RFC 6349 and TrueSpeed™: Experience Your Network As Your Customers Do

The Gap in Ethernet Turn-up Testing

- RFC 2544 and Y.1564 tests verify network performance in Layers 2/3, but customers still blame the network when their applications run slowly. Customer business applications run over TCP (Layer 4)—a layer typically not tested at turn-up.
- RFC 6349-compliant TrueSpeed, installed on a JDSU T-BERD®/MTS-6000A tester, bridges the testing gap by adding Layer 4 TCP tests during turn-up. This testing capability prevents costly problems that negatively affect the customer experience and increase churn.
- TrueSpeed lets the same technician complete additional TCP tests in 3 minutes, saving up to 30 percent in OpEx by preventing or quickly resolving painful finger-pointing scenarios.

RFC 6349 bridges the gap in service-activation testing.

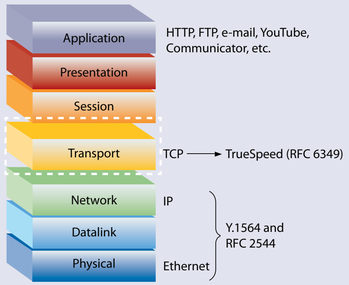

Typical TrueSpeed Turn-up Scenarios

TCP Problem #1: Misconfigured CPE End Host

After testing the provider's network, TrueSpeed verifies TCP throughput in an end-to-end manner. This test often shows that TCP performance issues reside in equipment configuration, such as servers and firewalls, at the customer site.

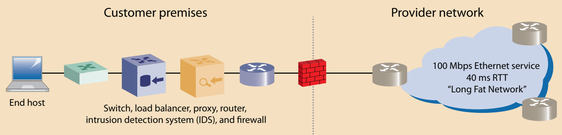

Problem	Description	Location/Responsibility	How can the problem be fixed?
Window size not optimized to network conditions.	High latency and/or high-bandwidth paths may require end-host TCP window adjustments for optimal TCP throughput.	End hosts, switches, routers, content caching, and inspection devices are the responsibility of the business customer.	The end customer may need to tune end-host devices and/or other content engines.

TCP Window Size

TCP window size is the amount of data sent over the network before it requires an acknowledgment (ACK) from the receiver. The TCP window size should equal or exceed the bandwidth delay product (BDP).

In this example, the BDP is 140 KB (bottleneck bandwidth x RTT/8 or 45 Mbps x 25 ms/8). To achieve full TCP throughput, end hosts must have TCP window sizes greater than the BDP. With 64 KB TCP windows, a typical default setting, only 20.5 Mbps TCP throughput is achieved.

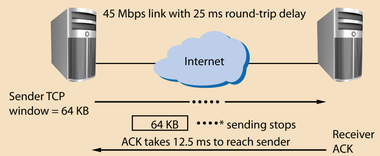

TCP Problem #2: Network Buffer Queues

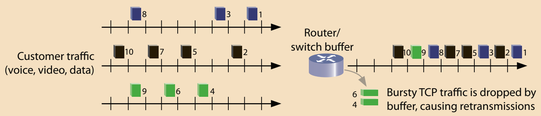

Problem	Description	Location/Responsibility	How can the problem be fixed?
Network queue buffer size and/or type may be inadequate for customer applications.	Because TCP traffic is bursty, it can sometimes exceed buffer queue capacity, causing drops and retransmissions.	The provider router and switches are the responsibility of either the provider or the business customer. Either party can increase device network buffers to alleviate packet drops.	A network provider can increase the sizes of the network buffer queues or use congestion management.

RFC 6349: TCP Throughput Test Methodology

TrueSpeed is the JDSU implementation of the new IETF RFC 6349, a practical methodology for measuring end-to-end TCP throughput in a managed IP network. The goal of RFC 6349 is to provide a better indication of the user experience by verifying TCP-layer performance. RFC 6349 also specifies TCP and IP parameters that optimize TCP throughput.

RFC 6349 recommends always conducting a Layer 2/3 turn-up test before TCP testing. RFC 6349 specifies these test steps to measure TCP throughput:

1. Path MTU detection (per RFC 4821) to verify the network maximum transmission unit (MTU) with active TCP-segment size testing to ensure that the TCP payload remains unfragmented.

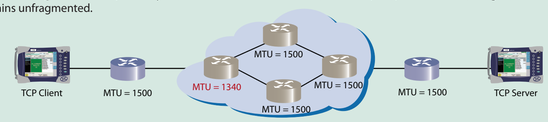

2. Baseline round-trip delay and bandwidth to predict the optimal TCP window size for automatically calculating the TCP BDP.

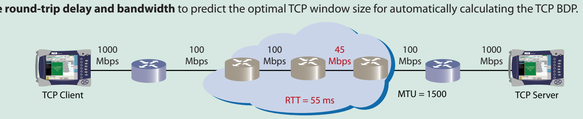

3. Single and multiple TCP-connection throughput tests to verify TCP window size predictions that enable automated "full-pipe" TCP testing.

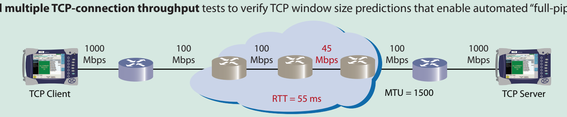

BDP = bottleneck bandwidth x RTT/8
= 45 Mbps x 55 ms/8
= 309,000 bytes

Using 5 TCP connections with 64 KB windows fills the BDP (5 x 64 KB); Layer 4 throughput should equal ~44 Mbps.

Each connection should achieve about 8.8 Mbps at Layer 4.

http://www.ietf.org/rfc/rfc6349.txt

RFC 6349 TCP Metrics

The **transfer time ratio** is the quotient of the actual TCP transfer time divided by the **ideal TCP transfer time**. Ideal TCP transfer time is derived from the network path bottleneck bandwidth and Layer 1, 2, 3, and 4 overheads.

Ideal TCP transfer time = 90 s
Actual TCP transfer time = 135 s
Transfer time ratio = 135/90 = 1.5

The **TCP Efficiency** metric is the percentage of bytes that did not have to be retransmitted and is defined as: $\frac{\text{transmitted bytes} - \text{retransmitted bytes}}{\text{transmitted bytes}} \times 100$

For example, if 100,000 bytes were sent and 1,000 had to be retransmitted, the TCP Efficiency is calculated as: $\frac{101{,}000 - 1{,}000}{101{,}000} \times 100 = 99\%$

Buffer Delay Percentage — TCP throughput is also affected by an increase in RTT, which can be caused by network congestion or buffer delay. The Buffer Delay Percentage is defined as: $\frac{\text{average RTT during transfer} - \text{baseline RTT}}{\text{baseline RTT}} \times 100$

For example, if the baseline RTT for a network path is 2 ms and the average RTT increases to 3 ms during the test, the percentage is calculated as: $\frac{3 - 2}{2} \times 100 = 50\%$

Ideal TCP Efficiency is 100%, indicating no retransmissions. Ideal Buffer Delay Percentage is 0%, indicating no loss in throughput due to congestive delay.

Turn-up Problems and Applicable Testing Standards

Problem	RFC 2544	Y.1564	RFC 6349
Single-service, Layer 2/3 SLA issues such as loss and jitter	✓	✓	N/A
Multiservice, Layer 2/3 SLA issues such as service prioritization, loss, and jitter	×	✓	N/A
Inadequate TCP window sizes (CPE issues)	×	×	✓
Inadequate device buffers (for bursty applications)	×	×	✓
Excessive retransmissions due to policing	×	×	✓

TrueSpeed
The industry's first RFC 6349-based TCP test

- **Reduce operating expenses** up to 30 percent by minimizing truck rolls
- **Complete fast**, repeatable, **automated TCP tests** in less than 5 minutes
- **Trust in the reliability** of a product developed by the lead author of RFC 6349
- **Verify results with an intuitive graphical user interface** that is easy to use by technicians at all skill levels

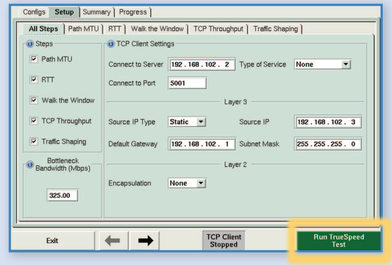

Using the JDSU T-BERD/MTS-6000A with the RFC 6349-compliant TrueSpeed option, TCP tests can be run in less than 5 minutes by novice technicians using a simple, push-button interface.

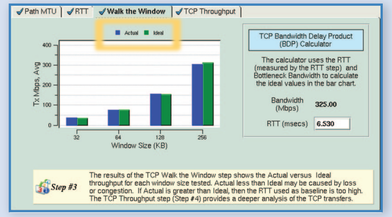

The Walk the Window step automatically tests four different TCP window sizes and highlights actual vs. ideal throughput

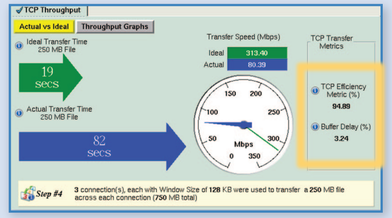

If actual TCP throughput does not equal the ideal, TrueSpeed conducts a more detailed TCP throughput test of a problematic window size. It also provides the following RFC 6349 metric results to assist in the diagnosis:

TCP Efficiency = loss metric

Buffer Delay Percentage = network buffer metric

TrueSpeed is currently available on these JDSU product platforms:

T-BERD/MTS-8000

T-BERD/MTS-6000A

To learn more, visit www.jdsu.com/go/rfc6349

2011 INTERNET TELEPHONY PRODUCT OF THE YEAR

JDSU

Note: Specifications, terms and conditions are subject to change without notice.
30173026 001 0112 TRUESPEED.PO.TFS.TM.AE