ExpertTCP™ - TCP Throughput Testing (RFC 6349)

(Test with multiple TCP streams up to 1 Gbps rate)

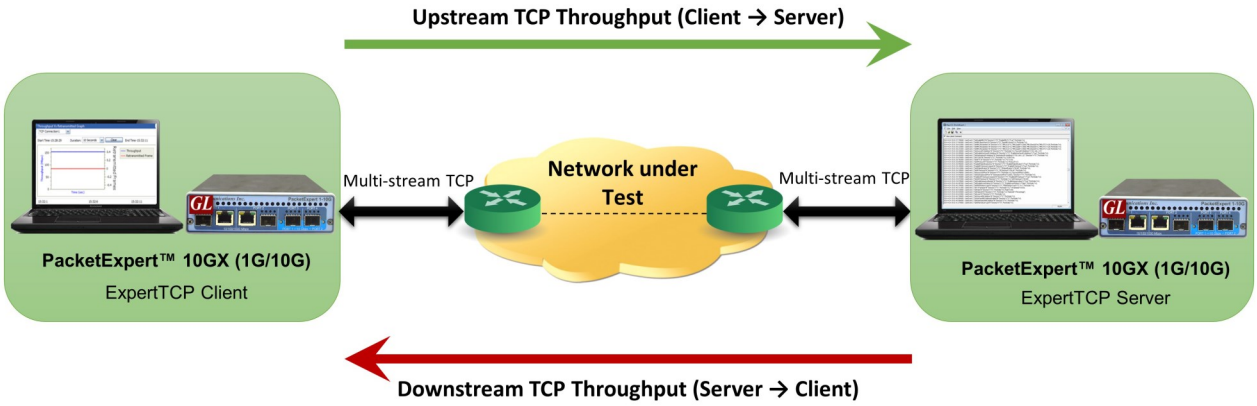

Overview

GL's **ExpertTCP™** is an optional application with PacketExpert™ 10GX platform. The application performs RFC 6349 based TCP Throughput testing on 10 Gbps, 2.5 Gbps and 1 Gbps networks. Although ExpertTCP™ is supported on 10 Gbps ports, the maximum TCP generation rate is limited to 1 Gbps due to constraints of the hardware-based FPGA TCP implementation. Many real-world networks are not symmetrical. There may be significant differences between upstream and downstream directions.

ExpertTCP™ testing is performed using the RFC 6349 standard. To conduct this test, users need **two PacketExpert™ devices** — one as the client and the other as the server. The ExpertTCP™ test covers both upload (Client to Server) and download (Server to Client), measuring TCP throughput and efficiency. The server at the remote location is completely controlled by the client side (located locally).

Supports unidirectional or simultaneous bi-directional testing, with results reported for each direction. The remote server is fully controlled by the local client, allowing users to configure both client and server from a single location. This eliminates the complexity of managing test setups across multiple sites. For bi-directional or unidirectional testing, two PacketExpert™ devices are required to conduct the test. One device acts as the TCP client and the other as TCP server.

PacketExpert™ 10GX (PXN100) has two 10/1 Gbps Optical/Electrical ports, and two 1000 Mbps Electrical/Optical ports. The 10 Gbps ports can be down-shifted to 1Gbps using Electrical to SFP transceivers, thus offering all 4 Electrical/Optical 1G ports for testing. The same two 10 Gbps ports can now be converted to 2.5 Gbps ports with appropriate SFP. It requires additional **PXN101** license installation to run the tests on 10G ports and 2.5G Ports. PacketExpert™ 10GX is available in portable as well as mTOP™ Rack-mount platforms. Each GigE port provides independent Ethernet/IP testing at wire speed for applications such as BERT, RFC 2544, SmartLoopback, Record Playback, PacketBroker, ExpertSAM™ and Multi-stream Traffic Generator and Analyzer and ExpertTCP™.

The current widely used standards for verifying Ethernet/IP based networks for SLA parameters are RFC 2544 or Y.1564. Both these standards are meant for testing at Layer2 or Layer3 (Ethernet or IP layers).

Though these tests are necessary, they are not sufficient, because they do not cover testing at TCP layer. Most web based applications like HTTP, FTP, E-mail etc. run over TCP. Many modern web applications such as Facebook, YouTube, and others use TCP.

Even if service provider networks are tested using RFC 2544 or Y.1564, customers may still face problems as the TCP throughput may not match the throughput at the Ethernet/IP layer. This is because TCP throughput depends on factors like the TCP Window Size, buffer size of intermediate network nodes etc. Also, impairments such as latency and packet drops causes TCP retransmissions, severely affecting the TCP throughput. So, there is a gap in the current testing methods and to cover the gap, RFC 6349 frame work has been devised.

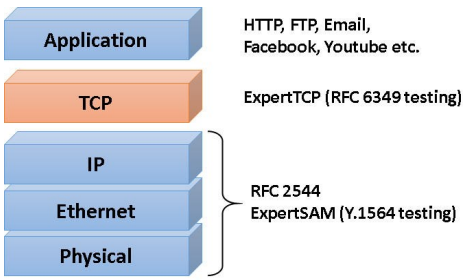

For more information, please visit ExpertTCP™ - TCP Throughput Testing (RFC 6349) webpage.

GL Communications Inc.

818 West Diamond Avenue - Third Floor, Gaithersburg, MD 20878, U.S.A

(Web) www.gl.com - (V) +1-301-670-4784 (F) +1-301-670-9187 - (E-Mail) info@gl.com

Features

- Supports Path MTU, Baseline RTT and TCP Throughput tests.
- Support for frame lengths from 64 bytes to Jumbo frames (up to 16000 bytes).
- Supports full duplex multiple TCP connections (up to 16 TCP connections on PXN100 and 8 TCP connections on PXE100).
- Test asymmetrical path with separate set of configurations for Upstream (Local to Remote), and Downstream (Remote to Local).
- Complete remote control - user needs to interact only with Local side - all results/statistics (both local and remote side) provided on local side.
- RFC 6349 specified metrics - TCP Efficiency, Buffer Delay Percentage, TCP Transfer Time ratio.
- With PXN101 licensing, the unit supports testing on 10G and 2.5G optical ports.

Functional Procedures

ExpertTCP™ (RFC 6349) specifies the TCP Throughput test to be conducted in 3 steps:

- Path MTU Discovery
- Determine Baseline RTT
- Conduct TCP Throughput test

ExpertTCP™ supports all 3 tests in a seamless one touch way. User can configure the parameters and simply start the test. ExpertTCP™ will run through all the three tests and report the results.

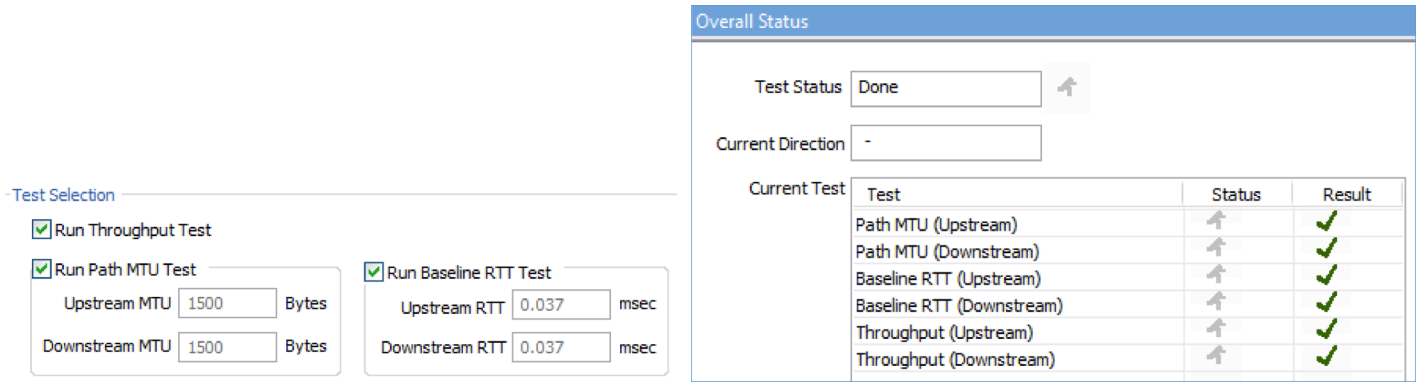

Figure: Test Selection and Status

Path MTU Discovery

As per RFC 6349, the first step is to discover the Path Maximum Transmission Unit (MTU). This is because the TCP Throughput test has to be conducted at the Path MTU, else the TCP segments can fragment, adversely affecting the test results.

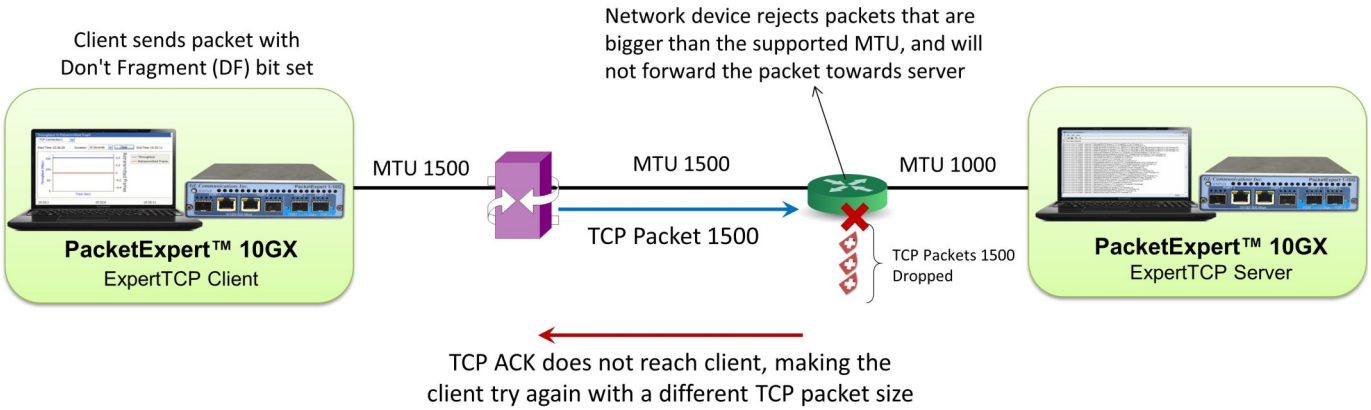

Path MTU Discovery

ExpertTCP™ discovers the Path MTU for both directions (Upstream/Downstream) separately. It follows a method similar to the one specified in the RFC 4821 (Packetization Layer Path MTU Discovery) standard, but uses TCP instead of ICMP.

Path MTU determines the Maximum Segment Size (MSS) that TCP can use during the test.

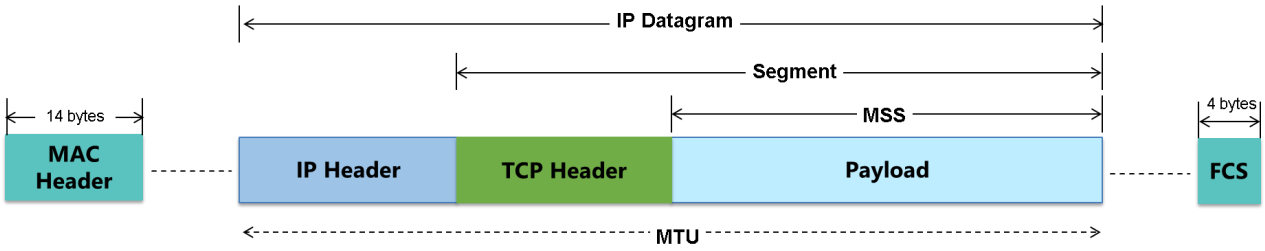

Determine Baseline RTT

This step establishes the inherent, non-congested Round-Trip Time (RTT) of the end-to-end network path. TCP RTT is the time taken for a TCP data packet to reach the other end, and for the corresponding ACK packet to reach the sender.

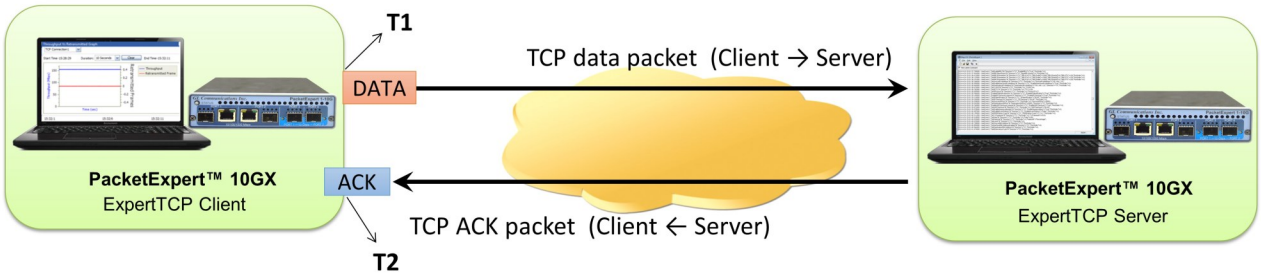

[RTT = T2 – T1]

RTT is an important metric in TCP, as the path RTT determines how much data can be sent out on the wire before an ACK can be received. This measurement is used to provide estimates of the TCP Receiver Window (TCP RWND) that should be used during subsequent Throughput test.

ExpertTCP™ performs Baseline RTT test separately for each direction and automatically calculates the optimum TCP RWND size based on the results. The Bandwidth Delay Product (BDP) and the RWND are displayed for user reference.

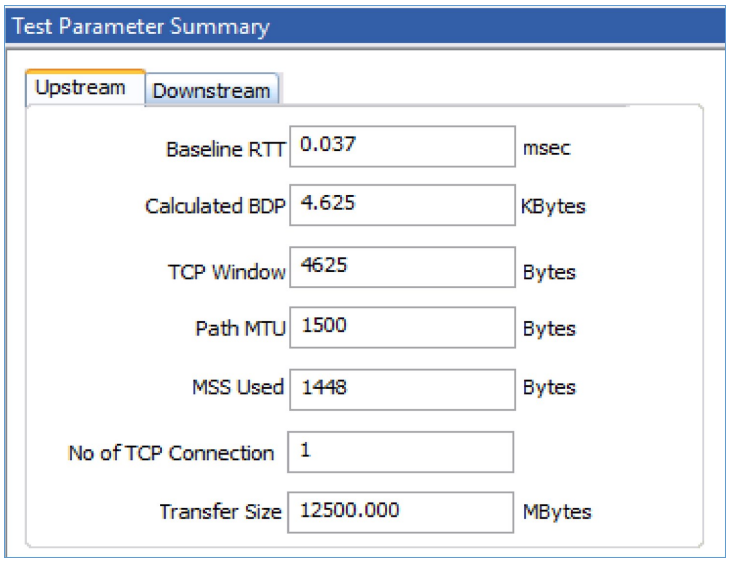

Upstream Downstream Test Parameter Summary

Conduct TCP Throughput Test

In this step, single/multiple TCP connection Throughput tests are conducted and the TCP Throughput is determined. The TCP RWND (Receiver Window) used during this step, is calculated from the Baseline RTT value measured during the previous Baseline RTT test. For multiple TCP connections, the calculated RWND is distributed among the connections.

Up to 8 TCP connections are supported on 1G interface (PXE100), whereas up to 16 TCP connections are supported on 10G interface (PXN100).

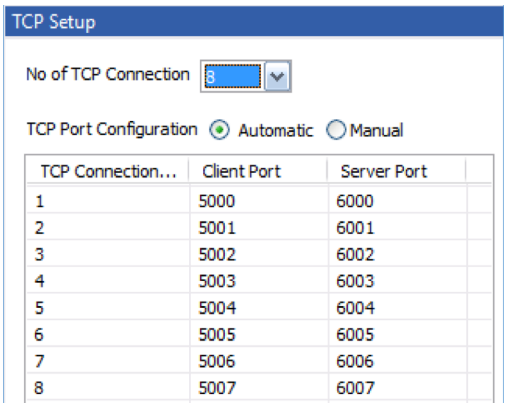

Statistics and Results

Various statistics at runtime such as minimum, maximum and average Throughput, and RTT measurements per connection provide detail insight into the performance.

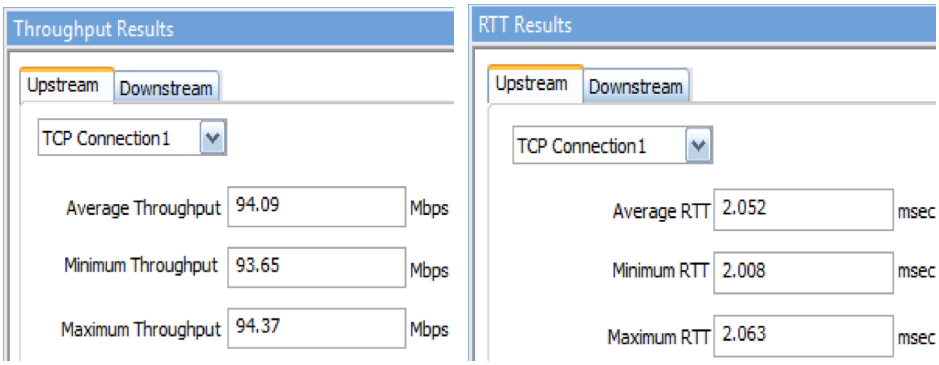

The statistics information on number of frames and bytes Transmitted, and Retransmitted per connection, gives a snapshot of the performance.

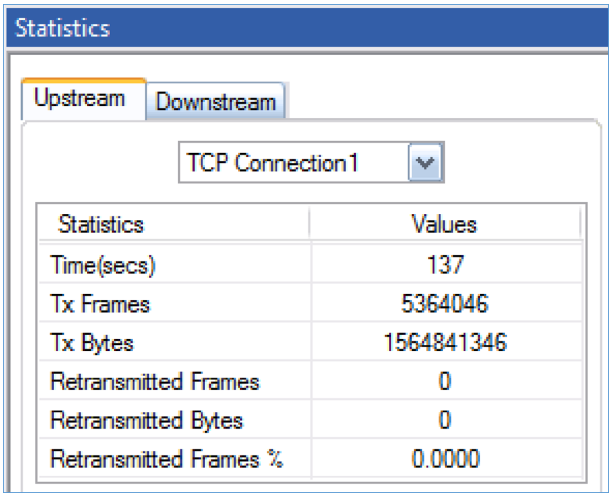

Statistics and Results

Final Results include the comparison of the actual Throughput/Transfer time with the ideal values. The three RFC 6349 TCP metrics defined in the specification - Transfer Time Ratio, TCP Efficiency and Buffer Delay are reported here.

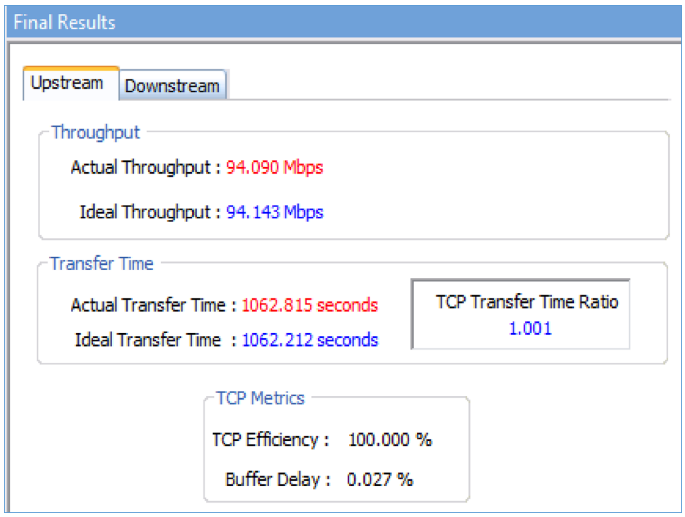

Upstream, Downstream Results

Graphs

Various graphs are available for visualizations. Throughput graph plots the Throughput over time. All connections are plotted in a single graph for a comprehensive view of the overall performance.

Throughput vs. Retransmissions graph provides insight into how Retransmissions are affecting the Throughput.

Throughput vs. RTT graph visualizes how the RTT variation affects the TCP Throughput.

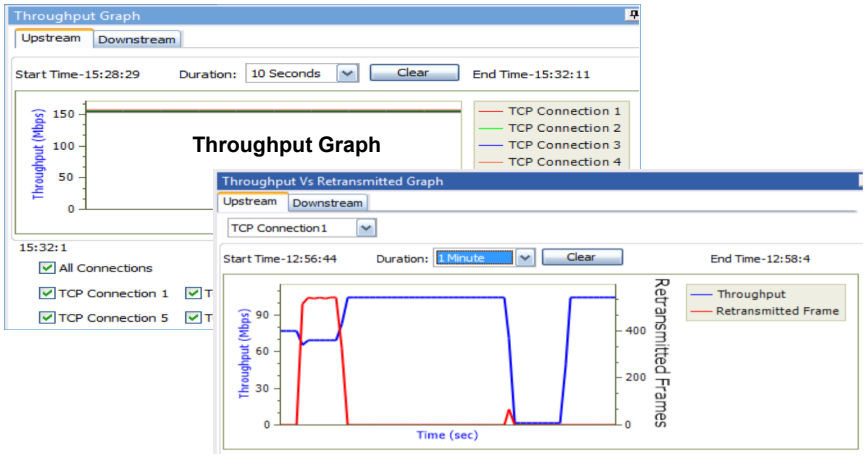

Throughput vs. Retransmissions Graph

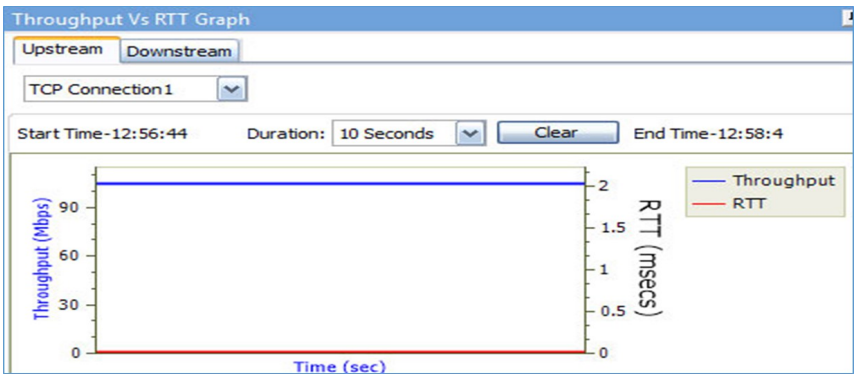

Throughput vs. RTT Graph

Hardware Specifications

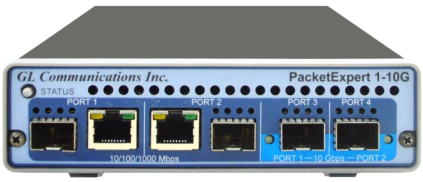

Portable 10GX Hardware Unit

1U mTOP™ PacketExpert™ 10GX Rack-mount(3 PXN100s)

Stacked 1U mTOP™ PacketExpert™ 10GX Rack-mount

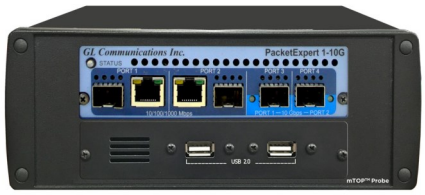

PacketExpert™ 10GX mTOP™ Probe

Physical Specification:
Length: 8.45 in. (214.63 mm)
Width: 5.55 in. (140.97 mm)
Height: 1.60 in (40.64 mm)
Weight: 1.713 lbs. (0.75 kg)

- **Dimension:** 1U/2U mTOP™ - 19" W x 16" L
- 1U mTOP™ Rack-mount Enclosure can support up to 3 PXN100s
- 2U mTOP™ Rack-mount Enclosure can support up to 6 PXN100s
- Optional 4 to 12 Port SMA Jack Trigger Board (TTL Input/Output)
- **Weight:** (not including the rails)
  1U with 3x PXN100 : 11 lbs
  2U with 6x PXN100 : 22 lbs

Physical Specification:
Length: 10.4 in. (264.16 mm)
Width: 8.4 in. (213.36 mm)
Height: 3.0 in. (76.2 mm)

Bus Interface: USB 3.0
External Power Supply:

- +12 Volts (Medical Grade), 3 Amps (For portable units having serial number ≥ 188400)
- +9 Volts, 2 Amps (For portable units having serial number < 188400)
- Optional 4-Port SMA Jack Trigger Board (TTL Input/Output)

SBC Specifications:

- Intel Core i3 or optional i7 Equivalent, Windows® 11 64-bit Pro OS
- USB 2.0 and USB 3.0 Hub, ATX Power Supply
- USB Type C ports, Ethernet 2.5GigE port
- 256 GB Hard drive, 8G Memory (Min)
- Two HDMI ports

SBC Specifications:

- Intel Core i3 or optional i7 NUC Equivalent, Windows® 11 64-bit Pro OS
- USB 2.0 and USB 3.0 Hub, Power Supply +12 Volts, 3 Amps
- USB Type C ports, Ethernet 2.5GigE port
- 256 GB Hard drive, 8G Memory (Min)
- Two HDMI ports

Temperature: Operating Temperature
0° C to +50° C (only up to operating altitude of 5000 feet, and for Optical SFPs only i.e. Non Electrical SFPs)
+5° to +40° C (for operating altitude up to 10,000 feet, and for both Electrical and Optical SFPs)
Non-Operating Temperature: -30° to +60° C

Humidity: Operating Humidity: 0% to 80% RH
Non-Operating Humidity: 0% to 95% RH

Altitude: Operating Altitude: up to 10,000 feet
Non-Operating Altitude: up to 50,000 feet

Interfaces:
4 x 1G Base-X Optical OR 10/100/1000 Base-T Electrical
2 x 100Mbps Base-FX Optical
2 x 2.5 Gbps Electrical/Optical Interface
2 x 10G Base-SR, -LR -ER Electrical/Optical Interface
Single Mode or Multi Mode Fiber SFP support with LC connector
Protocols:
IEEE 802.3ae LAN PHY compliance
RFC 2544 compliance

Pelican Carry Case

Buyer's Guide

Item No	Product Description
PXN108	Multi-Stream UDP/TCP Traffic Generator and Analyzer – for PXN100
PXN101	10G and 2.5G option for PXN100
CXN100	CLI Server for PXN100

Item No	Related Hardware
PXN100	PacketExpert™ 10GX
PXE100	PacketExpert™ 1G
PXN112G	PacketExpert™ 10GX (12-Port) - Rack-mount
PXN124G	PacketExpert™ 10GX (24-Port) - Rack-mount
MT001	mTOP™ 1U Rack-mount w/ SBC
MT002	mTOP™ 1U Rack-mount w/o SBC

Note: PCs which include GL hardware/software require Intel or AMD processors for compliance.

For more information, please visit ExpertTCP™ - TCP Throughput Testing (RFC 6349) webpage.

GL Communications Inc.
818 West Diamond Avenue - Third Floor, Gaithersburg, MD 20878, U.S.A
(Web) www.gl.com - (V) +1-301-670-4784 (F) +1-301-670-9187 - (E-Mail) info@gl.com