ATTO ATNETSTAT User Manual

macOS® Network Troubleshooting and Diagnostic Tool

ATTO FastFrame™ Network Interface Cards
ATTO ThunderLink® NS Devices

ATTO
The Power Behind the Storage

155 CrossPoint Parkway
Amherst, NY 14068

P. +1.716.691.1999
atto.com

1/2019 **PRMA-0475-000**

Contents

1 Overview 3

Atnetstat 3

Understanding the OSI Model and the Structured Layer Approach 3

2 OSI (Open Source Interconnection) 7 Layer Model 4

Troubleshooting the Lower Layers 5

Atnetstat Usage 5

Description of Atnetstat Statistics 6

NETSTAT 7

Using the Data Provided by Atnetstat and Netstat 8

Some Important Metrics to Consider 8

Packet Loss 8

Other Useful Commands 9

Appendix A Warranty 13

ATTO Technology, Inc. Limited Warranty 13

1 Overview

The following document provides a tutorial on how to use ATTO Atnetstat tool to tune for network performance and identify potential network issues when using the ATTO FastFrame™ network controller or the ATTO Thunderbolt to Ethernet controllers on macOS. It will also discuss how to utilize the logical framework of the OSI (Open Systems Interconnection) conceptual model to isolate problems to a particular functional layer.

Atnetstat

The ATTO Atnetstat tool is a command line utility (CLI) for macOS that was created to monitor Layer 2 Frame statistics and offer a view into Layer 1 to assist with troubleshooting specific issues and to offer insight into performance tuning opportunities for the ATTO FastFrame Network adapters and Thunderbolt to Ethernet adapters. Note that the Network Interface Card's Operating System device driver operates at Layer2. The lower layers of the stack are the foundation and are not natively visible through normal utilities within the macOS. If the foundation is experiencing issues, the whole stack will be affected. Without this tool problems at layer 4 cannot easily be confirmed, nor distinguished from problems originating in the Network Controller propagating up the stack to layer 4.

It is important to have a basic understanding of the networking model before discussing how, when and why to use Atnetstat.

Understanding the OSI Model and the Structured Layer Approach

The network uses a "stack" of layered protocols, one upon another. The OSI (Open Systems Interconnection) model is used to reduce the complexity as it breaks complex network interactions into simpler elements. The OSI Model is a way of thinking about how networks work. The model divides the network into a framework of 7 layers, or sets of related functions. Each layer communicates and supports the layer above/below it. Each layer is only responsible for the functions at that layer and then for passing the results on to the next layer. Layer 1 is foundational. Without proper Layer 1 performance, Layer 2 will not function properly. Without Layer 1 and 2, Layer 3 will not function, and so on. When troubleshooting, start at the bottom and work your way through the layers until you locate the problem layer.

Each Layer (1 through 7) of the OSI networking model encapsulates and addresses a different part of the needs of the communications, thereby reducing the complexity of the engineering solutions. This simplification enables the distilling of useful concepts and metaphors that may be easily and accurately applied to the task at hand.

Layer 1 - Physical Layer - Problems at this layer typically occur with cabling and media connector issues. Tools: Atnetstat

Layer 2 - Data Link Layer - Problems that can occur at this layer are MAC addressing errors, duplex errors, link, collisions, CRC frame errors. Tools: ifconfig & Atnetstat

Layer 3 - Network Layer – Routing and logical network address. Problems that can occur at this layer are network addressing issues and routing issues. Tools: Ping, traceroute, arp, netstat -s & netstat -r

Layer 4 - Transport Layer – End to end transmissions. Problems that can occur at this layer are Fragmentation, Flow control and congestion. Tools to identify potential issues: iperf, netstat –s

Layer 5 - Session Layer – Manages sessions and conversations.

Layer 6 - Presentation Layer – Used for data format, compression and encryption as well as graphics.

Layer 7 - Application Layer – Protocols at this layer include SMB, NFS, AFP, FTP and telnet. Tools: smbutil, nfsstat, dd

If a performance issue is caused by the network, it can often be found by focusing on the Layer 4 TCP layer. Be aware that lower layer issues like packet loss, will propagate up the stack and manifest themselves at layer 4. When testing performance, tests should be measured at Layer 4 using a tool like iperf (use the iperf defaults – the autotune features work well). Layer 7 application performance should also be measured (SMB, NFS, AFP, FTP etc.) using the appropriate application performance tool. If the performance at Layer 4 is good, you should have good but slightly reduced Layer 7 performance (due to protocol processing overhead). A large decrease at Layer 7 may give a hint at where the problem might reside. For instance, review the SMB options and look for SMB protocol issues

2 OSI (Open Source Interconnection) 7 Layer Model

LAYER	DATA	APPLICATION / EXAMPLE		CENTRAL DEVICE/PROTOCOLS			DOD4 MODEL
Application (7) Serves as the window for users and application processes to access the network services	**Data**	**End User Layer - Program that opens what was sent or creates what is to be sent** Resource Sharing - Remote file access - Remote printer access - Directory Services - Network management		**User Applications** SMTP		**GATEWAY**	**Process**
Presentation (6) Formats the data to be presented to the Application layer. It can be viewed as the "Translator" for the network	**Data**	**Syntax layer - encrypt & decrypt (if needed)** Character code translation - Data conversion - Data compression - Data encryption - Character Set Translation		JPEG/ASCII EBDIC/TIFF/GIF PICT			
Session (5) Allows session establishment between processes running on different stations	**Data**	**Synch & send to ports** (logical ports) Session establishment, maintenance and termination - Session support - perform security, name recognition, logging, etc.		**Logical Ports** RPC/SQL/NFS NetBIOS names			
Transport (4) Ensures that messages are delivered error-free, in sequence and with no losses or duplication	**Seg-ments**	**TCP** - Host to Host, Flow Control Message segmentation - Message acknowledgement - Message traffic control - Session multiplexing	**PACKET FILTERING**	TCP/UDP			**Host to Host**
Network (3) Controls the operations of the subnet, deciding which physical path the data takes	**Packets**	**Packets ("letter", contains IP address)** Routing - Subnet traffic control - Frame fragmentation - Logical-physical address mapping - Subnet usage accounting		**Routers** IP/IGMP/ICMP			**Internet**
Data Link (2) Provides error-free transfer of data frames from one node to another over the Physical layer	**Frames**	**Frames ("envelopes", contains MAC address) (NIC card -- Switch -- NIC card)** Establishes and terminates the logical link between nodes - Frame traffic control - Frame sequencing - Frame acknowledgment - Frame delimiting - Frame error checking - Media access control		**Switch** **Bridge** **WAP** OSPF/ARP	**Land Based Layers**		**Network**
Physical (1) Concerned with the transmission and reception of the unstructured raw bit stream over the physical medium	**Bits**	**Physical Structure** - Cables, hubs, etc. Data Encoding - Physical medium attachment - Transmission technique - Baseband or Broadband - Physical medium transmission Bits & Volts		**Hub**			

Troubleshooting the Lower Layers

ATTO's Atnetstat reports on the lower layers of the stack. Layer 1 (Physical) & Layer 2 (Data-Link). Monitoring with Atnetstat –s will display lower layer, frame level statistics. This distinct layering framework provides a structured approach to help in tuning and identifying problems.

Netstat is a standard command-line tool for checking your network configuration, network connections, routing tables, network protocol statistics and activity. It operates and reports protocol statistics, on Layer2 (ARP). Layer 3 (Network) and Layer 4 (Transport) of the stack.

Atnetstat Usage

Use ATTO's Atnetstat command line utility to display and reset Network statistics. The Atnetstat command works either globally (all ports) or on a selected channel. Each channel can be selected with the "–c" switch. In the case of two controllers, each with two ports, channels would be 1 through 4.

Statistics are maintained until a "–r" switch option resets them, or the host is rebooted. Upon execution of the Atnetstat command, the time frame of the collected statistics is reported followed by the actual statistics. Be sure to reset old statistics before beginning a fresh capture.

This utility can be found in the ATTO FastFrame™ driver package:

/Applications/ATTO/FastFrame (for the NIC)

/Applications/ATTO/ThunderLinkNC (for the Thunderbolt to 10GbE converter)

The FastFrame controller or Thunderbolt driver and the Atnetstat utility work as a matched pair. Atnetstat will only work with driver versions 2.X or 3.X and above. Previous drivers have no support for the feature.

Atnetstat Network Statistics Tool Options:

-c {channel} Select a specific controller channel for the operation, starts at 1, all channels are selected by default.

-h Display extended help

-l List the controllers in the system

-r Reset controller statistics

-s Display controller statistics

-v Display non-error messages

-z Suppress sub-counts for statistics that sum to zero

Note ***It may be useful to use a "While loop" to watch statistics over time. To watch the statistics once per second: # while true; do Atnetstat –s; sleep 1; echo;done***

SAMPLE OUTPUT:

"atnetstat –s" for a single channel

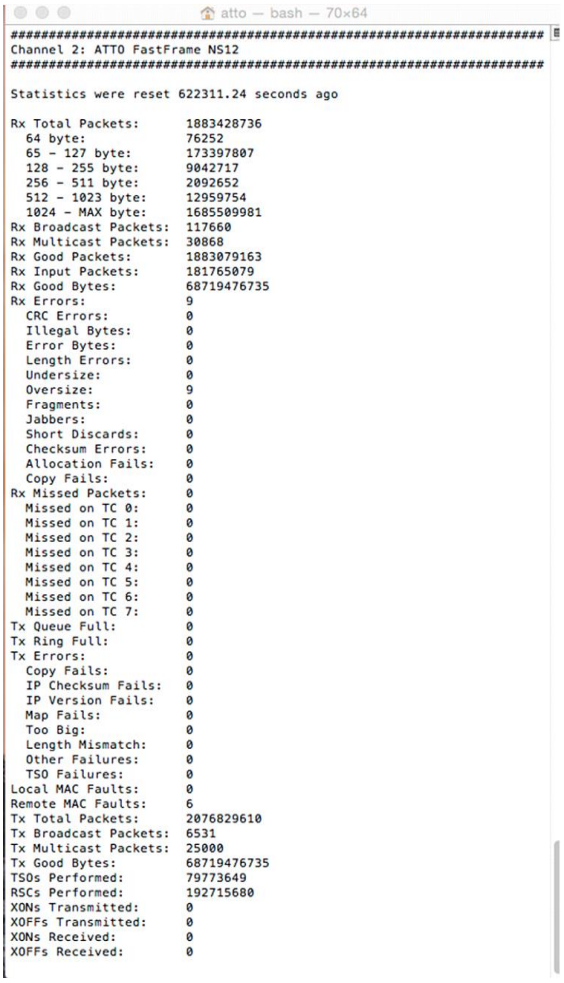

Description of Atnetstat Statistics

Rx Good Packets Number of good (non-erred) packets received that pass L2 filtering and have a legal length. Counts of good packets received are also displayed by packet size.

Rx Input Packets Number of good (non-erred) packets received that have been input to the network stack.

Rx Broadcast Packets Number of good (non-erred) broadcast packets received while the broadcast address filter is configured to allow reception of broadcast packets.

Rx Multicast Packets Number of good (non-erred) multicast packets received that pass L2 filtering, excluding broadcast packets and flow control packets.

Rx Total Packets Total number of all packets received (unicast, broadcast, multicast), regardless of length, errors, or L2 filtering, but excluding flow control packets.

Rx Good Bytes Total number of all bytes received in good (non-erred) packets from the <Destination Address> field through the <CRC> field, inclusively.

Rx Errors Total number of errors in packets received. When errors are displayed, check SFP, cable, MTU as well as local or remote interfaces.

CRC Errors - Number of packets received with CRC errors, not including packets whose length is less than 64 bytes (Fragments) or greater than the max packet size (Jabbers).

Illegal Bytes - Number of packets received with illegal byte errors, such as an illegal symbol in the packet.

Error Bytes - Number of packets received with error bytes, such as an error symbol in the packet.

Length Errors - Number of packets received whose packet length field in the MAC header doesn't match the actual packet length.

Undersize - Receive undersize errors: Received frames that are shorter than the minimum size (64 bytes) and have a valid CRC.

Oversize - Receive oversize errors: Received frames that are longer than the configured maximum packet size and have a valid CRC.

Fragments - Receive fragment errors: Received frames that are shorter than the minimum size (64 bytes) and have an invalid CRC.

Jabbers - Receive jabber errors: Received frames that are longer than the configured maximum packet size and have an invalid CRC.

Short Discards - Number of MAC short packet discard packets received.

Checksum Errors - Number of packets received that contain IPv4, TCP, UDP or SCTP checksum errors. Checksum errors are not counted when a packet has any MAC error (CRC, length, undersize, oversize, byte error or symbol error).

Allocation Fails - Number of packets that were dropped because of a memory allocation failure.

Copy Fails - Number of packets that were dropped because a memory copy operation unexpectedly failed.

Rx Missed Packets Number of packets received that were dropped because no buffer was available to receive the data. Check MBUF structures with netstat –m. Counts the total number of packets missed on all Traffic Classes (TC).

Local MAC Faults Count of faults detected in the local MAC. The occurrence of faults during link state transition is normal.

Remote MAC Faults Count of faults detected in the remote MAC. The occurrence of faults during link state transition is normal.

Tx Total Packets Total number of all packets transmitted, including standard, secure, FC, and manageability packets.

Tx Broadcast Packets Number of broadcast packets transmitted.

Tx Multicast Packets Number of multicast packets transmitted.

Tx Good Bytes Number of successfully transmitted bytes, including bytes from the <Destination Address> field through the <CRC> field, inclusively.

Tx Queue Full Number of times the transmit queue was full, resulting in a temporary transmit queue stall. Indication of possible dropped packets.

Tx Ring Full Number of times the transmit ring was full, resulting in a temporary transmit queue stall.

Tx Errors Total number of errors in packets transmitted - the sum of the following error counts:

Copy Fails - Number of packets that were dropped because a memory copy operation unexpectedly failed.

IP Checksum Fails - Number of packets that were dropped because of an error in the IP checksum.

IP Version Fails - Number of packets that were dropped because of an unexpected IP version.

Map Fails - Number of packets that were dropped because of an error mapping the packet memory.

Too Big - Number of packets that were dropped because they are too large for the configured MTU size.

Length Mismatch - Number of packets that were dropped because the packet length did not match the length indicated in the packet header.

Other Failures - Number of packets that were dropped due to a general failure.

TSO Failures - Number of packets that were dropped because a TSO was requested with invalid parameters.

TSOs Performed Number of Transmit Segmentation Offload operations attempted (including attempts that may have failed).

RSCs Performed Number of Received Side Coalescing operations attempted (including attempts that may have failed).

XON and XOFF Counts of Ethernet Pause Frames (Flow Control). Flow control is a Link layer attempt to relieve the pressure on queues to avoid congestion. When an Ethernet device gets congested or over loaded, flow control allows it to send PAUSE requests to the transmitter until the over loaded condition dissipates. If flow control is not enabled and an over loaded condition occurs, the device will drop packets. Dropping packets will impact performance.

NETSTAT

Netstat is a commonly available command-line tool for monitoring network protocol statistics and activity. Monitoring with netstat –s will display Protocol Statistics for IP Network Packets and TCP Segments organized and displayed as IP and TCP. The statistics are further sorted by send and receive.

EXAMPLE OUTPUT OF NETSTAT –S

Reduced (“SNIP”) number of fields for simplicity.

tcp:

157978 packets sent

9929 data packets (1926512 bytes)

40 data packets (6311 bytes) retransmitted

0 resends initiated by MTU discovery

115606 ack-only packets (83 delayed)

<SNIP>

375436 packets received

10818 acks (for 1926632 bytes)

285 duplicate acks

355817 packets (483863933 bytes) received in sequence

8916 out-of-order packets (12477569 bytes)

44 retransmit timeouts

83 correct ACK header predictions

349996 correct data packet header predictions

65 SACK recovery episodes

10 SACK options (SACK blocks) received

8557 SACK options (SACK blocks) sent

<SNIP>

ip:

478214 total packets received

0 bad header checksums

0 headers (0 bytes) checksummed in software

0 with size smaller than minimum

0 with data size < data length

1020 with data size > data length

0 with ip length > max ip packet size

0 with header length < data size

0 with data length < header length

0 with bad options

0 with incorrect version number

0 fragments received

0 dropped (dup or out of space)

0 dropped after timeout

0 reassembled ok

476975 packets for this host

<SNIP>

172247 packets sent from this host

0 output packets dropped due to no bufs, etc.

0 output packets discarded due to no route

0 output datagrams fragmented

0 fragments created

<SNIP>

Using the Data Provided by Atnetstat and Netstat

Some Important Metrics to Consider

Take time to observe and review the statistics. By observing all of the statistics one can get an understanding of what is occurring on the link. Identify where the Bottleneck exists by determining the problem layer, so you can concentrate on the real issue. Do not tune the Layer 4 sysctl variables when there is a Layer 2 issue. There are lots of knobs to use, the trick is finding them and learning how to use them. Use all the available tools to investigate issues and to gain an understanding of how the stack works.

Take note of the following:

Distribution of various packet sizes

Types of communications Broadcast, Multicast, Unicast

RSC and TSO offload efficiency

What type of

errors are being reported and how fast are they incrementing

Percentage of Good, compared to Total packets

Use a holistic approach including monitoring of switch and routers. External equipment will give additional data and may themselves, be the source of the problem. Do not ignore cabling switch/router port statistics etc. Once the actual issue is identified, tuning of the appropriate variables can be applied to attempt to remedy the situation.

Congestion is detrimental to the networks Quality of service. Increasing buffer sizes indiscriminately can result in “Buffer Bloat”, a phenomenon where excess buffering of packets causes high latency and packet delay variation (jitter). This results in a reduction in the overall network throughput.

Most default queues are FIFO (first in first out). When such a queue becomes full, all arriving traffic must be discarded. This is called “Tail Drop” This reduces the transmission rate. Multiple tail drop events can significantly reduce throughput and may lead to congestion collapse. Congestion is a very important subject but is beyond the scope of this manual, see the glossary for more on “Congestion”

Packet Loss

Look for evidence of packet loss. Netstat protocol statistics (retransmission and duplicate ack counters) are indications of loss. The Internet Standards treat packet loss and congestion as synonyms. The Layer 3 internet protocol (IP) is designed as a best-effort delivery service. Layer4 (TCP), provides guaranteed delivery for TCP “segments”, while UDP only error checks the “datagrams”. UDP is connectionless and does not have any concept of retransmissions. Packets may contain data corruption, arrive out of order, have duplicate arrivals or become lost (dropped/discarded). Both UDP and TCP are transport-layer protocols and provide multiplexing between processes on the same host implemented with port numbers. Routers discard incoming packets that can’t be stored or transmitted. Dropping of packets acts as an implicit signal that the network is congested, and may cause the transmitter to reduce the transmission rate resulting in lower performance. In the event of packet loss, the receiver notifies the sender of the loss. The sender automatically resends any segments that have not been acknowledged. This event is called a retransmission.

Three Duplicate ACKS triggers a fast retransmission (Layer 4 TCP function). In the case of Duplicate ACKs not returning from the Receiver, a retransmission timer (RTO - retransmission timeout) will expire on the transmitter, triggering a slow retransmission. Slow retransmissions (netstat –s “retransmission timeout”) are much more detrimental than Duplicate ACK fast retransmissions.

Since the Layer 4 Protocol UDP is connectionless and provides no recovery for packet loss, UDP Applications must define and implement their own mechanisms for handling packet loss. Drops can occur on ingress or egress. NICs will drop frames that are broken and not pass them up the stack to higher layers. Drops will result in Layer 4 issues like retransmissions.

On Ingress, broken frames will increment the appropriate counter in “RX Errors”

On Egress, drops are generally buffer exhaustion.

An incrementing “Duplicate Acks” count in netstat –s (tcp:) could indicate the receiver trying to notify the transmitter of lost packets. The receiver sends duplicate ACKs to notify the sender it is waiting on the next segment of data. The transmitter will then retransmit.

An incrementing “Retransmitted” packet and timeout count found in the netstat –s (tcp:) may indicate congestion. This is a count of the packets and bytes transmitted more than once.

An incrementing “Out of Order” count indicates possible congestion or alternate paths. These are detrimental to TCP processing as the stack must buffer and wait for the next segment before processing
the data.

Error counts should be considered as a percentage of the total packets. In other words 100 duplicates received would likely be insignificant for a million total received.

Use Ring/Queue adjustments when there is a Layer 7 slow application or a speed mismatch between hosts.

Netstat –an Useful for watching the send and receive queues

Netstat –w 1 Reports packets and bytes per 1 second. When streaming properly, we will observe consistent numbers between each 1 second sample.

Netstat –m Reports on the underlying memory structure, called MBUF. If the "requests for memory denied" value is nonzero, the mbuf and/or cluster pools may need to be adjusted

Other Useful Commands

The following commands may also be useful in examining network performance. To view the manual page for a command: “man ifconfig”

Ifconfig Used to display the current network configuration information

Arp (Address Resolution Protocol) is used for resolution of network layer addresses into link layer addresses. Use arp –an to view the resolution table.

Tcpdump Is a packet analyzer that runs from the command line. It allows the user to display TCP/IP and other packets being transmitted or received over a network.

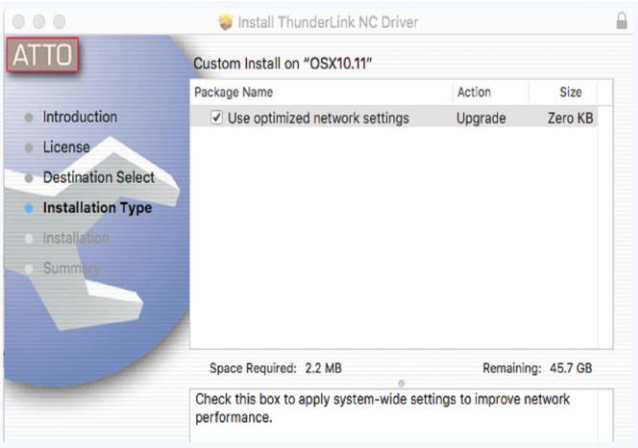

Iperf Is a network testing tool that can create Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) data streams and measure the throughput of a network that is carrying them. The default is for the stack to “autotune” TCP buffers. The autotune works well by figuring out the correct sizing of buffers and responds to changing network characteristics. The iperf tool allows you to override autotune with the –w and –l switches. Do not override autotune in iperf unless you fully understand the ramifications.

Sysctl Used to modify kernel parameters at runtime and can be applied to multiple layers of the stack.

sysctl -a complete list of all variables

sysctl -q queries a specific variable

sysctl -w writes a variable

sysctl net.inet.tcp lists all tcp variables

These settings are not persistent. Make your adjustments and test them. If you want to return to the defaults, simply reboot. When you are sure you have the ideal settings, override the defaults and preserve them across reboots by placing them in /etc/sysctl.conf

ATTO has spent considerable time analyzing optimal Sysctl settings for the ATTO 10GbE products and has created a

utility to easily change these settings for you. Select the “Use Optimized Network Settings” box when loading the ATTO NIC Operating System Device Driver to modify the default OSX sysctl settings. See figure below.

ACK	A flag or control bit that can be set indicating acknowledged data.
ARP	Address Resolution protocol. Mapping of MAC to IP addresses.
Autotune	TCP Auto Tuning enables TCP window scaling by default and automatically tunes the TCP receive window size for each individual connection based on the bandwidth delay product (BDP).
BDP	Bandwidth-Delay Product is the product of a data link's capacity (in bits per second) and its round-trip delay time (In seconds). Bits of data in transit between hosts = bottleneck link capacity (BW) * RTT**.** Throughput <= TCP buffer size / RTT. TCP window size >= BW * RTT.
Buffer Bloat	Phenomenon in packet-switched networks where excess buffering of packets causes high latency and packet delay variation (jitter) as well as reducing the overall network throughput.
Congestion Collapse	TCP settles into a stable state where traffic demand is high but little useful throughput is available. There are high levels of packet delay and loss (caused by discarded packets because output queues are too full).
Congestion	Network congestion occurs when a link or node is carrying so much data that its quality of service deteriorates. The loss of datagrams causes the TCP sender to enter slow-start, which reduces throughput in that TCP session until the sender begins to receive acknowledgements again and increases its congestion window. A more severe problem occurs when datagrams from multiple TCP connections are dropped, causing global synchronization; i.e. all of the involved TCP senders enter slow-start. This happens because, instead of discarding many segments from one connection, the router would tend to discard segments across all connections.
Congestion Control	A TCP congestion-avoidance algorithm, such as Reno, DCTCP or CUBIC (among others - each algorithm institutes different rules). Scheme to detect congestion, avoid and control congestion and to recover quickly from packet loss.
CWND	Congestion Window. The value of the CWND will be adjusted or increased with each acknowledgment received. Thus increasing the transmission rate.
Duplicate Ack	The receiver indicates lost data by acknowledging the same bytes of data. Three duplicate acks should trigger a retransmission (when using newer congestion control algorithms)
Fast Recovery	Response of some congestion avoidance algorithms such as Reno to resend lost data after 3 duplicate acks rather than wait for slow recovery as in Tahoe. Tahoe only uses a timeout for detecting congestion, while Reno uses both timeout and Fast-Retransmit.
Full-Duplex	A communication protocol which allows transmission in both directions at the same time. 10Gbe is a full duplex Link datagram.

FIFO	(First In First Out) Where the oldest (first) entry, or 'head' of the queue, is processed first. Packets leave the queue in the order in which they arrive. A full queue will drop during any attempts to put new data on the queue and this is called “Tail Drop” in a FIFO queue.
Flow Control	The process of managing the rate of data transmission between two nodes to prevent a fast sender from overwhelming a slow receiver. Measures are taken by the receiver to indicate to the sender the number of bytes it can receive beyond the last received TCP segment to avoid causing overrun and overflow in its internal buffers. This is sent in the ACK in the form of the highest sequence number it can receive without problems.
Frame	Layer 2 Data Link structure. Begins after the start frame delimiter with a frame header featuring source and destination MAC addresses.
IP	(Internet Protocol) Layer 3 Internetworking (data “packet”) that provides addressing and routing of packets across local area network boundaries. IP only provides best effort delivery and its service is characterized as unreliable. The upper layers are responsible for reliability.
Link	The physical and logical network component used to interconnect adjacent hosts in the network. A link protocol is a suite of methods and standards that operate only between those adjacent network nodes on a LAN (local area network) segment.
MAC	Media Access Controller.
MAC Address	A unique address assigned to an Ethernet device.
MSS	The largest amount of payload data (bytes) able to be sent in a single TCP packet. The value of MSS is negotiated between the endpoints during the 3 way handshake.
MTU	Maximum Transfer Unit is the largest possible frame size of a communications Protocol Data Unit (PDU) on an OSI Model Layer 2 data network.
OSI	Open Systems Interconnection.
Packet	Layer 3 structure. A packet consists of control information and user data, which is also known as the payload. Control information provides data for delivering the payload, for example: source and destination network addresses, error detection codes, and sequencing information. Typically, control information is found in packet headers and trailers.
PDU	Protocol Data Unit - the information delivered through a particular network layer. For each particular layer, a PDU is a complete message that implements the protocol at that layer.
PMTUD	Path MTU Discovery is a standard for determining the maximum transmission unit (MTU) size on the network path between two Internet Protocol (IP) hosts, with the goal of avoiding IP fragmentation.
PROTOCOL	An agreed-upon format for transmitting data between two devices.
Reliability	TCP assigns a sequence number to each byte transmitted, and expects a positive acknowledgment (ACK) from the receiving TCP. If the ACK is not received within a timeout interval, the data is retransmitted. The receiving TCP uses the sequence numbers to rearrange the segments when they arrive out of order, and to eliminate duplicate segments.

RETRANSMISSION	Resending of packets which have been either damaged or lost. Provides reliable transmission but impacts throughput negatively.
RTT	Round-Trip delay Time. It is used by TCP to adjust and manage connections. The ping command reports the RTT between nodes. RTT frequently changes over the duration of the session due to changing network conditions.
RTO	Retransmission Timeout is a timer based retransmission. Sometimes also called Slow Retransmission. A fast retransmission is triggered by 3 consecutive Duplicate Acks.
Segment	Layer 4 TCP structure used in end to end transport
Slow-start	Phase of congestion control strategy used by TCP to avoid sending more data than the network is capable of transmitting during the initiation of a connection or in response to congestion in some versions of congestion control algorithms.
Tail Drop	When a queue is filled to its maximum capacity, arriving packets are dropped until the queue drains enough to accept incoming traffic.
TCP	Transmission Control Protocol provides reliable, ordered, and error-checked delivery of a stream of octets between applications running at the Transport Layer on hosts communicating over an IP network. TCP works with “segments”.
UDP	User Datagram Protocol provides a connectionless datagram service at the Transport layer that emphasizes reduced latency over reliability. UDP works with structures called “datagrams”

Appendix A Warranty

ATTO Technology, Inc. Limited Warranty

ATTO Technology, Inc. (“ATTO”) warrants to the original purchaser of this product (“Product”) that the Product is free from defects in material and workmanship for the term described for this specific Product on ATTO's website (www.attotech.com). ATTO's liability shall be limited to replacing or repairing any defective product at ATTO's option. There is no charge for parts or labor if ATTO determines that this product is defective.

PRODUCTS WHICH HAVE BEEN SUBJECT TO ABUSE, MISUSE, ALTERATION, NEGLECT, OR THOSE PRODUCTS THAT HAVE BEEN SERVICED, REPAIRED OR INSTALLED BY UNAUTHORIZED PERSONNEL WILL NOT BE COVERED UNDER THIS WARRANTY. DAMAGE RESULTING FROM INCORRECT CONNECTION OR AN INAPPROPRIATE APPLICATION OF THIS PRODUCT SHALL NOT BE THE RESPONSIBILITY OF ATTO. LIABILITY UNDER THIS LIMITED WARRANTY IS LIMITED TO ATTO PRODUCT(S). DAMAGE TO OTHER EQUIPMENT CONNECTED TO ATTO PRODUCT(S) IS THE CUSTOMER'S RESPONSIBILITY. THIS LIMITED WARRANTY IS MADE IN LIEU OF ANY OTHER WARRANTIES, EXPRESS OR IMPLIED. ATTO DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT IMPLIED WARRANTIES CANNOT BE EXCLUDED, SUCH IMPLIED WARRANTIES ARE LIMITED IN DURATION TO THE EXPRESS WARRANTY PERIOD APPLICABLE TO THE PRODUCT. BECAUSE SOME STATES OR JURISDICTIONS DO NOT ALLOW LIMITATIONS ON THE DURATION OF IMPLIED WARRANTIES, THE ABOVE MAY NOT BE APPLICABLE. ATTO'S RESPONSIBILITY TO REPAIR OR REPLACE A DEFECTIVE PRODUCT IS THE SOLE AND EXCLUSIVE REMEDY PROVIDED TO THE CUSTOMER FOR BREACH OF THIS WARRANTY.

ATTO IS NOT RESPONSIBLE FOR DAMAGE TO OR LOSS OF ANY DATA, PROGRAMS OR ANY MEDIA. THE PRODUCTS ARE NOT INTENDED FOR USE IN: (I) MEDICAL DEVICES OR THE MEDICAL FIELD; OR (II) USE IN RUGGED APPLICATIONS.

ATTO IS NOT LIABLE FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, IRRESPECTIVE OF WHETHER ATTO HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. NO ATTO DEALER, AGENT OR EMPLOYEE IS AUTHORIZED TO MAKE ANY MODIFICATION, EXTENSION OR ADDITION TO THIS WARRANTY.

This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.