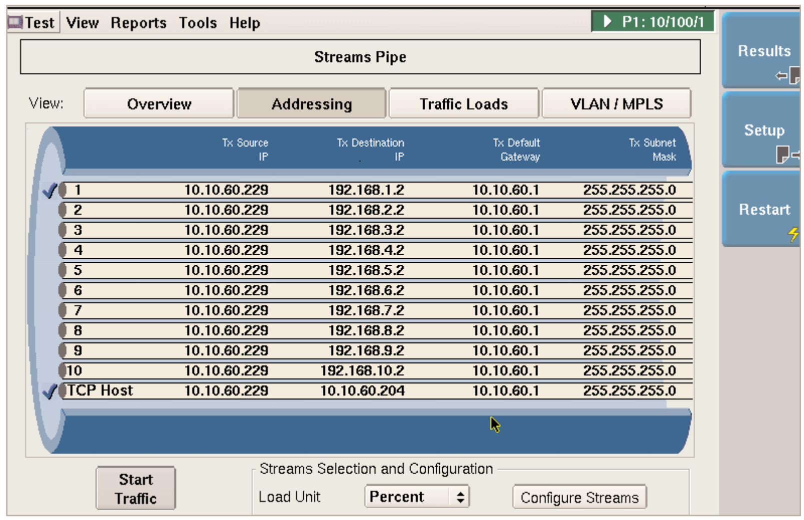

Figure 12: Network Pipe View of TCP and IP Streams Profile

Again, the primary goal of TCP and background testing is to verify that TCP throughput is sustainable under network stress conditions, which will detect incorrect CoS settings and router queuing issues. By viewing the graph of TCP stream versus the background IP stream throughput, users can easily determine whether the TCP throughput was preserved under network load. This network load may be caused by excessive UDP traffic, which is often the case when customers stream audio or video. In Figure 13, it is obvious that TCP throughput is maintained in the midst of increasing UDP traffic load.

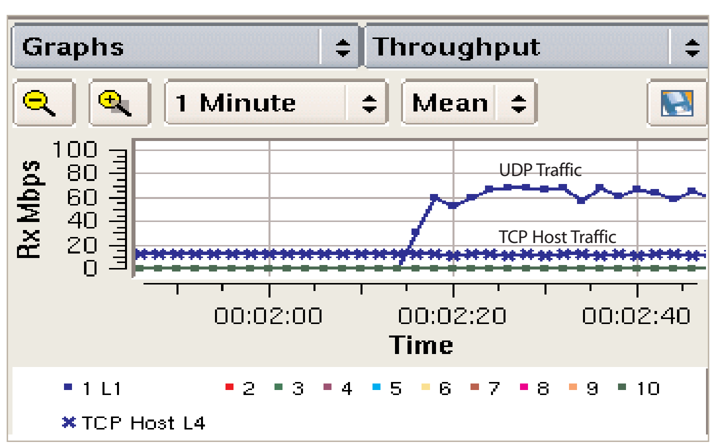

Figure 13: TCP Throughput Remains Stable as UDP Traffic Rises

Figure 14 demonstrates the condition where a router is not set to prioritize TCP versus UDP traffic. In this case, the TCP traffic experiences loss, which in turn degrades customer application performance due to TCP retransmissions.

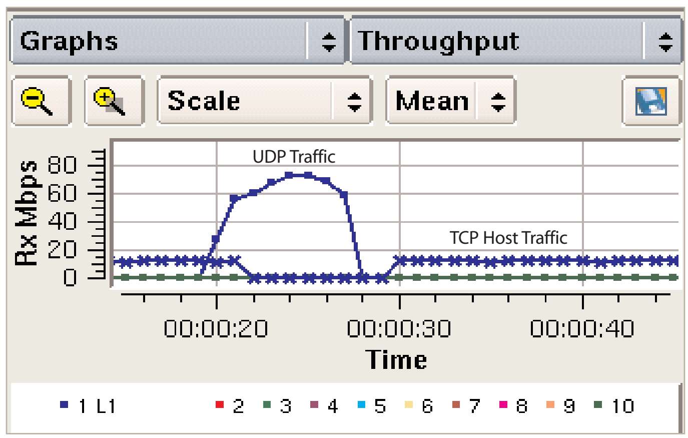

Figure 14: TCP Throughput Degrades as UDP Traffic Rises

When testing TCP with the JDSU T-BERD/MTS-8000 or T-BERD/MTS-6000, note that the tester can interoperate with standard servers, such as Windows or Linux computers, running popular open source TCP throughput packages such as iperf and ttcp. This testing lets users verify TCP throughput between a T-BERD/MTS tester and actual servers (test servers or other customer premises equipment) to diagnose TCP issues. The T-BERD/MTS can act as either TCP Client or Server for this type of test, which eliminates the need for a Tier 3 technician to conduct the same type of test with a laptop computer.

Now that L4 TCP performance has been verified and proper settings have been established, the application turn-up testing can commence up the stack to FTP and then HTTP.

FTP Throughput Testing

FTP is a very important application to verify because it places intensive bulk data transfer demands on the network. It is important to validate both download and upload file performance. The test tool must provide this FTP throughput test capability, which again eliminates the need for a technician to manually conduct tests with a laptop computer.

The JDSU T-BERD/MTS test platforms simplify FTP throughput testing with a wizard-like configuration interface. The user simply selects the desired FTP file sizes and the number of tests to conduct for each, as shown in Figure 15.

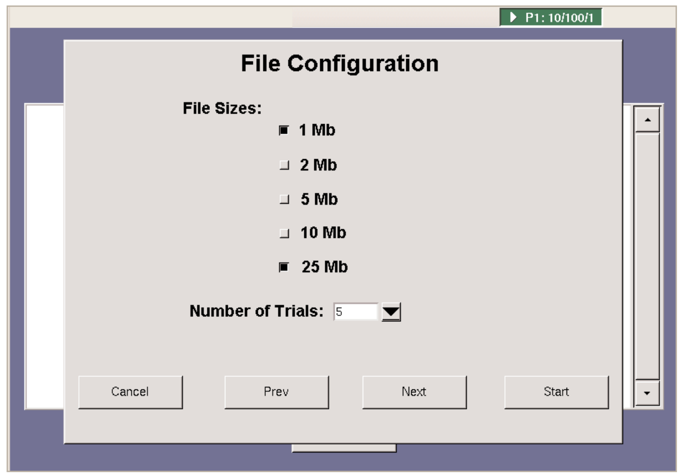

Figure 15: FTP Throughput Test Configuration

The test is then run between the two end-to-end JDSU T-BERD/MTS-8000 or T-BERD/MTS-6000 testers and upon completion, an easy-to-understand test report is presented to the user. As Figure 16 shows, the test report clearly highlights the file size versus throughput in both the FTP upload (PUT) and download (GET) directions.

Summary of all results:

File Size MB	PUT Mbps	GET Mbps
1	Min 24.39	Min 13.82
	Avg 24.39	Avg 13.82
	Max 24.39	Max 13.82
2	Min 16.0	Min 11.5
	Avg 16.0	Avg 11.5
	Max 16.0	Max 11.5
5	Min 11.08	Min 11.85
	Avg 11.08	Avg 11.85
	Max 11.08	Max 11.85
10	Min 5.926	Min 11.18
	Avg 5.926	Avg 11.18
	Max 5.926	Max 11.18
25	Min 5.263	Min 10.78
	Avg 5.263	Avg 10.78
	Max 5.263	Max 10.78

Figure 16: FTP Throughput Test Results

Combining a fully automated FTP throughput test with simple-to-understand results lets the user quickly verify the acceptability of the customer experience. This capability provides far more confidence in delivering acceptable customer performance than only L3 packet-level testing, which has no direct correlation to the customer experience.

Upon completion of the FTP testing, the final application turn-up test involves HTTP throughput testing, which is critical in terms of daily customer performance expectations.

HTTP Throughput Testing

Even with the structured process of testing the TCP (L4) performance and FTP (application layer) performance, HTTP throughput performance provides the most reflective test with respect to end customer experience. From casual Web browsing to Web-based applications, such as Simple Object Access Protocol (SOAP), HTTP is central to the customer experience.

The JDSU T-BERD/MTS test platforms also simplify HTTP throughput testing by providing a wizard-like configuration interface. The user simply selects from a list of Web pages (Universal Resource Locators, URLs) and the T-BERD/MTS-8000 or T-BERD/MTS-6000 automatically conducts an HTTP throughput test for the Web page. Figure 17 provides a screenshot of the JDSU T-BERD/MTS HTTP configuration wizard.

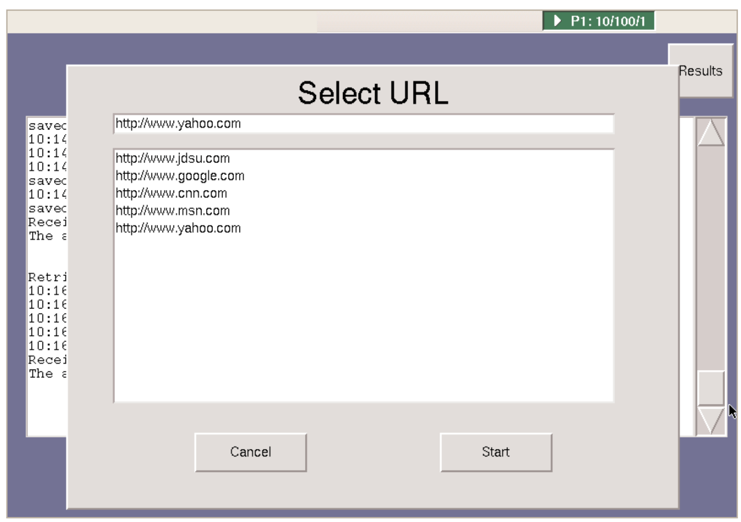

Figure 17: HTTP Throughput Test Configuration

The test is then run between a JDSU T-BERD/MTS-8000 or T-BERD/MTS-6000 and a live Web server, which can be a provider's test Web server or operational Web server. Upon completion of the entire HTTP test (users can continue to select URLs), an easy-to-understand test report is presented to the user that clearly highlights the Web page sizes versus throughput for each URL, as shown in Figure 18.

HTTP Throughput Test Report

Date	05/17/08
Time Start	10:19:50 AM
Time End	10:23:14 AM

Summary of Page 1:

URL	http://www.yahoo.com
Times visited	1
Website size	13790 bytes
Min Rate	3.169 Mbps
Avg Rate	3.169 Mbps
Max Rate	3.169 Mbps

Summary of Page 2:

URL	http://www.google.com
Times visited	2
Website size	25625 bytes
Min Rate	0.7091 Mbps
Avg Rate	0.7435 Mbps
Max Rate	0.7777 Mbps

Figure 18: HTTP Throughput Results – Summary after Completion

The results clearly indicate the achieved throughput for various sites. Results can also be output in PDF format (as is the case for all results highlighted in this application note).

Conclusion

As highlighted throughout this document, L2/L3 network turn-up is insufficient for verifying the ability of the network to carry customer application-level traffic. Specifically, customer applications (L4 and above) can perform poorly even over carrier networks that have been tested for proper L2/L3 performance, such as acceptable throughput, packet loss, delay, and jitter.

This application note provides an enhanced turn-up process for IP services and offers.

The following JDSU products provide application-centric turn-up capabilities.

T-BERD/MTS-8000 Transport Module

T-BERD/MTS-6000 Compact Optical Test Platform

HST-3000 Handheld Services Tester

The following links reference other useful testing-related application notes and white papers on the JDSU Web site:

White Paper: Fundamentals of Ethernet

White Paper: Understanding Ethernet and Fibre-Channel Standard-Based Test Patterns

White Paper: Verifying Metro Ethernet Quality of Service

Poster: Carrier Ethernet Service Testing

Test & Measurement Regional Sales

NORTH AMERICA	LATIN AMERICA	ASIA PACIFIC	EMEA	WEBSITE: www.jdsu.com
TEL: 1 866 228 3762	TEL: +1 954 688 5660	TEL: +852 2892 0990	TEL: +49 7121 86 2222	
FAX: +1 301 353 9216	FAX: +1 954 345 4668	FAX: +852 2892 0770	FAX: +49 7121 86 1222	

 30149464 001 1008 TURNUP.WP.ACC.TM.AE October 2008